

STEREOGRAPHICS[®]

**DEVELOPERS'
HANDBOOK**

**Background on Creating Images for
CrystalEyes[®] and SimulEyes[®]**

StereoGraphics Developers' Handbook

Table of Contents

CHAPTER 1	1
Depth Cues 1	
Introduction	1
Monocular Cues	3
Summing Up	4
References	5
CHAPTER 2	7
Perceiving Stereoscopic Images	7
Retinal Disparity	7
Parallax	8
Parallax Classifications	9
Interaxial Separation	10
Congruent Nature of the Fields	10
Accommodation/Convergence Relationship	11
Control of Parallax	12
Crosstalk	12
Curing the Problems	13
Summary	13
References	13
CHAPTER 3	15
Composition	15
Accommodation/Convergence	16
Screen Surround	17
Orthostereoscopy	18
ZPS and HIT	20
Viewer-Space Effects	21
Viewer Distance	22
Software Tips	22
Proposed Standard	23
Summary	24
References	26

CHAPTER 4	27
Stereo-Vision Formats	27
What's a Format?	27
Interlace Stereo	28
Above-and-Below Format	28
Stereo-Ready Computers	30
Side-by-Side	31
White-Line-Code	31
Summing Up	33
CHAPTER 5	35
Parallelel Lens Axes	35
The Camera Model	36
The Parallax Factor	37
Windows and Cursors	39
References	40
CHAPTER 6	41
Creating Stereoscopic Software	41
CREATING STEREOSCOPIIC PAIRS	42
Perspective Projections	42
Stereo Perspective Projections	43
Parallax: Too Much of a Good Thing?	44
Putting the Eyes Where They Belong	45
Interacting with the User	49
THE ABOVE-BELOW FORMAT	49
Initialization and Interface Issues for Above-Below Stereo	51
THE SIMULEYES VR STEREOSCOPIIC FORMAT	52
SimulEyes VR White Line Code	52
Initialization and Interface Issues for SimulEyes VR Stereo	52
Pseudoscopic Images	53
The Page-Swap Alternative	54
GLOSSARY	55
INDEX	59


Depth Cues

Introduction


This handbook is about getting the most out of StereoGraphics products by acquiring a background in the fundamentals: how people see three-dimensionally and stereoscopically (they're not necessarily the same thing, as we shall learn), how electro-stereoscopic displays work, and how they interface with the visual system. While guidelines for creating good-looking stereoscopic images are given, and aspects of the subject may be codified, the creation of a stereoscopic image, like the creation of any image, is an art.

The purpose of this guide is to enhance the appearance and usefulness of stereoscopic images which may be displayed by a monitor or a projector. The images may be created with a graphics workstation or personal computer.

A stereoscopic image presents the left and right eyes of the viewer with different perspective viewpoints, just as the viewer sees the visual world. From these two slightly different views, the eye-brain synthesizes an image of the world with stereoscopic depth. You see a single — not double — image, since the two are fused by the mind into one which is greater than the sum of the two.



View-Master stereoscope.



Binocular perspectives. Each eye sees a slightly different perspective view, shown by the projection of the pyramid onto the plane surfaces.

A familiar example of a stereoscopic display is the View-Master[®] viewer illustrated here. If you close one eye when looking into the View-Master stereoscope, the image looks “flat.” With both eyes open, you see an image with stereoscopic depth. This kind of image used by the View-Master or by our **CrystalEyes** or **SimulEyes VR** products is called a *plano-stereoscopic* image, because it uses two planar (flat or two-dimensional) images.


A 3-D database may be displayed stereoscopically and many graphics packages or games may be upgraded to produce a true binocular stereoscopic image. In most cases, this is best left to the software developer who has access to the source code. Moreover, the developer is best able to support his or her product. Experience has shown that adding stereoscopic capability to an existing package is straightforward — a word not used lightly in a world where software development often requires heroic efforts. All computer-generated images produced from a three-dimensional database require the computation of an image from a single perspective point of view or location in space. A stereoscopic image differs only in that the image must be generated from two locations.

As an alternative, the image may be captured from the real world by means of a camera. In other words, it may be photographed or videographed. Whether the image is computed or photographed, the generation or capture of images from two perspectives — a left and a right — is required. If a single perspective viewpoint can be produced, it is conceptually straightforward to create the additional perspective. The creation of such images is the subject matter of this handbook.

People in the field of interactive computer graphics have also been playing with language by using the terms “3-D” or “three-dimensional” to mean a “realistic”-looking image which may be appreciated with one eye. Unfortunately, this is confusing nomenclature because most people who speak English outside of the field of computer graphics use the term “3-D,” and on occasion “three-dimensional,” when they mean a *stereoscopic* image. Such an image requires the beholder to use **both** eyes. The usual so-called 3-D computer image loses nothing when you close one eye, but a true stereoscopic image loses its *raison d’être* when viewed with one rather than two eyes.

Monocular Cues


The monocular, or extrastereoscopic, depth cues are the basis for the perception of depth in visual displays, and are just as important as stereopsis for creating images which are perceived as truly three-dimensional. These cues include light and shade, relative size, interposition, textural gradient, aerial perspective, motion parallax and, most importantly, perspective. A more complete description of depth perception may be found in a basic text on perceptual psychology.¹




Light and shade.




Relative size.



Interposition.



Textural gradient.



Aerial perspective.

Images which are rich in the monocular depth cues will be even easier to visualize when the binocular stereoscopic cue is added.

Light and shade provide a basic depth cue. Artists learn how to make objects look solid or rounded by shading them. Cast shadows can make an object appear to be resting on a surface.

Bright objects appear to be nearer than dim ones, and objects with bright colors look like they're closer than dark ones.

Relative size involves the size of the image of an object projected by the lens of the eye onto the retina. We know that objects appear larger when they are closer, and smaller when they are farther away. Memory helps us to make a judgment about the distance of familiar objects. A person seen at some great distance is interpreted to be far away rather than small.


Interposition is so obvious it's taken for granted. You perceive that the handbook you are now looking at is closer to you or in front of whatever is behind it, say your desk, because you can't see through the book. It is interposed between you and objects which are farther away.

Textural gradient is the only monocular depth cue articulated by a psychologist in modern times. The other cues were known and used by painters by the time of the Renaissance. A textured material, like a grassy lawn or the tweed of a jacket, provides a depth cue because the texture is more apparent as the object is closer to the observer.


Aerial perspective is the diminution in visibility of distant objects caused by intervening haze. Often, distant vistas will pick up a bluish haze because of the scattering of red light by the intervening atmosphere. In thick fog or haze, objects may not be all that distant.

Motion parallax is a depth cue which can be provided in the visual world or in moving-image displays, and can be achieved, for example, through the rotation of an object. A familiar example of this cue is seen from a moving car. Everyone has noticed that telephone poles move past more rapidly than the distant hills.

Not all graphic compositions can include motion parallax, but the use of motion parallax can produce a good depth effect, no matter what the spatial orientation of the axis of rotation. This axial independence indicates that motion parallax is a separate cue from stereopsis. If the two were related, only rotation about a vertical axis would provide a depth effect, but even rotation around a horizontal axis works. The point is made because some people think stereopsis and motion parallax are manifestations of the same entity. Motion parallax — as in the rotation of an object, for example — is a good cue to use in conjunction with stereopsis because it, like other monocular cues, stresses the stereoscopic cue. However, details or features of rotating objects are difficult to study.



Perspective.



Depth cuing.

Perspective, sometimes called “geometric,” “rectilinear,” or “photographic” perspective, is the most important extrastereoscopic depth cue for our purposes, since it scales the stereoscopic depth cue.² Computer-generated images with a strong perspective cue look deeper and may be easier to view, as we shall learn in later chapters. *Perspective* is the relationship between foreground and background objects. If it is exaggerated, or if there are perspective cues such as lines receding to a vanishing point, the image’s depth will be enhanced.

Depth Cuing is the name of a technique used in computer graphics. Depth cuing reduces the intensity of the object in proportion to the distance from the viewer.


Summing Up

Monocular depth cues are part of electronic displays, just as they are part of the visual world. While the usual electronic display doesn’t supply the stereoscopic depth cue, it can do a good job of producing a seemingly three-dimensional view of the world with monocular cues. These depth cues, especially perspective and motion parallax, can help to heighten the stereoscopic depth cue.


References

1. Kaufman, Lloyd. *Sight and mind: An introduction to visual perception*. New York: Oxford University Press, 1974.
2. MacAdam, D.L. Stereoscopic perceptions of size, shape, distance and direction. *SMPTE Journal*, 1954, 62:271-93.

Perceiving Stereoscopic Images



Wheatstone's stereoscope. Your nose is placed where mirrors L and R join, and your eyes see images at L1 and R1.




Single thumb. The eyes are converged on the thumb, and (with introspection) the background is seen as a double image.

Until relatively recently, mankind was not aware that there was a separable binocular depth sense. Through the ages, people like Euclid and Leonardo understood that we see different images of the world with each eye. But it was Wheatstone who in 1838 explained to the world, with his stereoscope and drawings, that there is a unique depth sense, *stereopsis*, produced by retinal disparity. Euclid, Kepler, and others wondered why we don't see a double image of the visual world. Wheatstone explained that the problem was actually the solution, by demonstrating that the mind fuses the two planar retinal images into one with stereopsis ("solid seeing").¹


A stereoscopic display is an optical system whose final component is the human mind. It functions by presenting the mind with the same kind of left and right views one sees in the visual world.

Retinal Disparity


Try this experiment: Hold your finger in front of your face. When you look at your finger, your eyes are *converging* on your finger. That is, the optical axes of both eyes cross on the finger. There are sets of muscles which move your eyes to accomplish this by placing the images of the finger on each fovea, or central portion, of each retina. If you continue to converge your eyes on your finger, paying attention to the background,



Double thumb. The eyes are converged on the background, and (with introspection) two thumbs are seen.



Retinal disparity. The image that would be seen on overlaid left and right retinae (provided the eyes were not converged on the goat).



Lenticular stereoscope. Invented by Sir David Brewster in 1844, and named to distinguish it from the mirror stereoscope. The eyes view the left and right images (stereopairs) through lenses which help them focus.

you'll notice the background appears to be double. On the other hand, when looking at the background your eyes are converging on *it*, and your finger, with introspection, will now appear to be double.


If we could take the images that are on your left and right retinae and somehow superimpose them as if they were Kodachrome slides, you'd see two almost-overlapping images — left and right perspective viewpoints — which have what physiologists call *disparity*. *Disparity* is the distance, in a horizontal direction, between the corresponding left and right image points of the superimposed retinal images. The corresponding points of the retinal images of an object on which the eyes are converged will have *zero disparity*.

Retinal disparity is caused by the fact that each of our eyes sees the world from a different point of view. The eyes are, on average for adults, two and a half inches or 64 millimeters apart. The disparity is processed by the eye-brain into a single image of the visual world. The mind's ability to combine two different, although similar, images into one image is called *fusion*, and the resultant sense of depth is called *stereopsis*.

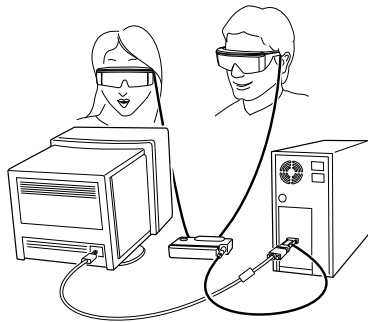
Parallax

A stereoscopic display is one that differs from a planar display in only one respect: It is able to display *parallax values* of the image points. *Parallax*, as we shall see, produces disparity in the eyes, thus providing the stereoscopic cue. The stereoscope is a means for presenting disparity information to the eyes. In the century and a half since Wheatstone, it has been improved and, as mentioned in the preface, the View-Master is one such device. In the View-Master stereoscope there are two separate optical systems. Each eye looks through a magnifying lens to see a slide. The slides are taken with a pair of cameras mounted on a base, to replicate the way we see the world with two eyes. Each camera sees a different perspective view of the world because it is offset horizontally from the other. When the two slides are viewed in the View-Master the result is the perception of an image with stereoscopic depth, resulting from the parallax information incorporated in the two photos. The parallax values of image points in the pair of slides may be measured just as we conceived of measuring retinal disparity — by laying the slides on top of each other and using a ruler between corresponding points.

Electro-stereoscopic displays provide parallax information to the eye by using a method related to that employed in the stereoscope. In the **CrystalEyes** or **SimulEyes VR** display, the left and right images are alternated rapidly on the monitor screen. When the viewer looks at the



CrystalEyes. An emitter, sitting on the monitor, broadcasts an infrared signal (dotted lines). The eyewear receives the signal and synchronizes its shutters to the video field rate.



SimulEyes VR

screen through shuttering eyewear, each shutter is synchronized to occlude the unwanted image and transmit the wanted image. Thus each eye sees only its appropriate perspective view. The left eye sees only the left view, and the right eye only the right view.


If the images (the term “fields” is often used for video and computer graphics) are *refreshed* (changed or written) fast enough (often at twice the rate of the planar display), the result is a flickerless stereoscopic image. This kind of a display is called a field-sequential stereoscopic display. It was this author’s pleasure to build, patent, and commercially exploit the first flickerless field-sequential electro-stereoscopic display, which has become the basis for virtually every stereoscopic computer graphics product.²

When you observe an electro-stereoscopic monitor image without our eyewear, it looks like there are two images overlaid and superimposed. The refresh rate is so high that you can’t see any flicker, and it looks like the images are double-exposed. The distance between left and right corresponding image points (sometimes also called “homologous” or “conjugate” points) is *parallax*, and may be measured in inches or millimeters.

Parallax and disparity are similar entities. Parallax is measured at the display screen, and disparity is measured at the retinae. When wearing our eyewear, parallax becomes retinal disparity. It is parallax which produces retinal disparity, and disparity in turn produces stereopsis. Parallax may also be given in terms of angular measure, which relates it to disparity by taking into account the viewer’s distance from the display screen.


Because parallax is the entity which produces the stereoscopic depth sensation, it’s important to study it. Below is a classification of the kinds of parallax one may encounter when viewing a display screen.

Parallax Classifications




Zero parallax.


Four basic types of parallax are shown in the drawings. In the first case, zero parallax, the homologous image points of the two images exactly correspond or lie on top of each other. When the eyes of the observer, spaced apart at distance t (the *interpupillary* or *interocular* distance, on average two and a half inches), are looking at the display screen and observing images with *zero parallax*, the eyes are converged at the plane of the screen. In other words, the optical axes of the eyes cross at the plane of the screen. When image points have zero parallax, they are said to have *zero parallax setting (ZPS)*.



Positive parallax.



Divergent parallax.



Negative parallax.

The drawing illustrates uncrossed or *positive parallax*. In one type of positive parallax, the axes of the left and right eyes are parallel. This happens in the visual world when looking at objects a great distance from the observer. For a stereoscopic display, when the distance between the eyes t equals the parallax t , the axes of the eyes will be parallel, just as they are when looking at a distant object in the visual world. We'll learn that having parallax values equal to t , or nearly t , for a small screen display will produce discomfort.


Any uncrossed or positive value of parallax between t and zero will produce images appearing to be within the space of the cathode ray tube (CRT), or behind the screen. We will say that such objects are *within CRT space*.

Another kind of positive parallax is shown here, in which images are separated by some distance greater than t . In this case, the axes of the eyes are *diverging*. This divergence does not occur when looking at objects in the visual world, and the unusual muscular effort needed to fuse such images may cause discomfort. There is no valid reason for divergence in computer-generated stereoscopic images.

In the drawing, we give the case in which the eyes' axes are crossed and parallax points are said to be crossed, or *negative*. Objects with negative parallax appear to be closer than the plane of the screen, or between the observer and the screen. We say that objects with negative parallax are *within viewer space*.

Interaxial Separation

The distance between the lenses used to take a stereoscopic photograph is called the *interaxial separation*, with reference to the lenses' axes. The axis of a lens is the line which may be drawn through the optical center of the lens, and is perpendicular to the plane of the imaging surface. Whether we're talking about an actual lens or we're using the construct of a lens to create a computer-generated image, the concept is the same. If the lenses are close together, the stereoscopic depth effect is reduced. In the limiting case, the two lenses (and axes) correspond to each other and a planar image results. As the lenses are located farther apart, parallax increases, and so does the strength of the stereoscopic cue.



Interaxial separation, t_c .

Congruent Nature of the Fields


The left and right image fields must be identical in every way except for the values of horizontal parallax. The color, geometry, and brightness of

the left and right image fields need to be the same or to within a very tight tolerance, or the result will be “eyestrain” for the viewer. If a system is producing image fields that are not congruent in these respects, there are problems with the hardware or software. You will never be able to produce good-quality stereoscopic images under such conditions. Left and right image fields congruent in all aspects except horizontal parallax are required to avoid discomfort.³


Accommodation/Convergence Relationship

The eyes' axes will be converged as if portions of a stereoscopic image are at different distances, but they remain focused (*accommodated*) on the plane of the screen. This is the only significant manner in which an electro-stereoscopic display differs from the way we see objects in the visual world. In the example given earlier (page 7), in which your eyes were asked to converge on your finger, they will also *focus* on the finger. When the eyes converge on the background, they will also *focus* on the background. This accommodation/convergence relationship is a habitual response learned by everyone: Invariably, in looking at objects, accommodation and convergence correspond. However, looking at an electro-stereoscopic image results in an exception to this relationship.

Because the action of the muscles controlling convergence and the muscles controlling focusing depart from their habitual relationship, some people may experience an unpleasant sensation when looking at stereoscopic images, especially images with large values of parallax. Thus, attention needs to be paid to parallax. Experience teaches that it's better to use the lowest values of parallax compatible with a good depth effect in order to minimize the breakdown of accommodation/convergence. Low parallax values will help to reduce viewer discomfort.



To clarify the concept, look at the sketch which gives the case of ZPS (zero parallax setting) for image points, in which the eyes are converged *and* accommodated on the plane of the display screen. This is the only case in which the breakdown of accommodation/convergence does not occur when looking at a projected plano-stereoscopic display. (An electronic display screen may be considered to produce a rear-projected image, in contrast to a stereoscope display which places the two views in separate optical systems.) Accordingly, low values of parallax will reduce the breakdown of accommodation/convergence, and large values of parallax will exacerbate it. This subject will be discussed in greater detail in the next chapter.



Maximum parallax. Left and right image points separated by half an inch, viewed from a distance of 18", shown here at almost half scale. If you're farther from the screen, more distance between the image points is acceptable.

Control of Parallax

The goal when creating stereoscopic images is to provide the deepest effect with the lowest values of parallax. This is accomplished in part by reducing the interaxial separation. In only one case, the ZPS condition, will there be no breakdown of the accommodation/convergence relationship. If the composition allows, it is best to place the principal object(s) at or near the plane of the screen, or to split the difference in parallax by fixing the ZPS on the middle of an object so that half of the parallax values will be positive and half negative.

As a rule, don't exceed parallax values of more than 1.5°. That's half an inch, or 12 millimeters, of positive or negative parallax for images to be viewed from the usual workstation distance of a foot and a half. This conservative rule was made to be broken, so let your eyes be your guide. Some images require less parallax to be comfortable to view, and you may find it's possible to greatly exceed this recommendation for negative parallax values or viewer-space effects.

Expressing parallax in terms of angular measure directly relates it to retinal disparity. For example, a third of an inch of parallax produces the same value of retinal disparity at three feet as two-thirds of an inch of parallax at six feet. Keep in mind the distance of viewers from the screen when composing a stereoscopic image.

Crosstalk

Crosstalk in a stereoscopic display results in each eye seeing an image of the unwanted perspective view. In a perfect stereoscopic system, each eye sees only its assigned image. That's what happens in a stereoscope, because each slide is optically separated from the other, as shown on page 8. Our technology, although good, is imperfect. In particular, there are two opportunities for crosstalk in an electronic stereoscopic display: departures from the ideal shutter in the eyewear, and CRT phosphor afterglow.⁵

If the electro-optical shutters were ideal, they would block all of the unwanted image. The shutters in our products are so good that no unwanted image will be perceived because of incomplete shutter occlusion. All the crosstalk is produced by another phenomenon — the afterglow of the display monitor's CRT phosphors.

In an ideal field-sequential stereoscopic display, the image of each field, made up of glowing phosphors, would vanish before the next field was

written, but that's not what happens. After the right image is written, it will persist while the left image is being written. Thus, an unwanted fading right image will persist into the left image (and vice versa).

The term *ghosting* is used to describe perceived crosstalk. Stereoscopists have also used the term "leakage" to describe this phenomenon. The perception of ghosting varies with the brightness of the image, color, and — most importantly — parallax and image contrast. Images with large values of parallax will have more ghosting than images with low parallax. High-contrast images, like black lines on a white background, will show the most ghosting. Given the present state of the art of monitors and their display tubes, the green phosphor has the longest afterglow and produces the most ghosting. If an image is creating ghosting problems, try reducing the green in the image.

Curing the Problems

The breakdown of the accommodation/convergence relationship and ghost images produced by crosstalk can detract from the enjoyment of a stereoscopic image. Fortunately, the cure for both problems is the same: Use the lowest values of parallax to produce a good stereoscopic effect.

As corresponding points approach the ZPS, both problems go away! As corresponding points move closer together, accommodation and convergence approach their usual relationship and the ghost image also diminishes, since it will increasingly correspond with the wanted image.

Summary

We have seen that a stereoscopic display differs from a planar display by incorporating *parallax*. Parallax is classified into *positive parallax*, which produces an image distant from the viewer (in CRT space), and *negative parallax*, which produces off-screen effects or images in viewer space. The perceptual artifacts introduced by the limitations of display tubes and the human visual system may be held to tolerable limits by properly adjusting the parallax, a subject which will be addressed further in the next chapter.

References

1. Wheatstone, Charles. On some remarkable, and hitherto unobserved, phenomena of binocular vision (Part the first). *Philosophical Transactions of the Royal Society of London*, 1838, 371-94.

2. Lipton, Lenny, et al. *Stereoscopic Television System*. U.S. Patent No. 4,523,226, June 11, 1985.
3. Lipton, Lenny. Binocular symmetries as criteria for the successful transmission of images. *Processing and Display of Three-Dimensional Data II*, SPIE Vol.507, 1984.
4. Valyus, N.A. *Stereoscopy*. London, New York: Focal Press, 1962.
5. Lipton, Lenny. Factors affecting ghosting in a time-multiplexed plano-stereoscopic CRT display system. *SPIE Proceedings*, Vol.761, 1987.

Composition

Electro-stereoscopic images have an enhanced quality, but does the addition of stereopsis signify that such an image is more “realistic” than the planar version of that image? The idea that stereoscopic images are more realistic than their planar counterparts is accepted by many people. One might say that this has become the commonsense point of view. Although the addition of stereopsis to a display adds another depth sense, it is arguable whether this addition makes the display more realistic.

I’m content with the idea that there is an enhancement, but that this is merely an addition of another depth cue to the display. Stereopsis adds information in a form that is both sensually pleasing and useful. If you want to say it’s more realistic that’s your choice, but technology hasn’t reached the point where anybody is fooled into mistaking the display for the world itself.


As we shall see, a stereoscopic image which has a one-to-one correspondence or is *isomorphic* with the visual world may be uncomfortable to look at, and of questionable value for scientist, engineer, technician, and artist. To understand why, and how to control the unique variables of stereoscopic composition, we must first understand them.

There are several ways in which a stereoscopic image may depart from being isomorphic with the visual world. We’ll discuss three of these conditions which involve, in this order, psycho-physics, the psychology of

depth perception, and geometrical or optical considerations. The three conditions are: breakdown of accommodation/convergence, screen surround conflicting cues, and orthostereoscopic conditions.

Accommodation/Convergence

As stated in the previous chapter, when looking at an object in the visual world our eyes focus (*accommodate* by changing the shape of the eyes' lenses) and *converge* (rotate toward or away from each other). Accommodation and convergence are linked by the necessities of the visual world. When you look at an object, your eyes focus on it so it may be seen clearly, and your eyes also converge on it so that the object may be seen singly. The latter process is called *fusion*. Different sets of muscles control these functions, and the neurological pathways for these two processes are also separate. Nonetheless, we have grown accustomed or habituated to the linked response because of a lifetime of visual experience.



Zero parallax. In this case there's no breakdown of accommodation/convergence when viewing a plano-stereoscopic display.

When looking at a stereoscopic display, the eyes accommodate on the plane of the screen but converge based on screen parallax. Thus there is a breakdown of the habitually tied-together responses of two separate mechanisms. For some, but not all, people, this is perceived as discomfort. For the first time in their visual experience, the two sets of muscles are used separately and the unexpected effort may be unpleasant. Many individuals, especially children, don't have any problem with the breakdown of accommodation and convergence. In addition, with practice (which involves no conscious effort), many people become accustomed to viewing stereoscopic images.

The problem is exacerbated for the small screens (only a foot or two across) viewed at close distances (a foot and a half to five feet) that are characteristic of electro-stereoscopic workstation displays. Large-screen displays, viewed at greater distances, may be perceived with less effort. There is evidence to indicate that the breakdown of accommodation and convergence is less severe in this case.¹ When the eyes are accommodated for distances of many yards the departures from expected convergence may not be troublesome. This conforms to my experience, and I've been amazed at my agreeable response to stereoscopic motion pictures, typically at theme parks, employing gigantic positive parallax.

The breakdown of the relationship between accommodation and convergence is one important way in which a plano-stereoscopic display cannot be isomorphic with the visual world. This psycho-physical fact is one that users, software designers, and stereovideographers must keep in

mind, because a plano-stereoscopic display can never overcome this artifact. A direct consequence of this is that, as a rule of thumb, one should keep parallax to the minimum required for a satisfactory depth effect.

Accommodation and convergence don't provide the mind with depth information.² This statement is at odds with the belief of many stereoscopists. In the case of convergence, the evidence that it provides any depth information has not been clearly demonstrated. For accommodation, it may be the perception of blur (rather than muscles changing the shape of the eyes' lenses) that provides depth information. Therefore, the breakdown of accommodation/convergence probably doesn't result in a conflict of depth cues. While it may seem like a fine distinction, the conflict is actually based on a departure from the habituated relationship of two sets of neural pathways and the muscles they control.



The accommodation/convergence relationship is habitual and learned; it is not a physiological given. People who are experienced in looking at stereoscopic images may no longer have the usual accommodation/convergence response. This can be a disadvantage, because stereoscopic images so composed may exceed the limits of what is comfortable for untrained people. Bear this in mind and avoid composing images which tax other people's visual systems. It's a good idea to try your images out on people who are not adept.

Screen Surround

Looking at a television or monitor screen is like looking through a window, and the plane of the display screen is called the *stereo window*. The vertical and horizontal edges of the screen are called the *surround*. As discussed earlier, a stereo image can be thought of as existing in two regions: one within the cathode ray tube (CRT space) and one between it and the viewer (viewer space). When composing a stereoscopic image, there's a basic decision to be made: where to position the scene or the object with respect to the screen surface or stereo window (the boundary layer separating CRT space from viewer space).

It's best if the surround doesn't cut off a portion of an object with negative parallax. (By convention, as explained, an object with negative parallax exists in viewer space.) If the surround cuts off or touches an object in viewer space, there will be a conflict of depth cues which many people find objectionable. Because this anomaly never occurs in the visual world, we lack the vocabulary to describe it, so people typically

call the effect “blurry,” or “out of focus.” For many people, this conflict of cues will result in the image being pulled back into CRT space, despite the fact that it has negative parallax.³



Interposition and parallax cues. No conflict of cues: the top drawing shows the globe, in viewer space with negative parallax, centered within the screen surround. Conflict of cues: the lower drawing shows the globe cut off by the surround. In this case the interposition cue tells the eye-brain the globe is behind the surround, but the stereopsis cue tells the eye-brain the globe is in front of the surround.

Here's what's happening: One cue, the cue of interposition, tells you the object must be behind the window since it's cut off by the surround. The other cue, the stereoscopic cue provided by parallax, tells you that the object is in front of the window. Our experience tells us that when we are looking through a window, objects cannot be between the window and us. You can accept the stereo illusion only if the object is not occluded by a portion of the window.

Objects touching the horizontal portions of the surround, top and bottom, are less troublesome than objects touching the vertical portions of the surround. Objects can be hung from or cut off by the horizontal portions of the surround and still look all right.

In this regard, the vertical portion of the surround leads to problems. That's because of the paradoxical stereoscopic window effect. When looking outdoors through a real window — for example, out the left edge of the window — the right eye sees more of the outdoor image at the edge than the left eye sees. The same kind of experience occurs for the right edge of the window. The difficulty arises in stereoscopic imaging from the following fact: The right eye, when looking at the left edge of the screen for objects with negative parallax, will see less image, and the left eye will see more. This is at odds with what we see when we look through a real window. This departure from our usual experience may be the cause of the exacerbated difficulty associated with the perception of images at the vertical surrounds.

There's an exception to this which is especially important for games developers. Objects which are rapidly racing off the screen may have very large negative parallax values. They tend to look just fine. It's when such objects linger at the screen surround that many people have perceptual problems.

Orthostereoscopy


There are three conditions to be fulfilled for a stereoscopic image to be *orthostereoscopic*.⁴ Such an image would be isomorphic — in terms of the stereoscopic depth cue and perspective considerations — with the visual world. First, images of very distant objects must cause the viewer's eyes to have parallel lens axes. Second, the distance between

the perspective viewpoints used to compute or capture the two views must be equal to the viewer's interocular distance. And third, the perspective constraint must be fulfilled. That is, the image must be viewed from a distance (V) that is equal to the product of image magnification (M) of the object and the focal length (f) of the lens. This is written as:

$$V = Mf$$

This last constraint is well known to lens designers and photographers⁵ and holds true for all photographically-created images or images which are computer-generated using a camera model, as described in the next two chapters.

If we attempted to fulfill the first condition for distant image points, their parallax values would be equal to the interpupillary distance, and that's much too large for comfortable viewing. It's true that viewing distant objects in the visual world will produce parallel axes for the eyes' lenses, but the resultant breakdown of accommodation and convergence for a workstation screen will be uncomfortable.



$V = Mf$. When this condition is fulfilled, the angle α , subtended by the projected image of the object at the observer's eye, equals that of the object subtended by the camera's lens on the film.

If we tried to fulfill the second condition by specifying the distance between perspective views to be equal to the distance between the eyes, for many applications (such as molecular modeling or aerial mapping) the result would be useless. Molecules are very small, requiring a perspective separation computed on the order of the size of atoms, and aerial mapping often requires a stereo base of hundreds or thousands of feet.

Moreover, the distance between human eyes, or the interocular, varies between 50 and 75 millimeters. A view computed for one set of eyes wouldn't produce an orthostereoscopic result for most other eyes.



If we were to take the third constraint (the perspective constraint) seriously, we'd have to change our distance to the screen for every new object that appeared. Nobody is going to do this. People have learned how to make a mental adjustment for images produced with wide-angle or telephoto lenses.

Although people have asked about the orthostereoscopic condition over the years, I can't recall anybody needing to use it — with the exception of specialized photogrammetric and virtual-reality applications. It's interesting to keep in mind if only to see how one is departing from it. Since you now know the ortho constraints, try an experiment. Create an


image conforming to these constraints. How does it look? Now change the compositional parameters in accordance with the advice given here.

If the image looks good, that's good enough.

ZPS and HIT



Convergence. The eyes rotate to fuse images of an object at different distances.



HIT. Horizontal image translation of the left and right fields will control parallax.

The parallel lens axes algorithm will be described in Chapter 5. In this software approach, the two views are considered to be created by cameras whose lens axes are parallel. The zero parallax setting (ZPS) is produced through *horizontal image translation* (HIT) of the two perspective views. When the image fields are shifted so that the object, or a portion of the object, is at the ZPS, that portion of the object or scene will appear to be in the plane of the screen. Composing the image this way helps reduce the breakdown of accommodation and convergence.

Historically the term “convergence” has been applied to the process used to achieve the zero parallax condition for stereo imaging. The use of the term in this context creates confusion with regard to the physiological process of convergence in the visual world — a process requiring rotation of the eyes. There is ambiguity with regard to which process is meant — *physiological* or *stereoscopic* convergence.

A further complication arises from this usage: While a portion of the image is said to be “converged” when the corresponding points coincide, the term “plane of convergence” has been used for the *zero parallax* condition which supposedly exists for all image points located on a plane at a given distance (the distance on which the camera is converged). Unfortunately, the locus of such points will not be a plane if rotation is used for convergence.⁴ Rather, it will be a complex but continuous surface. (However, a distortion-free plane of zero parallax will exist when using the parallel lens axes algorithm described and explored in the next two chapters.)


Therefore, a break with prior usage is called for. The eyes converge on an object in the visual world, or on an image point in a stereoscopic display. Image points in a stereoscopic display appear in the plane of the screen when at the ZPS which, for a computer-generated image using the algorithm described here, is achieved through *HIT*.

We have the ability to control the location of the image in CRT or viewer space, or at the plane of the screen, through setting the HIT. For example, one can put the entire scene or object in front of the window, but it will probably be difficult to look at because of the breakdown of

accommodation and convergence, or because of the conflict of cues at the screen surround. Much of the time you're better off making the image appear to be seen through a window. This is especially true if you're creating a scenic view or an interior, as we will discuss in a moment.

On the other hand, many computer-generated images involve a single object — for example a molecule or a machine part. For this class of compositions try placing the center of the object in the plane of the screen with some of it in viewer space and some of it in CRT space. This tends to reduce the breakdown of accommodation and convergence. Where you place the object is also an aesthetic decision.

For compositions resembling scenes in the visual world, most of the stereoscopic image should lie within CRT space. Imagine the following southwestern scene: A cactus is in the foreground, a horse is in the middle distance, and at a great distance we can see mesas. The cactus touches the bottom edge of the screen surround. Let's give it zero parallax where it touches. To give it negative parallax would lead to a conflict of cues, because the stereopsis cue would indicate that the image was in front of the screen while the interposition cue would indicate it was behind. (In this particular example a little negative parallax for the cactus, placing it in viewer space, won't hurt because the conflict of cues for the horizontal edge of the surround is relatively benign.) The most distant objects in the scene, the mesas, will have the maximum positive parallax. Everything else in the scene has between zero parallax and the maximum value of positive parallax.




Parallax values. The cactus is at the plane of the screen; the horse and the mesas have positive parallax.

Viewer-Space Effects

Images with negative parallax values should be carefully treated, and a number of empirical rules apply. Such rules can be broken, but it may save time to learn one approach before trying everything under the sun. One of the most pleasant things about designing stereoscopic computer-generated images is that you can observe what you are doing as you are doing it. If there are troubles with your composition, you can correct them and observe how variations lead to heightened or diminished stereoscopic effects, and how changes may affect visual comfort. Most people enjoy off-screen effects, so it's a technique worth mastering. Some people, vulgarians or the young-at-heart, don't consider an image to be truly three-dimensional unless it pokes them in the nose.

An important variable to manipulate for strong off-screen effects is not parallax, but rather the perspective cue. It is the signature of the amateur



An image elongates with distance.

to use huge parallax values to obtain strong viewer-space effects. Stressing perspective is a better approach, and it's accomplished by using a wide-angle view, at a bare minimum 40 degrees horizontal, and by scaling the object to fill the screen. This will be discussed below and more fully in the following two chapters.

Viewer Distance

The farther the observer is from the screen, the greater is the allowable parallax, since the important entity is not the distance between homologous points but the convergence angle of the eyes required to fuse the points, as discussed on page 12. A given value of screen parallax viewed from a greater distance produces a lower value of retinal disparity.

A stereoscopic image looks more three-dimensional or deeper the farther the observer is from the display screen. Distant objects look farther away, and objects in viewer space look closer to the observer. This subject is treated in more detail elsewhere^{3,4} but, simply put, the proportional distances of parts of an image remain constant for any viewer distance from the screen. An object that looks like it's one foot from you if you're two feet from the screen will look like it's two feet from you when you move to a position four feet from the screen. Therefore, off-screen effects will be exaggerated for more distant viewers.

Images of objects may, for some compositions, have maximum negative parallax larger than that acceptable for positive parallax. Objects can have large values of negative parallax and be cut off by the vertical surround if they're held on the screen for a short period of time. As mentioned earlier, objects which are racing swiftly off or onto the screen can pass through the vertical surround without being difficult to view. Only experiments will tell you what works best. Remember that in this context the vertical edges of the surround produce more of a perceptual conflict than the horizontal edges.

Software Tips

The perspective cue weights or scales the stereoscopic cue. If there is a strong perspective cue produced by a wide-angle perspective viewpoint, there will be a strong stereoscopic effect. In order to maintain the image size, the object must be brought closer to the camera when using a wide-angle view. By this means perspective will be stressed. Given a strong

stereoscopic effect, the values of parallax may be lowered by reducing the interaxial distance. The stereopsis depth cue may be reduced in this case because the extrastereoscopic depth cue of perspective has been strengthened. The reduction in parallax will diminish the breakdown of convergence and accommodation and ghosting, thereby producing an image more pleasing to view.


Here's more advice for software developers who are dealing with images of single objects centered in the middle of the screen: When an image initially appears on the screen, let's allow image points at the center of the object to default to the ZPS condition. Moreover, let's maintain the ZPS within the object even during zooming or changing object size. Some people don't like this notion. They think that the ZPS ought to be in front of the object when it is far away (small) and behind the object when it is close (big), to simulate what they suppose happens in the visual world. But in the visual world the eyes remain converged on the object as it changes distance; hence, retinal disparity, the counterpart of screen parallax, remains constant.

The depth cue of relative size is sufficient to make the object appear to be near or far, or changing distance, and often we're only trying to scale the image to make it a convenient size. Its center ought to remain at the plane of the screen for easiest viewing.

One thing that's fun about stereoscopic displays has nothing to do with the depth effect. Rather it has to do with glitter, sparkle, and luster. These effects are related to the fact that certain objects refract or reflect light differently to each eye. For example, looking at a diamond, you may see highlight with one eye but not the other. Such is the nature of glitter. This is a fun effect that can be added to a stereoscopic image. It also is an effect that can be used for fire to make it look alive. For example, if you use bright orange for one view and red for the other, you will create the effect of scintillation.

Proposed Standard

Here is a proposed standard for stereoscopic software to help users control the two basic variables of stereoscopic composition: interaxial distance (t_c) and HIT. The arrow keys **N**, **S**, **E**, and **W** should be used to control these functions. **N** and **S** are devoted to the interaxial setting — **N** for increasing the distance between camera axes, and **S** for decreasing the distance.



Arrow keys to control parameters. Changes in distance between centers of perspective should be controlled by N and S (8 and 2) keys. HIT should be controlled by E and W (4 and 6) keys.

For controlling HIT there are two possible ways to change the relative direction of the image fields, and it doesn't matter which is assigned to the **E** and **W** keys. That's because the image fields approach each other until ZPS is reached and then pull away from each other after passing through ZPS. Thus, which way the fields will move relative to each other is dependent upon where the fields are with respect to the ZPS. This parameter is best controlled interactively without using eyewear because the relative position of the two fields is most easily observed by looking at the "superimposed" left and right images.

A thoughtful touch is to maintain the ZPS when the interaxial is varied. In this way, the user can concentrate on the strength of the stereoscopic effect without having to be involved in monitoring the ZPS and working the HIT keys.

Summary

Stereoscopic composition for electronic displays has been discussed by taking into consideration facts about the perception of depth. We have seen that physiology and psychology play an important part in guiding us toward aesthetically-pleasing solutions to the challenges of electro-stereoscopic composition. This in turn controls the design of software. I've stated an opinion that the addition of stereopsis to a display does not necessarily make it more realistic. The addition of stereopsis to a visual display is, in my view, the addition of useful and potentially pleasurable information. Whether or not such a display is more realistic than its planar counterpart is for you to decide.

Here are some points to remember, based on what's appeared (and will appear) in this and other chapters in this handbook:

Points to Remember

- *The concept of orthostereoscopy, in which the image is congruent or isomorphic with the visual world, is rarely of practical value.*
- *Create images with the lowest parallax to give the desired effect.*
- *Stress the perspective cue, by using a wide-angle point of view, to allow lower parallax.*

- *Use ZPS for the center of an image of a single object to reduce the breakdown of accommodation/convergence.*
- *Avoid touching the surround with images having negative parallax.*
- *When changing the scale of an image of a single object, it may be best to allow the ZPS to remain at the object's center.*
- *Remember that the image should be easy to look at initially. Build in default conditions so that the image boots up with a properly placed ZPS, and maintains the constancy of ZPS with changes of scale, angle of view, and changes of the interaxial distance.*
- *Use the parallel axes camera model (to be discussed in the next two chapters). Do not use rotation for the creation of stereopairs.*
- *Don't exceed 20 millimeters for positive parallax on a display screen. This is a rough rule of thumb. Judge for yourself.*
- *Observe the image without glasses to see where you have converged and to see the maximum values of parallax.*
- *For scenic views, try putting the ZPS on the closest foreground object, and be careful to control the maximum positive values of parallax produced by distant image points. You can achieve the latter by reducing the interaxial separation and/or increasing the angle of view.*
- *If you see annoying crosstalk, and you have adjusted the parallax to low values, try changing the color of the image. Because of long phosphor persistence, green tends to ghost the most.*
- *Once you know the basics, composing stereoscopic images is an art, not a science.*

There are tens of thousands of people in the world who are using electronic stereoscopic displays. A few years ago there weren't even a hundred. In the next years a great deal will be learned about the composition and aesthetics of such images. It may well be that some of the solutions suggested here will be improved or even overthrown. That's fine, because we're at the beginning of the exploration of the medium — an exploration that is likely to continue for decades.


References

1. Inoue, T., and Ohzu, H. Measurement of the human factors of 3-D images on a large screen. *Large-Screen Projection Displays II*, SPIE Vol.1255, 1990.
2. Kaufman, Lloyd. *Sight and mind: An introduction to visual perception*. New York: Oxford University Press, 1974.
3. Lipton, Lenny. *Foundations of the stereoscopic cinema*. New York: Van Nostrand Reinhold, 1982.
4. Spottiswoode, Raymond, and Spottiswoode, Nigel. *The theory of stereoscopic transmission and its application to the motion picture*. Berkeley: University of California Press, 1953.
5. Lipton, Lenny. *Independent filmmaking*. New York: Simon & Schuster, 1983.

Stereo-Vision Formats

What's a Format?


There are several distinct means for preparing time-shared multiplexed images for stereo-vision electronic displays. Furthermore, the concept of format is distinct from the selection technique, or the means for providing each eye with its needed image (and rejecting the unnecessary image). Format and image selection must work together, and in designing such a system the selection technique may determine the format. For example, if the image is to be viewed in a stereoscope similar to the popular View-Master viewer (as mentioned in Chapter 1), the left and right images may be formatted so that they are side-by-side on the display screen.




CrystalEyes eyewear. Infrared (IR) emitter (top of monitor) sends sync information to electronic eyewear with LC shutters.

In this chapter we are concerned with formats which work in conjunction with eyewear (also known as selection devices) and use the time-multiplexing approach. Such eyewear may be active, using liquid crystal (LC) shutters like StereoGraphics' **CrystalEyes** or **SimulEyes** **VR** products, or they may be passive polarizing spectacles as are used in our projector **ZScreen**[®]. In fact, even when using the passive eyewear with our projection **Zscreen**, the user is looking through a shutter, but one whose components are distributed between the **Zscreen** electro-optics and the polarizing eyewear.

The most widespread method for viewing electronic stereo-vision images uses **CrystalEyes** on workstations for scientific visualization applications such as molecular modeling. In the most general terms, the format used by **CrystalEyes** is called the alternate field, field-sequential, or time-multiplexed technique.



ODD



EVEN

Odd and even fields. In planar video, fields combine to produce a complete interlaced picture. For stereo, odds are used for right images and lefts for even (or vice versa).

Interlace Stereo

The original stereo-vision television format used interlace to encode left and right images on odd and even fields. It's a method that's used today, and it has the virtue of using standard television sets and monitors, standard VCRs, and inexpensive demultiplexing equipment — a simple field switch that shunts half the fields to one eye and half to the other by syncing the shutters' eyewear to the field rate. The method results in flicker — which some people find more objectionable than others. One way to mitigate the flicker when viewing a monitor is to reduce the brightness of the image by adding neutral density filters to the eyewear. Another problem is that each eye sees half the number of lines (or pixels) which are normally available, so the image has half the resolution.

The interlace approach or, to be more precise, the time-multiplexed low field rate approach, has turned out to have an interesting application these days when used in conjunction with an HMD using LC displays. Because of the long persistence of LC displays, a low number of longer-lasting fields will produce a more or less flicker-free image. If a field switch is used to alternate fields to one and then the other LC display, the result can be a stereo image with barely detectable flicker.


Above-and-Below Format

Our concern has been to create stereo-vision formats and add-on selection devices which operate within the existing infrastructure of computer graphics and video hardware, without modification to the infrastructure hardware or basic working procedures. The original method we invented has survived for computer graphics, but has disappeared for video. The method uses two subfields arranged above and below each other in a single standard field. The images in these subfields are squeezed top to bottom by a factor of two.

At the standard 60 fields per second, it takes half the duration of an entire field, or 1/120th second, to scan a subfield. When played back on a monitor operating at 120 fields per second, the subfields which had been juxtaposed spatially become juxtaposed temporally. Therefore, each eye

of the beholder, when wearing the proper shuttering eyewear, will see 60 fields of image per second, out of phase with the other 60 fields prepared for the other eye. Thus, it is possible to see a flicker-free stereoscopic image because each eye is seeing a pattern of images of 1/120th second followed by 1/120th second of darkness. When one eye is seeing an image, the other is not, and vice versa.

Today there are many models of high-end graphics monitors which will run at field rates of 120 or higher. Provided that a synchronization pulse is added between the subfields in the subfield blanking area, such a monitor will have no problem in displaying such images. The monitor will automatically "unsqueeze" the image in the vertical so the picture has the normal proportions and aspect ratio. The term "interlace", when applied to such a display, is irrelevant because the stereo-format has nothing to do with interlace, unlike the odd-even format explained above. The above-and-below format can work for interlace or progressive scan images.



Above-and-below images. Squeezed in the vertical to compress images. When displayed at 120 fields per second a stereo image results.

StereoGraphics' sync doubling emitter for the above-and-below format is available for PCs. It adds the missing synchronization pulses to the vertical blanking for a proper video signal. The shutters in the eyewear are triggered by the emitter's infrared signal. This emitter is used to take a PC with a video board running at a standard rate, nominally 60 fields per second, and double it for the desired flicker-free rate. The emitter also provides an IR signal for our **CrystalEyes** eyewear whose electronics sync to the emitter's signal, which corresponds to the sync pulses (both original and added subfield pulses).

If the image has high enough resolution to begin with or, more to the point, enough raster lines, the end result is pleasing. Below 300 to 350 lines per field the image starts to look coarse on a good-sized monitor viewed from two feet, which seems quite close. However, that's about the distance people sit from workstation or PC monitors. This sheds light on why this approach is obsolete for video. NTSC has 480 active video lines. It uses a twofold interlace so each field has 240 lines. Using the subfield technique, the result is four 120-line fields for one complete stereoscopic image. The raster looks coarse, and there is a better approach for stereo video multiplexing as we shall shortly explain.

At the frequently used 1280x1024 resolution, an above-and-below formatted image will wind up at 1280 by about 500 pixels per eye, which looks good. Even from the workstation viewing distance of two feet, the image is extraordinary, and most people would agree that this is near photographic quality.

Stereo-Ready Computers


These days most graphics computers, like those from Silicon Graphics (SGI), Sun, DEC, IBM, and HP, use a double buffering technique to run their machines at a true 120 fields per second rate. Each field has a vertical blanking area associated with it which has a synchronization pulse. These computers are intrinsically outputting a high field rate and they don't need the above-and-below solution to make flicker-free stereo images, so they don't need a sync-doubling emitter to add the missing sync pulse. These computers are all outfitted with a jack that accepts the StereoGraphics standard emitter, which watches for sync pulses and broadcasts the IR signal with each pulse. Most of these machines still offer a disproportionately higher pixel count in the horizontal compared with the vertical.

Some aficionados will insist that square pixels are *de rigueur* for high-end graphics, but the above-and-below format (or most stereo-ready graphics computers) produces oblong pixels — pixels which are longer in the vertical than they are in the horizontal. The popular 1280x1024 display produces a ratio of horizontal to vertical pixels of about 1.3:1, which is the aspect ratio of most display screens, so the result is square pixels. But in the above-and-below stereo-vision version for this resolution, the ratio of horizontal to vertical pixels for each eye is more like 2.6:1, and the result is a pixel longer than it is high by a factor of two.

SGI's high-end machines, the Onyx[®] and Crimson[®] lines, may be configured to run square-pixel windowed stereo with separate addressable buffers for left and right eye views, at 960x680 pixels per eye (108 fields per second). This does not require any more pixel memory than already exists to support the planar 1280x1024. Some SGI high-end computers have additional display RAM available, and they can support other square-pixel windowed stereo resolutions, though higher resolutions come at the expense of field rate. For example, such high-end high-display-memory SGI systems support 1024x768 pixels per eye but only at 96 fields per second, which is high enough to have a flicker free effect for all but very bright images.

For most applications, having pixels which are not perfectly square can result in a good-looking picture. The higher the resolution of the image, or the more pixels available for image forming, the less concern the shape of the pixels is.

Side-by-Side



Side-by-side format. Images squeezed as shown, so left and right views fit into a standard video field.

We developed the side-by-side technique to cure a significant problem of the above-and-below technique as applied to video: not enough raster lines.


While the above-and-below solution is a good one for computer graphics applications because computer displays tend to output more raster lines than television, our video product uses different techniques to create stereo formats. First, for real-time viewing this is what we do: The left and right images from the two video camera heads making up a stereoscopic camera are fed to our View/Record box, and for viewing real-time the images are stored and then played back at twice the rate at which they were read.

In addition, the fields are concatenated or alternated to achieve the necessary left-right pattern. The result is an over-30-KHz or twice normal video bandwidth signal which preserves the original image characteristics but, in addition, is stereoscopic. What we've described here is for real-time viewing using a graphics monitor with 120 fields per second capability, and it is the function of the View section of the View/Record box to produce such a signal.

Once it becomes necessary to interface with the existing television infrastructure, a problem arises: NTSC recorders must be used if we are to have a generally useful system, and that means containing the image to within the NTSC specification of about a 15 KHz line rate. Thus the Record section of the View/Record box serves to compress the left and right images so they occupy the normal NTSC bandwidth. It does this by squeezing the images horizontally so that they occupy a single standard field. The resultant signal is in fact an NTSC signal which may be recorded on an NTSC recorder (we also make a PAL version). When played back the side-by-side image is digitized by our Playback Controller, and formatted for stereo viewing. The result is an image which has characteristics which are similar to the real-time image described above.

White-Line-Code

The White-Line-Code (WLC) system is used for multi-media PCs and it offers a high-quality but low-cost solution to the problem of stereo-vision imaging. This format doesn't care if the left and right fields are in interlace or progressive scan modes, and it doesn't care about the field rate. WLC was created to offer the most flexible possible stereo-vision




White-Line-Code. Fields are tagged so each eye sees the image meant for it, and code triggers the eyewear's shutters.

system for the content providers, developers, and users. The hardware components of the WLC allow for rapid user installation.

On the bottom of every field, for the last line of video, white lines are added to signify whether the field is a left or a right, as shown in the illustration. The last line of video was chosen because it is within the province of the developer to add the code in this area immediately before the blanking area (which is not accessible to the developer). When our electronics see the white line, it is prepared to shutter the eyewear once the vertical sync pulse is sensed.

The WLC is universal in the sense that it simply doesn't care about interlace, or progressive scan, or field rate, or resolution. If the WLC is there, the eyewear will shutter in synchrony with the fields and you'll see a stereoscopic image. Our **SimulEyes VR** product is used with the WLC in PCs, and the WLC is an integral part of the **SimulEyes VR** product. **SimulEyes VR** chips are used on video boards for driving our new low-cost eyewear; for users who wish to retrofit their PCs for stereo-vision, we have a solution with our **SimulEyes VR** controller.

The two most popular modes of operation for WLC are the page-swapping mode, which is used most often for DOS action games running at either 70 or 80 fields per second (we've worked out a driver that allows for 80 fields per second which reduces flicker considerably); the other mode is that used most often for multi-media in Windows, the IBM standard 8514A, which is a 90 fields per second interlace mode.



SimulEyes system. System hooks up to multi-media PC and uses the WLC to view in stereo.

New video boards are coming along for the PC and these use a technique which is similar to that used in stereo-ready workstations. They assign an area in the video buffer to a left or right image with a software "poke."

With the passage of time it's our expectation that the board and monitor manufacturers will allow their machines to run at higher field rates, say 120 fields per second, as is the case for workstations. When this happens the difference between PCs and workstations will entirely vanish vis-à-vis stereo-vision displays. Until that time, the lower 80 and 90 fields per second rates serve nicely to provide good-looking images in most situations.

Summing Up

We have seen that there is a wide variety of stereo formats that are used on different computer and video platforms. All of the time-multiplexed formats discussed here are to a greater or a lesser extent supported by StereoGraphics. For example, our high-end video camera may be used to shoot material for the interlace format. PC users may find themselves using either above-and-below format with the sync-doubling emitter and **CrystalEyes**, or the WLC format with **SimulEyes VR**.

Users preparing computer generated animation to be used in videos may find themselves viewing in above-and-below or the workstation's stereo-ready format. However, for recording on a standard VCR, the user may actually format the image in the side-by-side format.

Format Attributes

Format	Fields/Sec¹	Medium	Viewing Hardware
Interlace, NTSC	60	NTSC ²	Eyewear or HMD ³
Interlace, PAL	50	PAL ²	Eyewear or HMD ³
Side-by-side, NTSC	120	NTSC ²	Playback, Emitter, CE2 ⁴
Side-by-side, PAL	100	PAL ²	Playback, Emitter, CE2 ⁴
Above-and-Below	120	PC	Sync-Doubling Emitter, CE-PC ⁴
Stereo-Ready	120	Workstation	Emitter, CE2 ⁴
White-Line-Code	70-90	Multi-Media PC	SimulEyes VR

Footnotes:

1 Divide by two for field rate for one eye.


2 For broadcast, VCR, laser disc, or DVD.

3 StereoGraphics manufactures OEM eyewear for this application.

4 CE 2 is the latest model of our CrystalEyes product. CE-PC is our product for the PC.

Parallel Lens Axes

Computing a stereoscopic image involves drawing two monocular viewpoints from different positions. In the past, people have created computer-generated stereoscopic pairs by rotating an object through a few degrees, but we advise against this procedure. The use of rotation to generate stereoscopic pairs will result in vertical misalignment of corresponding left and right image points. Such a vertical misalignment causes the eye muscles to work in an extraordinary way which most people experience as discomfort. If the eyes try to fuse vertical parallax, which contains no depth information, the result can be painful.




Rotation produces distortion.

The rectangle $ABCD$, on the left, has its vertical axis indicated by the dotted line. When rotated about the axis, as shown, corner points A' , B' , C' , and D' are now higher or lower than corresponding points A , B , C , and D . Despite the fact that this example involves a simple figure, it's typical of what happens when rotation is employed for producing stereoscopic pairs. Instead of rotation, we recommend producing the two views as a horizontal perspective shift along the x axis, with a horizontal shift of the resulting images (HIT — horizontal image translation) to establish the ZPS (zero parallax setting). We will now describe a mathematical model of a stereo camera system, using the method of horizontal perspective shifts without rotation to produce the proper stereoscopic images.

The Camera Model

The geometry of the basic algorithm for producing computer-generated electro-stereoscopic images is illustrated here, showing left and right cameras located in data space. Their interaxial separation is given by t_c . The camera lens axes are parallel to each other in the z direction. The distance from the cameras to the object of interest is d_o .



Distortion-free images. The camera axes (or the axes of the centers of perspective) must be parallel.

We have chosen a camera model because of the diverse approaches and nomenclatures that people use for computer graphics. It's our belief that everybody understands something about photography these days because of the wide acceptance of 35-millimeter and video cameras. Imagine that the cameras shown are two still or video cameras whose lens axes are parallel, mounted on a flat base. t_c is the interaxial separation or distance between the cameras or, more accurately, the centers of the lenses. The cameras are mounted so that the lenses' axes — lines passing through the centers of the lenses and perpendicular to the imaging planes — always remain parallel. Imagine that we're looking at an object in this data space. The two cameras will produce two different perspective viewpoints, because they are horizontally shifted the distance t_c .

The two cameras are using lenses whose focal lengths are identical. Short-focal-length lenses produce a wide-angle view, and long-focal-length lenses produce a narrow angle of view. For example, wide-angle lenses for 35-millimeter photography are generally anything under 50 millimeters, and long-focal-length or telephoto lenses are generally anything above 50 millimeters.

Let's assume we have two video cameras with lenses looking straight ahead, each seeing its perspective viewpoint of an object in space. Let's switch the video signal between the cameras at field rate.¹ If we were to take a look at the images on a television monitor, we would see two images which are not exactly superimposed, because the cameras are t_c apart. Since the cameras are pointing straight ahead with parallel lens axes, they will see different parts of the object. Let's horizontally shift these two images relative to each other, so that one portion of the image is superimposed. Wherever the image is superimposed the parallax is zero, and that portion of the object (or scene) appears at the plane of the screen.

This has been called "convergence," but that term is easily confused with human physiology and the rotation of the eyes needed for fusion, as discussed on page 20. Therefore, we use the term *HIT* (horizontal image translation), which is used to achieve *ZPS* (zero parallax setting).

If the images are horizontally shifted so that some part of the image is perfectly superimposed, that part of the image is at ZPS. Remember that our object exists not only in the x and y axes, but also in the z axis. Since it is a three-dimensional object, we can achieve ZPS for only one point (or set of points located in a plane) in the object. Let's suppose we use ZPS on the middle of the object. In this case, those portions of the object which are behind the ZPS will have positive parallax, and those which are in front of the ZPS will have negative parallax.

This parallel lens axes algorithm for stereoscopic computer-generated images uses two camera viewpoints, with parallel lens axes, set some distance t_c apart. Both viewpoints, or camera lenses, have the same angle of view. The degree to which we horizontally shift or translate (HIT) the images is more than a matter of taste. Earlier we explained the advantages of creating images with low parallax. We must think of using ZPS to produce the best compromise of parallax for the entire image.

This approach involves no rotation and leads to geometric-distortion-producing vertical parallax, but we will have achieved the two perspective viewpoints needed for a stereoscopic image.

If HIT is used as described, with other conditions (angle of view of lenses, distance of object from the lenses, t_c) kept constant, there is no change in depth content. It will take a moment for your eyes to reconverge for the different values of parallax, but the total depth content of the image remains the same. After your eyes readjust, the image will appear to be just as deep as it was before. However, you may have produced, through HIT, images which appear to be entirely in front of or behind the screen. As discussed earlier, it's best to place the ZPS at or near the center of the object to reduce the breakdown of accommodation and convergence and to reduce the ghosting artifact.

The Parallax Factor

We suggest the use of small t_c values and a wide angle of view, at least 40 degrees horizontal. To help understand why, let us take a look at the relationship:

$$P_m = Mf_c t_c \left(\frac{1}{d_o} - \frac{1}{d_m} \right)$$

This is the depth range equation used by stereographers, describing how the maximum value of parallax (P_m) changes as we change the camera setup. Imagine that by applying HIT, the parallax value of an object at distance d_o becomes zero. We say we have achieved ZPS at the distance

d_o . An object at some maximum distance d_m will now have a parallax value of P_m .

The aim is to produce the strongest stereoscopic effect without exceeding a maximum of 1.5° of parallax. The form of this equation is helpful in understanding this. The value of magnification (M) will change the value of P_m . For example, the image seen on a big screen will have greater parallax than the image seen on a small screen. A 24-inch-wide monitor will have twice the parallax of a 12-inch monitor, all things being equal. Reducing t_c will reduce the value of screen parallax. We also note that reducing lens focal length f_c (using a wide-angle lens) also reduces P_m .

The most important factor which controls the stereoscopic effect is the distance t_c between the two camera lenses. The greater t_c , the greater the parallax values and the greater the stereoscopic depth cue. The converse is also true. If we are looking at objects that are very close to the camera — say coins, insects, or small objects — t_c may be low and still produce a strong stereoscopic effect. On the other hand, if we're looking at distant hills t_c may have to be hundreds of meters in order to produce any sort of stereoscopic effect. So changing t_c is a way to control the depth of a stereoscopic image.

Small values of t_c can be employed when using a wide angle of view (low f_c). From perspective considerations, we see that the relative juxtaposition of the close portion of the object and the distant portion of the object, for the wide-angle point of view, will be exaggerated. It is this perspective exaggeration that scales the stereoscopic depth effect. When using wide-angle views, t_c can be reduced.²

The reader ought to verify the suggestions we're making here. Producing computer-generated stereoscopic images is an art, and practice makes perfect. It is up to you to create stereoscopic images by playing with these parameters. As long as you stick with the basic algorithm we've described (to eliminate geometric distortion), you have a free hand in experimenting with t_c , d_o , and angle of view (f_c). This parallel lens axes approach is particularly important when using wide-angle views for objects which are close. In this case, if we use rotation, the geometric distortion would exacerbate the generation of spurious vertical parallax.

Remember that the strength of the stereoscopic effect is controlled by the interaxial distance t_c , which controls the parallax values of the object or the scene. Once the interaxial setting is established, HIT is used to control the ZPS of the images, and is often achieved in software by moving two portions of the buffer relative to each other, as discussed in

the next chapter. The horizontal movement of the left and right images in the upper and lower portions of the buffer produces HIT, which is used to control ZPS.

In terms of producing user-friendly software, consideration should be given to the fact that when an image is booted up, it should have the proper ZPS. For a single-object image, a good place for ZPS is the middle of the object. If this suggestion is not followed, the user will have a difficult time because of the breakdown of accommodation and convergence, and because of the concomitant crosstalk. It's important to implement the ZPS condition in software as a default condition. In addition, when changing the size of the image it will also be desirable to maintain the ZPS, and this can also be implemented in software. Given our parallel lens axes camera model, when t_c is changed the ZPS will also change; and it is desirable to implement, in software, a constancy of ZPS even when t_c is varied. This effort will help the user produce a better-quality, easier-to-look-at stereoscopic image.

Windows and Cursors

One important concept for windows is that alphanumeric must appear in both the left and right fields, and it's often appropriate for this material to be at ZPS. Alphanumerics (and the vertical and horizontal lines that make up the window) should appear in both the left and right fields in the same x,y locations. There may be occasions when it is desirable to have windows which are in CRT or viewer space. This may be appropriate when multiple windows are used.

Cursors have the same benefits as they have in the planar display mode but, in addition, for a stereoscopic display they can also provide z -axis-locating with a great degree of accuracy. People who have been active in photogrammetry use stereoscopic cursors, typically in the form of a cross, in order to measure height or depth. A cursor appearing in the two subfields (see the next chapter) in exactly the same location will have zero parallax, but it's also important to be able to move the cursor in z -space. By moving one or both of the cursors in the left and right subfields in the horizontal direction (the x direction), the cursor will appear to be moving in z -space. By counting x pixels, the exact location of the cursor can be determined. If additional information about the capture or photography of the image is provided, one will then be able to translate the counting of x pixels into accurate z measurements. The depth sense of binocular stereopsis is exquisitely sensitive, and a skilled operator will be able to place a cursor to an accuracy of plus-or-minus one pixel.

References

1. Lipton, Lenny, and Meyer, Lhary. A time-multiplexed two-times vertical frequency stereoscopic video system. *International Symposium Digest*, SID 1984, Vol.XV.
2. MacAdam, D.L. Stereoscopic perceptions of size, shape, distance and direction. *SMPTE Journal*, 1954, 62:271-93.


Creating Stereoscopic Software

In this chapter we take a more mathematical approach than in the early chapters of the book. There are two important requirements for conveying stereoscopic three-dimensional images on a single screen:

- Software must perform an offset perspective projection for each eye, thus simulating what each eye would see if it were immersed in the three-dimensional virtual world that the software renderings are based on, and
- The left-eye and right-eye renderings must be formatted to the computer display system in a fashion that allows the stereoscopic hardware system to separate out left-eye and right-eye information.

This chapter is divided into three parts. The first, creating stereoscopic pairs, deals with the first requirement. The second, the above-below format, discusses how to meet the second requirement for **CrystalEyes** hardware. Finally, the third, the **SimulEyes VR** format, discusses how to meet the second requirement for **SimulEyes VR** hardware that uses the WLC stereoscopic format.

CREATING STEREOSCOPIC PAIRS



Perspective projection of two points onto *xy* plane.


Perspective Projections

These left- and right-eye projections are similar to the single-perspective projection normally used for converting a three-dimensional data set to a monoscopic rendering. The reader will note that, to be consistent with most programming environments, a “right-handed” coordinate system is used — with the positive *z*-axis pointing out from the screen — throughout this chapter. Please keep in mind that the positive nature of the *z*-axis may create confusion because image points along this portion of the axis are defined, per optical terminology, to have negative parallax.

There are several common methods for drawing three-dimensional figures onto a two-dimensional screen. The simplest, *orthographic projection*, involves merely throwing away the troublesome *z*-coordinate.

For a three-dimensional effect (whether monoscopic or stereoscopic), a *perspective projection* is the proper way to render coordinates onto a planar screen. This projection yields images with a good sense of depth by placing both the screen’s surface and the viewer’s eye within the coordinate system. If a line from the eye through an arbitrary point in space intersects the screen, that intersection is where the point will be plotted on the screen.

By drawing a line through the virtual eye that lies in the coordinate system and a 3-D point, and finding the intersection of that line with the portion of the *xy* plane that represents the screen (if such an intersection exists), the resulting intersection represents the screen location where that point is plotted. Points *A* and *B* in the drawings lie in the 3-D coordinate system. Their intersections with the *xy* plane, *A*₁ and *B*₁, indicate where points *A* and *B* will be projected on the screen.



Perspective projection changes as viewpoint changes.

Assuming that the screen lies in the *xy* plane, and that the eyepoint (technically called the center of projection) lies on the positive *z*-axis at $(0,0,d)$, an arbitrary point (x,y,z) will be plotted at

$$\left(\frac{xd}{d-z}, \frac{yd}{d-z} \right) \tag{1}$$

provided that the point is in front of rather than behind the eye. What happens when the eye is moved to the right, from $(0,0,d)$ to, say, $(t_c,0,d)$? When moving to the right in real life, the scene changes with the new viewpoint. In fact, the same is true within a perspective projection.

The drawing shows that the eye was moved from position 1 a distance of t_c to the right, to position 2. Note how the projection of point A moved to the right, and the projection of point B moved to the left. At the new eye position of $(t_c, 0, d)$, an arbitrary point (x, y, z) will now be plotted at

$$\left(\frac{xd - zt_c}{d - z}, \frac{yd}{d - z} \right) \quad (2)$$

The same result can be achieved by subtracting t_c from the x -coordinate of the arbitrary point, performing the simple z -axis perspective projection using equation 1, and adding t_c to the resulting x -coordinate.

Stereo Perspective Projections

Left-eye and right-eye images are created for stereoscopic viewing by calculating projections from two different eye positions onto a single screen. Rather than offsetting the right eye by t_c as shown above, both eyes are offset by $t_c/2$ in opposite directions. The twin centers of projection straddle the negative z -axis. Point (x, y, z) will be plotted at

$$\left(\frac{xd + \frac{zt_c}{2}}{d - z}, \frac{yd}{d - z} \right) \quad (3)$$

for the left eye, and at

$$\left(\frac{xd - \frac{zt_c}{2}}{d - z}, \frac{yd}{d - z} \right) \quad (4)$$

for the right eye, where t_c is the eye separation, and d is the distance from the eyes to the viewing plane. Note that the transformed y coordinates for both left and right eyes are the same.

Since equations 3 and 4 are based on a model in which both the screen and the eyes lie in the same space as the 3-D objects to be drawn, both t_c and d must be specified in that same coordinate system. Since a person's eyes are about 64 millimeters apart, and are positioned about 60 centimeters from the screen, one might assume it to be best to just plug in these "true" values. However, if the screen is displaying a view of the solar system, for example, the values of t_c and d will be about 12 orders

of magnitude too small. Even if, coincidentally, a 30-centimeter-wide screen is displaying a 30-centimeter-wide scene, it is still not necessarily best to use the “true” t_c and d values. The image might look better with, say, t_c set to three centimeters and d set to 50 centimeters. We want to set the values of t_c and d such that the resulting stereoscopic image will convey a sense of depth, but not so much depth that the image will be unpleasant or uncomfortable to view. Setting a value for t_c is discussed in Chapter 3.

Parallax: Too Much of a Good Thing?

In the drawing, A_{left} , the left-eye projection of A , lies to the left of the right-eye projection, A_{right} . If we subtract the x -coordinate of A_{left} from that of A_{right} , the resulting parallax will be a positive number, which we call *positive parallax*. To the viewer of the projected image, point A will appear to lie behind the surface of the screen. On the other hand, point B has *negative parallax*: the x -coordinate of B_{right} minus that of B_{left} is a negative number. To the viewer, point B will appear in front of the screen. If a point happens to lie right on the projection plane, the left- and right-eye projections of the point will be the same; they’ll have *zero parallax*.

Parallax can also be described in terms of viewing angle. If the right- and left-eye projections of a point appear P centimeters apart, and the user’s eyes are d centimeters away from the screen, then the parallax angle is

$$\beta = 2 \arctan \frac{P}{2d} \quad (5)$$

As discussed on page 12, parallax angle β values exceeding 1.5° tend to make viewing uncomfortable.

Using other techniques in rendering the image (i.e. hidden-surface removal, shading, depth cueing, etc.) will compliment the depth effect, and allow one to play with the value of β_{max} . Finally, perception depends on the individual; some people can handle more parallax on a screen than others.

If β_{max} is set to 1.5° , you may want to set β_{min} to -1.5° , thus getting as much depth as possible both in front of and behind the surface of the screen by placing the ZPS in the center of the object. However, it may be better to set β_{min} closer to 0° , or perhaps even at 0° . Why do that? If an object that appears to sit in front of the screen is cut off by the left or

right edge of that screen, the resulting effect of a far object (the screen edge) obscuring a near object (the thing displayed in negative parallax) will interfere with the user's ability to make sense of the image, as discussed on page 17.

Putting the Eyes Where They Belong

In order to obtain a stereoscopic image with enough parallax to convey a sense of depth, but not so much parallax that the effect bothers the viewer, we must place the two centers of projection, which represent the two eyes, at an appropriate distance from the projection plane (d), and with an appropriate separation from each other (t_c). Also, if we are going to use negative parallax, the z -coordinate of the plane of zero parallax (z_{zps}) must be selected.

There are two possible approaches for determining the value of z_{zps} (aside from the no-negative-parallax situation, where z_{zps} simply equals z_{max}):

- 1) Specify a z_{zps} value in advance, and calculate the eye separation based on that, and
- 2) calculate the z_{zps} value that optimizes use of the parallax range from β_{min} to β_{max} .

The first method is best where one wants the plane of zero parallax at a z -coordinate value of 0. The second method is best where software must choose z_{zps} automatically for a variety of possible coordinate ranges, and it is desirable to make use of the full allowable parallax range from β_{min} to β_{max} .

Given that

θ = the field-of-view angle ($\theta = 40^\circ$ to 50°),

x_{min}, x_{max} = the expected range of world x -coordinate values that will be projected onto the screen (the x range specifically refers to the x -coordinate range, on the plane of zero parallax, that ought to appear in the screen viewport),

z_{min}, z_{max} = the expected range of world z -coordinate values that will be projected onto the screen (this range is likely to be a narrower one than that specified by the far and near clipping planes),

β_{\max} = the maximum positive parallax angle (1.5° is a good value),

β_{\min} = the minimum parallax angle (generally equal to $-\beta_{\max}$ if negative parallax permitted, or 0° if not),

P_{\max} = the maximum parallax specified in world coordinates, and

P_{\min} = the minimum parallax specified in world coordinates,

here are the formulas for determining d , t_c , and z_{zps} in world coordinates:

$$\Delta x = x_{\max} - x_{\min} \quad (6)$$

$$\Delta z = z_{\max} - z_{\min} \quad (7)$$

$$d = \frac{\Delta x}{2 \tan \frac{\theta}{2}} \quad (8)$$

$$P_{\max} = \frac{\tan \frac{\beta_{\max}}{2}}{\tan \frac{\theta}{2}} \Delta x \quad (9)$$

$$P_{\min} = \frac{\tan \frac{\beta_{\min}}{2}}{\tan \frac{\theta}{2}} \Delta x \quad (10)$$

If $\beta_{\min} = 0^\circ$, then $z_{zps} = z_{\max}$ and:

$$t_c = P_{\max} \left(1 = \frac{d}{\Delta z} \right) \quad (11)$$

To specify a z_{zps} value in advance, and calculate the eye separation based on that:

$$t_{c1} = -P_{\min} \left(\frac{d}{z_{\max} - z_{zps}} - 1 \right) \quad (12)$$

$$t_{c2} = P_{\max} \left(\frac{d}{z_{zps} - z_{\min}} + 1 \right) \quad (13)$$

$$t_c = \text{the smaller of } t_{c1} \text{ and } t_{c2} \quad (14)$$

Alternatively, to calculate the t_c and z_{zps} values that optimize use of the parallax range from β_{\min} to β_{\max} :

$$n = \frac{d(P_{\max} - P_{\min})}{2\Delta z} + \frac{P_{\max} + P_{\min}}{2} \quad (15)$$

$$t_c = n + \sqrt{\left(n^2 - P_{\max}P_{\min}\right)} \quad (16)$$

$$z_{zps} = z_{\min} + \frac{d P_{\max}}{t_c - P_{\min}} \quad (17)$$

Having defined the necessary environment variables, we can apply equations 3 and 4 to convert any point (x,y,z) in the acceptable three-dimensional range to two-dimensional plane coordinates. However, equations 3 and 4 assume that:

- 1) The projection plane is the xy plane ($z_{zps} = 0$), and
- 2) The eye viewpoints straddle the z -axis ($x_{\text{mid}} = 0$ and $y_{\text{mid}} = 0$,

$$\text{where } x_{\text{mid}} = \frac{x_{\min} + x_{\max}}{2} \text{ and } y_{\text{mid}} = \frac{y_{\min} + y_{\max}}{2}.$$

Accepting that z_{zps} , x_{mid} , and y_{mid} may have any values, the arbitrary point (x,y,z) now plots to

$$\left(\frac{d(x - x_{\text{mid}}) - \frac{t_c}{2}(z_{zps} - z)}{d + z_{zps} - z} + x_{\text{mid}}, \frac{d(y - y_{\text{mid}})}{d + z_{zps} - z} + y_{\text{mid}} \right) \quad (18)$$

for the left-eye view, and

$$\left(\frac{d(x - x_{\text{mid}}) + \frac{t_c}{2}(z_{zps} - z)}{d + z_{zps} - z} + x_{\text{mid}}, \frac{d(y - y_{\text{mid}})}{d + z_{zps} - z} + y_{\text{mid}} \right) \quad (19)$$

for the right-eye view.

These equations can also be specified as transformation matrices. For the left-eye view:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ \frac{x_{\text{mid}} - \frac{t_c}{2}}{-d} & -\frac{y_{\text{mid}}}{d} & 1 & -\frac{1}{d} \\ -\frac{z_{\text{zps}}}{d} \left(x_{\text{mid}} - \frac{t_c}{2} \right) & -\frac{z_{\text{zps}} y_{\text{mid}}}{d} & 0 & 1 - \frac{z_{\text{zps}}}{d} \end{bmatrix} \quad (20)$$

and for the right-eye view:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ \frac{x_{\text{mid}} + \frac{t_c}{2}}{-d} & -\frac{y_{\text{mid}}}{d} & 1 & -\frac{1}{d} \\ -\frac{z_{\text{zps}}}{d} \left(x_{\text{mid}} + \frac{t_c}{2} \right) & -\frac{z_{\text{zps}} y_{\text{mid}}}{d} & 0 & 1 - \frac{z_{\text{zps}}}{d} \end{bmatrix} \quad (21)$$

On a graphics workstation, it is likely that the projections specified above can be efficiently implemented by specifying the equivalent graphics manipulations. For the left-eye view:

$$\text{Translate by } \left(-x_{\text{mid}} + \frac{t_c}{2}, -y_{\text{mid}}, z_{\text{zps}} \right)$$

Do projection from $(0,0,d)$ onto xy plane

$$\text{Translate by } \left(x_{\text{mid}} - \frac{t_c}{2}, y_{\text{mid}}, -z_{\text{zps}} \right)$$

and for the right-eye view:

$$\text{Translate by } \left(-x_{\text{mid}} - \frac{t_c}{2}, -y_{\text{mid}}, z_{\text{zps}} \right)$$

Do projection from $(0,0,d)$ onto xy plane

$$\text{Translate by } \left(x_{\text{mid}} + \frac{t_c}{2}, y_{\text{mid}}, -z_{\text{zps}} \right)$$

Interacting with the User

Some graphics applications give the user more freedom to manipulate images than others. If the user is able to change the viewing configuration, the application needs to make the necessary adjustments so that the user will continue to get a high-quality stereoscopic image, with comfortable parallax.

Some examples:

- If an x - or y -coordinate translation occurs, the two centers of projection (eyes) need to remain centered about a point perpendicular to the center of the screen.
- If a z -coordinate translation occurs, the z -coordinate of the ZPS (z_{zps}) will need to be changed.
- If there is a scaling, or some other operation that changes the x , y , or z coordinate ranges, then z_{zps} , as well as the distance from the eyes to z_{zps} (d), and the eye separation (t_c) may need to be changed.

Be sure to test your application to insure that stereoscopic scenes remain comfortable and effective throughout a user's interaction with it.

THE ABOVE-BELOW FORMAT

Above-Below Subfields

The standard **CrystalEyes** system displays stereoscopic images by doubling the screen refresh rate, thus causing left- and right-eye stereo pair elements to overlap. An infrared link synchronizes liquid crystal shutter glasses, allowing the user to perceive true stereoscopic depth.

In order to be compatible with this stereoscopic hardware system, graphics software must run full-screen, using the Above-Below stereoscopic format. This stereoscopic format consists of mapping a vertically-compressed left-eye subfield to the top half of the display, and a vertically-compressed right-eye subfield to the bottom half of the display. This technique is described on page 28.

There is one more important detail to add: Every vertical scan of the monitor screen includes hundreds of visible horizontal scan lines, followed by 20 to 40 invisible ones. The monitor normally “draws” these invisible ones while the electron beam which scans the image returns to the top of the screen. When writing to a screen that will be drawn at double the normal refresh rate, there needs to be a blank horizontal area with a vertical thickness (exactly the same as that of the normal vertical blanking interval) inserted between the subfields. When drawn at double the normal refresh rate, the top of the blanking interval will represent the bottom of the first (left-eye) subfield, and the bottom of the blanking interval the top of the second (right-eye) subfield.

Where the normal screen resolution is $XWIDTH$ -by- $YHEIGHT$ visible pixels, and the normal vertical blanking interval is $YBLANK$, the blanking separation between where the left-eye and right-eye views are mapped in normal screen coordinates will also be $YBLANK$. The left-eye and right-eye views will each be mapped to screen coordinate regions $XWIDTH$ pixels wide and

$$YSTEREO = \frac{(YHEIGHT - YBLANK)}{2} \quad (22)$$

pixels high. The bottom row of the lower right-eye view will map to row

$$YOFFSET_{right} = 0 \quad (23)$$

and the bottom row of the upper left-eye view will map to row

$$YOFFSET_{left} = YSTEREO + YBLANK \quad (24)$$

A point with coordinates (x,y) in the right-eye view would be mapped to screen coordinates

$$\left(XWIDTH \frac{x - x_{min}}{x_{max} - x_{min}} , YSTEREO \frac{y - y_{min}}{y_{max} - y_{min}} \right) \quad (25)$$

and a point with coordinates (x,y) in the left-eye view would be mapped to screen coordinates

$$\left(XWIDTH \frac{x - x_{min}}{x_{max} - x_{min}} , YSTEREO \frac{y - y_{min}}{y_{max} - y_{min}} + YOFFSET_{left} \right) \quad (26)$$

where (x_{min}, y_{min}) are the world coordinates that correspond to the lower

left corner of the viewport, (x_{\max}, y_{\max}) are the world coordinates that correspond to the upper right corner, and where screen coordinates $(0, 0)$ represent the lower left corner of the display.

While it is generally easy for the software to find out the resolution (*XWIDTH* and *YHEIGHT*) of the current display, depending on the environment, software will probably need to obtain the value for the vertical blank interval, *YBLANK*, indirectly. The “up” and “down” arrow keys could be used to line up two horizontal line segments to set the correct value of *YBLANK* as described on page 23.

Initialization and Interface Issues for Above-Below Stereo

- Because this stereoscopic hardware system alters the sync rate of the monitor, converting the entire display to stereo mode, Above-Below stereoscopic applications must run full-screen.
- Be sure to prompt the user for blank interval information. Ideally, there should be some kind of calibration function available. Note that the blank interval may be different at different display resolutions.
- The allowable vertical range of the cursor will need to be changed from its full range (0 to $YHEIGHT - 1$) to the range for one eye's subfield (0 to $YSTEREO - 1$). This will result in a cursor that is visible in only one eye, which is preferable to a cursor that wraps from one eye's subfield to the other. If the programming environment permits two cursors, one could not only map the cursor for both eyes, but also use parallax to position the cursor in the z direction.
- Interface elements, such as scroll bars, menus, and dialog boxes either need to be drawn to both stereoscopic subfields, or they will be visible in only one eye.
- Just before exiting, be sure to undo any environment changes made upon initialization.

THE SIMULEYES VR STEREOSCOPIC FORMAT

SimulEyes VR White Line Code

The new **SimulEyes VR** stereoscopic hardware system does not double the existing sync rate of the display in order to display stereoscopically. Instead, it changes to a higher sync rate, in a page-swapping or interlaced video mode, and the liquid crystal shutter glasses synchronize with alternating screen refreshes.

The simplest way to be compatible with this stereoscopic hardware system is to use the WLC stereoscopic format, placing left-eye information on even-numbered pixel rows (where the top pixel row is row #0, an even-numbered row), and right-eye information on odd-numbered pixel rows.

This stereoscopic layout must be accompanied by a two-row bitmap, the WLC, filling the display's bottom two pixel rows. The upper row of this two-row bitmap should be pure white for the leftmost 25% of the pixel-row, and pure black for the remainder of that row. The lower row should be pure white for the leftmost 75% of the pixel-row, and pure black for the remainder of that row. (If these two pixel-rows are switched, with the 75% white pixel-row appearing on top, the display's even-numbered pixel rows will be conveyed to the right eye, and odd-numbered pixel-rows to the left eye, opposite from normal.)

In order for this Interline mapping to be compatible with the stereoscopic hardware, the display must be in an interlaced video mode. Switching to interlace mode may be done manually, but ideally this switch should be done automatically by the application software. StereoGraphics provides software development tools that make it easy to include this functionality in your application.

Initialization and Interface Issues for SimulEyes VR Stereo

- Unlike Above-Below format stereoscopy, **SimulEyes VR** format stereoscopy works well with windowed application environments, and does not require that applications run full-screen. Application windows may be full-screen, or they may be windowed with other stereoscopic or non-stereoscopic applications.
- Note that it is important to place left-eye information on even pixel rows, and right on odd pixel rows (unless reversed using the appropriate variation of the WLC), *with respect to the*

overall display, not with respect to the application window. For example, if the application window begins on the display's pixel row #101, an even-numbered pixel row of the application window will actually be an odd-numbered pixel row with respect to the overall display.

- In order for the stereoscopic hardware system to work, the display needs to be running in an interlace video mode. The Software Development Kit includes library routines to initialize the hardware system for stereoscopy, and to restore the system when the program is about to exit. Consult the Software Development Kit documentation for further details.
- The WLC bitmap must not be confined to a window, but should fill the two bottom pixel-rows of the overall display. In addition, care should be taken to insure that the WLC bitmap is not covered up by other display elements such as scroll-bars, windows, or icons. StereoGraphics provides software tools for easy implementation of the WLC bitmap in your application.
- Non-stereoscopic interface elements, such as cursors, menus, and dialog boxes should work reasonably well within stereoscopic applications. Greater visual clarity can be achieved by using larger-than-average cursors and fonts.
- Just before exiting, be sure to undo any environment changes made upon initialization.

Pseudoscopic Images

Pseudoscopic stereo occurs when the right eye sees the image intended for the left eye, and vice versa. This might occur for a number of reasons-- for example, if your application allows a user to import an interlined stereoscopic image in which it is unknown whether the top row is left-eye or right-eye information. In such a case there will be a conflict between the stereopsis cue and the monocular cues. Some users will find the effect disturbing and some won't seem to notice it.

To avoid this, it is important to keep track of which eye's view is being sent to which overall even/odd sets of pixel rows, and be sure to use the correct variation of the "white code" when displaying stereoscopic scenes. Pseudoscopic stereo may be corrected by moving the overall stereoscopic image either up or down by one pixel row, or by switching the "white code."

The Page-Swap Alternative

An alternative to the Interline stereoscopic format, for use with the same **SimulEyes VR** hardware system, is to use page-swapped video display. This method has the advantage of not requiring Interline mapping of the stereoscopic subfields. Page-swapped stereoscopy may be preferable to the Interline format when using low-resolution video modes that permit high-speed page-swapping.

Using this method, left- and right-eye subfields should be drawn to separate video buffers, which then must be page-swapped, with such video page-swaps synchronized to the display's vertical refresh events. The overall vertical sync rate should be significantly higher than 60 Hz (usually either 70 Hz or 80 Hz) in order for the page-swapped display to not have an uncomfortable flicker effect.

To implement the “white code” bitmap in a page-swapped stereoscopic application, you need to fill in the bottom pixel-row of the left- and right-eye subfields as follows:

- The bottom pixel-row of the left-eye subfield should be pure white for the left-most 25% of the pixel-row, and pure black for the remainder of the row.
- The bottom pixel-row of the right-eye subfield should be pure white for the leftmost 75% of the pixel-row, and pure black for the remainder of the row.

Glossary

Accommodation. Focusing of the eyes.

Accommodation/Convergence Relationship. The habitual relationship established between the focusing and converging of the eyes when looking at an object in the visual world.

Binocular. Two eyes. The term “binocular stereopsis” is used in some psychology books for the depth sense simply described as “stereopsis.”

Conjugate Points. See **Corresponding Points**.

Convergence. The rotation of the eyes, in the horizontal direction, producing fusion. The term has also been used to describe the movement of left and right image fields or the “toe-in” or rotation of camera heads.

Corresponding Points. The image points of the left and right fields referring to the same point on the object. The distance between the corresponding points is defined as **parallax**.

Crosstalk. Incomplete isolation of the left and right image channels. Crosstalk is a physical entity, whereas **ghosting** is a psychophysical entity.

CRT Space. The region appearing to be within a CRT display screen or behind the surface of the screen. Images with negative parallax will appear to be in CRT space.

Depth Range. A term that applies most appropriately to stereoscopic images created with cameras. It is the distance from the ZPS point to the background point having parallax which is acceptable in terms of ease of viewing.

Disparity. The distance between conjugate points on overlaid retinae, sometimes called **retinal disparity**. The corresponding term for the display screen is **parallax**.

Extrastereoscopic Cues. Those depth cues appreciated by a person using only one eye. Also called **monocular cues**.

Field-Sequential. In the context of electro-stereoscopy, the rapid alternation of left and right perspective viewpoints as fields are written on the video display screen.

Fusion. The combination, by the mind, of the left and right images — seen by the left and right eyes — into a single image.

Ghosting. The perception of **crosstalk** is called **ghosting**.

HIT. Horizontal image translation. The horizontal shifting of the two image fields to change the value of the parallax of corresponding points. The term **convergence** has been used to denote this concept.

Homologous Points. See **Corresponding Points**.

Infrared (IR). Invisible radiation used as the communications link for **CrystalEyes**.

Interaxial Distance. Also **interaxial separation**. The distance between camera lenses' axes. See *t*.

Interocular Distance. See **Interpupillary Distance**.

Interpupillary Distance. Also **interpupillary** or **interocular separation**. The distance between the eyes' axes. See *t*.

Monocular Cues. See **Extrastereoscopic Cues**.

Multiplexing. The technique for placing the two images required for an electro-stereoscopic display within an existing bandwidth.

Parallax. The distance between conjugate points. It may be measured with a ruler or, given the distance of an observer from the screen, in terms of angular measure. In the latter case, the parallax angle directly provides information about disparity.

Planar. Flat. Two-dimensional. A planar image is one contained in a two-dimensional space, but not necessarily one which appears flat. It may have all the depth cues except stereopsis.

Plano-Stereoscopic. A stereoscopic display which is made up of two planar images.

Retinal Disparity. See **Disparity.**

Selection Device. The hardware used to present the appropriate image to the appropriate eye and to block the unwanted image.

Side-by-Side. The format used by StereoGraphics' video system to squeeze left and right images into a single video field.

Stereo. Short for **stereoscopic**. I've limited its usage because the term is also used to denote stereophonic audio.

Stereoplexing. Stereoscopic multiplexing. A means to incorporate information for the left and right perspective views into a single information channel without expansion of the bandwidth.

Stereopsis. The binocular depth sense. Literally, "solid seeing."

Stereoscope. A device for viewing plano-stereoscopic images. It is usually an optical device with twin viewing systems.

Stereoscopy. The art and science of creating images with the depth sense **stereopsis**.

Surround. The vertical and horizontal edges of the screen immediately adjacent to the display.

***t*.** In stereoscopy, *t* is used to denote the distance between the eyes, called the **interpupillary** or **interocular** distance. t_c is used to denote the distance between stereoscopic camera heads' lens axes.

Viewer Space. The region between the display screen surface and the viewer. Images will appear in this region if they have negative parallax.

White Line Code. The last line of the video signal used to tag left and right perspectives and trigger the **SimulEyes VR** eyewear shutters to sync with the fields.

Window. The stereo window corresponds to the screen surround given the algorithms described in these pages.

ZPS. Zero parallax setting. The means used to control screen parallax to place an object, or a portion of an object or picture, in the plane of the screen. ZPS may be controlled by HIT. When image points have ZPS, prior terminology would say that they were “converged.” That term is not used here in that sense, because it may be confused with the convergence of the eyes, and because the word implies rotation of camera heads or a software algorithm employing rotation. Such rotation invariably produces geometric distortion.

Index

- Above-and-below technique, 28-29, 33, 49-51
- Accommodation/convergence relationship, 11, 12, 13, 16-17, 19, 20-21, 23, 25, 37, 39, 55
- Active eyewear. *See* Eyewear
- Aerial mapping, 19
- Aerial perspective as a depth cue, 3
- Afterglow, 12-13, 25
- Akka, Robert, 41
- Alphanumerics, 39
- Alternate field technique. *See* Time-multiplexed technique
- Angle of view, 25, 36, 37, 44
- Animation, 33
- Arrow keys, 23-24
- Axes
 - of eyes, 11
 - of lenses, 10, 23

- Beta (β) (parallax angle), 44-47
- Binocular perspectives, 2, 55
- Blanking intervals, 50-51
- Brewster, David, 8
- Broadcast video, 33

- Camera model, 36-37
- Cathode ray tube. *See* CRT
- CE-PC eyewear. *See* CrystalEyes display and eyewear
- CE2 eyewear. *See* CrystalEyes display and eyewear
- Color of the image, 25
- Composition of stereoscopic images, 15-26
- Computer animation, 33
- Computer-generated images, 19, 20-21, 35-40
- Conjugate image points (corresponding image points), 9, 55

Convergence
 angle, 22
 defined, 55
 of the eyes, 7-8, 11, 20, 23, 36, 37
 See also Accommodation/convergence relationship
 Crimson (SGI computers), 30
 Crosstalk (ghosting), 12-13, 23, 25, 37, 39, 55
 CRT phosphor afterglow, 12-13, 25
 CRT space, 10, 17-18, 20-21, 39, 56
 CrystalEyes display and eyewear, 2, 8-9, 27-28, 29, 33, 49-51
 Cursors, 39, 51, 53

d (distance from viewer's eyes to screen), 42-49
 Databases, three-dimensional, 2
 Depth cues, 1-5, 15
 aerial perspective, 3
 conflict of, 17, 18, 21, 53
 extrastereoscopic, 3-4, 23, 56
 interposition, 3, 18, 21
 light and shade, 3
 monocular, 3-4, 23, 56
 motion parallax, 4
 perspective, 4, 19, 20, 21-23, 24
 relative size, 3, 23
 stereopsis. *See* Stereoscopic cues
 textural gradient, 3
 Depth cuing, 4
 Depth range, 56
 Dialog boxes, 51, 53
 Digital Equipment Corporation (DEC), 30
 Discomfort when viewing stereoscopic images, 11, 16, 17, 35
 Disparity. *See* Retinal disparity
 Distance (from camera to object photographed). *See* d_m and d_o

 Distance (from eyes to screen), 42-49
 Distance (viewer), 19
 Distortion, 35, 36, 37, 38
 Divergence, 10
 Divergent parallax, 10
 d_m (maximum distance of object from camera), 37-38
 d_o (distance of object from camera), 36, 37-38
 DOS games. *See* Games
 DVD applications, 33

Electro-optical shutters, 12, 27, 32
 Emitter, 27, 29, 30, 33
 Euclid, 7
 Extrastereoscopic depth cues. *See* Depth cues
 Eye muscles, 11, 16, 35
 Eyestrain. *See* Discomfort
 Eyewear, 27-28, 32, 52
 CrystalEyes. *See* CrystalEyes display and eyewear
 electro-optical shutters in, 12
 passive, 27

 f_c (focal length of lens), 19, 37-38
 Field-of-view angle (Θ), 45-46
 Field-sequential technique. *See* Time-multiplexed technique
 Filters, neutral-density, 28
 Fire (photographing in stereo), 23
 Flicker, 28, 54
 Focal length of lens. *See* f_c
 Focusing of the eyes, 11
 Fonts, 53
 Formats for stereo, 27-33
Foundations of the Stereoscopic Cinema (Lipton), 26
 Fusion, 8, 16, 56

 Games, 18, 32
 Ghosting. *See* Crosstalk
 Glasses. *See* Eyewear
 Glitter (of the stereoscopic image), 23
 Graphics applications, 30, 31, 48-49

 Head-mounted display (HMD), 28, 33
 Hewlett-Packard (HP), 30
 HIT. *See* Horizontal image translation
 HMD. *See* Head-mounted display
 Homologous image points, 9, 56
 Horizontal image translation (HIT), 20-21, 23-24, 35, 36-37, 38-39, 56, 58
 Horizontal parallax, 10-11

 IBM computers, 30, 32
 Icons, 53
 Image magnification. *See* M
 Image selection techniques, 27
Independent Filmmaking (Lipton), 26
 Infrared emitter. *See* Emitter

Initialization, 51, 53
 Inoue, T., 26
 Interaxial distance (t_c), 10, 12, 23-24, 25, 36, 37-38, 39, 42-48, 56, 57
 Interaxial separation. *See* Interaxial distance
 Interface elements, 51, 53
 Interlace technique. *See* Time-multiplexed technique
 Interlace video mode, 52, 53
 Interline stereoscopic format, 52, 54
 Interocular distance (t), 9, 10, 19, 56, 57
 Interposition as a depth cue. *See* Depth cues
 Interpupillary distance. *See* Interocular distance
 Introspection, 7-8

Kaufman, Lloyd, 5, 26
 Kepler, Johannes, 7

Laser disc applications, 33
 LC shutters, 27
 Leakage. *See* Crosstalk
 Left and right image fields, congruence of, 10-11
 Left image field, 10-11
 Lens axes, 10, 36
 parallel, 18, 19
 Lenses
 focal length. *See* f_c
 long-focal-length (telephoto), 19, 36
 short-focal-length (wide-angle), 19, 24, 36
 Lenticular stereoscope, 8
 Leonardo da Vinci, 7
 Light and shade as depth cues, 3
 Lipton, Lenny, 14, 26, 40
 Liquid crystal shutters, 27
 Luster (of the stereoscopic image), 23

M (image magnification), 19, 37-38
 MacAdam, D.L., 5, 40
 Menus, 51, 53
 Meyer, Lhary, 40
 Mirror stereoscope, 8
 Molecular modeling, 19, 28
 Monocular depth cues. *See* Depth cues
 Motion parallax as a depth cue, 4

Negative (crossed) parallax. *See* Parallax
 NTSC equipment and specifications, 29, 31, 33

OEM products, 33
 Ohzu, H., 26
 Onyx (SGI computers), 30
 Orthographic projection, 42
 Orthostereoscopy, 18-20, 24

P (distance apart of right- and left-eye projections), 44
 Page-swapping mode, 32, 52, 54
 PAL equipment and specifications, 31, 33
 Parallax, 8-10, 11, 12, 13, 16, 17, 24, 56, 57
 angle (β), 44-47
 divergent parallax, 10
 horizontal parallax, 10-11
 motion parallax as a depth cue, 4
 negative (crossed) parallax, 10, 13, 17-18, 21, 22, 25, 42, 44
 P_m or P_{max} (maximum value of parallax), 37-38, 46
 positive (uncrossed) parallax, 10, 13, 16, 21, 22, 25, 44
 vertical parallax, 35, 37, 38
 zero parallax, 9, 16, 20-21, 36, 39, 44, 45
 Parallel lens axes. *See* Lens axes
 Parallel lens axes algorithm, 20, 25, 35-40
 Passive eyewear, 27
 PCs, 29, 31-32, 33
 Personal computers. *See* PCs
 Perspective as a depth cue. *See* Depth cues
 Perspective projections, 42-48
Philosophical Transactions of the Royal Society of London , 13
 Phosphor afterglow. *See* CRT phosphor afterglow
 Photogrammetry, 19, 39
 Pixels, 30
 Plane of convergence, 20
 Plano-stereoscopic image, 2, 11, 57
 Playback Controller, 31, 33
 P_m or P_{max} (maximum value of parallax). *See* Parallax
 Polarizing eyewear, 27
 Positive (uncrossed) parallax. *See* Parallax
 Pseudoscopic stereo images, 53

RAM, 30
Raster, 29, 31
Real-time viewing, 31
Relative size as a depth cue, 3, 23
Retinal disparity, 7, 9, 12, 22, 23, 55
Right image field, 10-11
Rotation, for creation of stereopairs, 25, 35, 37, 38

Scintillation (of the stereoscopic image), 23
Screen surround, 17-18, 21, 22, 25, 57
Scroll bars, 51, 53
Selection devices. *See* Eyewear
Selection techniques for stereoscopic images, 27
Shade and light as depth cues, 3
Shutters. *See* Electro-optical shutters
SID, 40
Side-by-side format, 27, 31, 33, 57
Sight and Mind: An Introduction to Visual Perception (Kaufman),
5, 26
Silicon Graphics (SGI), 30
SimulEyes VR product, 2, 8-9, 27, 32, 33, 52-54
SMPTE Journal, 5, 40
Society for Information Display, 40
Society of Motion Picture and Television Engineers. *See* SMPTE
Society of Photo-optical and Instrumentation Engineers. *See* SPIE
Software, stereoscopic
 creating, 41-54
 standard for, 23
Software Development Kit, 53
Sparkle (of the stereoscopic image), 23
SPIE, 14, 26
Spottiswoode, Nigel, 26
Spottiswoode, Raymond, 26
Stereo pairs, creating, 42-49
Stereo-ready computers, 30, 33
Stereo window, 17-18, 20, 39, 58
Stereoplexing, 57
Stereopsis, 7, 8, 9, 15, 57
Stereopsis depth cue. *See* Stereoscopic cues
Stereoscope, 7, 8, 11, 12, 27, 57
Stereoscopic cues, 8, 21, 22, 23
Stereoscopic images, defined, 1, 2
Stereoscopy (Valyus), 14
Sun Microsystems, 30
Surround. *See* Screen surround

t (interocular distance). *See* Interocular distance
 t_c (interaxial distance). *See* Interaxial distance
Telephoto lenses. *See* Lenses, long-focal-length (telephoto)
Textural gradient as a depth cue, 3
The Theory of Stereoscopic Transmission... (Spottiswoode & Spottiswoode), 26
Theta (Θ) (field-of-view angle), 45-46
Three-dimensional databases, 2
Time-multiplexed technique for formatting stereoscopic images, 9, 27, 28, 33, 52, 56

V (viewing distance), 19
Valyus, N.A., 14
VCR recording, 33
Vertical parallax. *See* Parallax
View-Master viewer, 2, 8, 27
View/Record box, 31
Viewer space, 17, 20-22, 39, 58
Virtual reality, 19

Wheatstone, Charles, 13
 stereoscope, 7
White-Line-Code (WLC) system, 31-32, 33, 52-53, 58
Wide-angle lenses. *See* Lenses, short-focal-length (wide-angle)
Window. *See* Stereo window
Windowed application environment, 32, 52-53
Workstations, 16, 19, 28, 29, 33

z -axis and z coordinates, 36, 37, 39, 42, 45-49, 51
Zero parallax. *See* Parallax
Zero parallax setting (ZPS), 9, 11, 12, 13, 20-21, 23, 24, 25, 35, 36-37, 38-39, 44-49, 56, 58
ZScreen, 27
Zooming, 23
ZPS. *See* Zero parallax setting