Displaying 3D Images: Algorithms for Single-Image Random-Dot Stereograms

Harold W. Thimbleby, University of Stirling

Stuart Inglis and Ian H. Witten, University of Waikato

A new, simple, and symmetric algorithm can be implemented that results in higher levels of detail in solid objects than previously possible with autostereograms.

We perceive depth in various ways: perspective, shadows, color and intensity effects, hazing, and changes in size. Artists know these effects well and can manipulate them in pictures to give an impression of depth — indeed, pictures of solid objects drawn without perspective look odd.

In three dimensions, additional methods are available for perceiving depth. Because our eyes are separated, each has a slightly different view of a distant object, and the angular difference varies with distance. When both eyes look at an object, they turn toward it (convergence) and the retinas receive slightly different images (stereo disparity). Also, the eyes focus differently for objects at different distances.

The fusing of the two eyes' images to produce the effect of depth is called *stereopsis*. Stereopsis requires the cooperation of both eyes to construct a single mental picture. This article discusses how to use stereopsis to re-create the appearance of depth from a purely flat picture. To understand how this is done, we need only consider the geometry of stereo vision, not its physiological or psychological basis.

Stereopsis occurs naturally when viewing solid objects with both eyes. It is possible with two-dimensional pictures only when each eye receives a separate image, corresponding to the view it would have of the actual 3D object depicted. There are various ways to produce stereo images; special equipment is normally required, either for making the image or for viewing it. In a stereoscope,[1,2] an optical instrument similar to binoculars, each eye views a different picture and thereby receives the specific image that would have arisen naturally. An early suggestion for a color stereo computer display involved a rotating filter wheel held in front of the eyes.[3]

In contrast, this article describes a method for viewing on paper or on an ordinary computer screen without special equipment, although it is limited to the display of 3D monochromatic objects. (The image can be colored, say, for artistic reasons, but the method we describe does not allow colors to be allocated in a way that corresponds to an arbitrary coloring of the solid object depicted.) The image can easily be constructed by computer from any 3D scene or solid object description.

Background and overview

Julesz and Miller[4] were the first to show clearly that a sense of depth could arise purely from stereopsis, without relying on other cues such as perspective or contours. They used random patterns of dots that were meaningless when viewed by a single eye but nevertheless created an impression of depth when viewed in a stereoscope.

It might seem that stereopsis necessarily requires two separate pictures, or at least some method of splitting a single picture to give each eye a separate view (using, for instance, red/green filters, polarized light, or interference, as in holograms). Recently, however, Tyler and Clarke[5] realized that a pair of random-dot stereograms can be combined, resulting in a "single-image random-dot stereogram" or, more generally, an autostereogram. Essentially, this is an overlay of the two separate random-dot patterns, with the dots carefully placed so that each serves simultaneously for two parts of the image. All that is necessary is to suitably constrain the dot pattern. Ordinary stereograms, such as photographs or wireframe models, cannot be combined into a single image, since the result is merely a double picture. (An effect called the "wallpaper illusion" occurs when lines of horizontally repeating patterns are perceived to lie at different depths; however, since the patterns repeat monotonously in wallpaper, they convey no useful information.)

It turns out that very convincing images with a vivid effect of depth can be constructed in this way, and the advantage of this ingenious approach is that no special viewing equipment is required. It takes a little practice to see depth in the pictures, but once experienced it is very satisfying.

Tyler and Clarke[5] described a simple but asymmetric algorithm, which meant that some people could see the intended effect only when the picture was held upside down. This article presents a new, simple, and symmetric algorithm for generating single-image stereograms from any solid model. It also avoids several minor problems with the original algorithm. Further background on stereograms and the development of our work is available in Thimbleby and Neesham.[6]

The literature on the psychology of stereopsis is vast. For example, Marr and Poggio[7,8] discuss computational models

Figure 1. A solid object can be simulated on a 2D image plane for stereoscopic viewing.

of the visual processes involved in interpreting random-dot stereograms. Gulick and Lawson[9] offer an excellent general survey of the psychological processes and history of stereopsis. Such background is useful, though it does not bear directly on the technique described in this article.

Stereo vision and autostereograms

Figure 1 shows an image plane placed between the eyes and a solid object. Imagine that it is a sheet of glass. (In fact, it could be a computer screen or a piece of paper.) Light rays coming from the object pass through the image plane and enter each eye. As far as the eyes are concerned, two rays pass through each point in the image plane, one for each eye. If both rays are the same color and intensity, they can be conveniently reproduced by a single light source in the image plane. Hence, although the object can be seen stereoscopically, there need be only one image in the image plane, not two, and it can be shared by both eyes. This solves the problem of seeing a stereoscopic picture without any special equipment.

The problem of generating the autostereogram amounts to illuminating the screen in such a way that it simulates a pattern of light that could have come from a solid object lying behind the screen. In general, each point of the object will map into two points on the image plane. Now, if two locations on the solid object are chosen carefully, as shown in Figure 1, and both are black dots, then it can be arranged that they generate just

three black images on the plane, two of the images coinciding. Notice that the distance between each pair of dots is different: The farther the corresponding point on the object is behind the image plane, the farther apart are its two image points.

The central dot shown on the image plane in the figure represents two separate locations on the object. Therefore, these two locations must have the same color. In turn, then, the other two dots shown in the image plane must be the same color. Overall, of course, some dots must be of different colors, or else the image plane would appear uniform and not present any useful information about the object lying behind it. Such considerations constrain the surface coloring of the object. It is sufficient to use only two colors (for example, black and white), but gray levels and a full spread of colors can be used too. There is considerable flexibility in choosing the palette, and colors may be chosen for artistic impact.

Figure 1 also illustrates the task of viewing autostereograms, of seeing depth in the initially meaningless arrangement of random dots. Suppose the image plane is transferred to a sheet of opaque paper. If the eyes converge to view the paper in the normal way, they are not converged in a way that allows reconstruction of the solid image. The same effect occurs when you look at a mark on a window: Objects behind the window appear double. In the case of random-dot stereograms, seeing the solid image "double" is tantamount to not seeing it at all. To view it correctly, one must deliberately deconverge the eyes, as explained in the sidebar "Viewing a single-image random-dot stereogram."

An algorithm for generating single-image random-dot stereograms

In Figure 1 we can see that constraints affect only points along a line that lies in the same plane as the two eyes. This provides a clue to making the algorithm efficient: It can construct the image line by line. The inevitable disadvantage is that there are no constraints at any other angle, and therefore the stereo effect is achievable only when the picture is upright (or upside down). Tyler and Clarke[5] briefly discuss the possibility of having orthogonal lines of constraints, but we will not pursue it here.

Our algorithm is based on the geometry shown in Figure 2. The object to be portrayed lies between two planes called the near and far planes. The latter is chosen to be the same distance D behind the image plane as the eyes are in front of it. This is a convenient value because when viewing the autostereogram the eyes should converge on the far plane, and you may be able to catch your reflection in the screen and use this to assist the convergence process. Since at first you don't know what you are looking for, it helps if the eyes can initially converge on a large equidistant target. The near plane is a distance μD in front of the far plane, and in the program μ is chosen to be 1/3. The separation between the near and far planes determines the depth of field (not the depth of focus, for all dots actually lie on the image plane and both eyes should focus there). Increasing the depth of field by increasing μ brings the near plane closer to the screen and causes greater difficulty in attaining proper convergence of the eyes.

It is convenient to define the "image stereo separation" of a point on the surface of the solid object viewed in 3D to be the distance between its image points lying in the image plane. This quantity relates directly to the conventional measure of stereo disparity, which is the difference in angle subtended at the eyes. Since the background is taken to be the same distance behind the screen as the screen is from the eyes, the separation for the far plane is half the distance between the eyes.

Figure 2 shows the stereo separation s for a point with a specified z-coordinate. The z-values range from 0 (which corre-

Viewing a single-image random-dot stereogram

Figure A1 is a random-dot stereogram of a large annulus, exactly the same size and texture as that of Figure A2, floating above the middle of the picture. The challenge of successful viewing is that the viewer must decouple the eyes' convergence from their focusing. Since this is never done in normal circumstances (except perhaps when tired), it may be difficult to achieve a successful 3D effect at first. The trick is to converge the eyes at a distance well beyond the paper on which the dots lie, yet focus them on the dots themselves. Working in bright light can help, as contraction of the iris causes the eyes' depth of field to increase, and as a result they rely less on focusing.

Altering the eyes' convergence is not easy, but the two circular marks near the bottom of the picture can help. By blinking (or covering) the eyes alternately, you can make the marks produce four images, two with the right eye and two with the left. Now deconverge your vision so that all four images can be seen with both eyes open. The aim is to move the two pairs of images closer together until the inner two coincide and only three distinct marks are seen. The center mark is binocular, while the outer two are seen monocularly (with only one eye each). Fixate on the center mark to stabilize the image, then carefully allow your gaze to wander around the whole picture. Placing the convergence marks in the center of the picture instead of near the bottom can make for easier viewing but usually spoils the picture.

When viewing the autostereogram, you should strive to converge your eyes on the far plane. In Figure A1 this is the same distance beyond the paper as the paper is from the eyes. Copying the stereogram onto a transparency greatly aids viewing. Mark an X on a plain sheet of paper and hold the transparency about halfway between your eyes and the paper. View the paper through the transparency and converge your eyes on the X. Carefully allow your gaze to wander across the whole picture. Once you have some experience, you can dispense with the transparency and work directly from the printed image.

Do not be content with an inferior 3D effect. You may become dimly aware of some circularity in the dot pattern, along with some tenuous 3D effects, but this is not the full illusion. Viewed properly, the annulus should leap out of the plane at you, and the effect should be stable and persist through sustained and detailed examination of each part of the image, both foreground and background. Stereopsis has hysteresis: Once you have begun to fuse an image, it becomes clearer as you look. The farther your eyes are from the paper, the greater the depth effect. If you see two interlocking rings instead of a single ring, your eyes have converged incorrectly, at twice the distance apart that they should be. (This and other artifacts are described in the article.) Reexamine the convergence dots and strive to make the four images into three. If you see the annulus behind the plane rather than in front of it, you are viewing the stereogram with your eyes crossed, that is, converging on a point in *front* of the paper. When properly diverged, the eyes converge on a point *behind* the paper. (The latter is sometimes referred to as "wall-eyed" or "boss-eyed.") Some people can cross their eyes or diverge them at will; others can only diverge them.

Even if you can capture the 3D picture without using the convergence marks, it is worth learning to use them properly. If you wear spectacles to correct for short-sightedness, it is advantageous to remove them: Your natural focus will be closer than the patterns appear to be when the eyes are correctly converged. If your spectacles also correct for other sight defects, such as astigmatism, removing them may not help.

Allow plenty of time: It is often necessary to work for some minutes to achieve the full 3D effect for the first time. Viewing becomes much easier with practice.

If viewing difficulty persists, photocopy two identical transparencies from Figure A1, line them up carefully, and slide one horizontally over the other until the annulus appears. Although the picture is not in stereo, this will at least serve to convince your eyes that there *is* an annulus hidden in the dots. Two identical transparencies are recommended, rather than one and the original, since the photocopier will tend to stretch the transparencies and the annulus may only be visible if both images are distorted similarly.
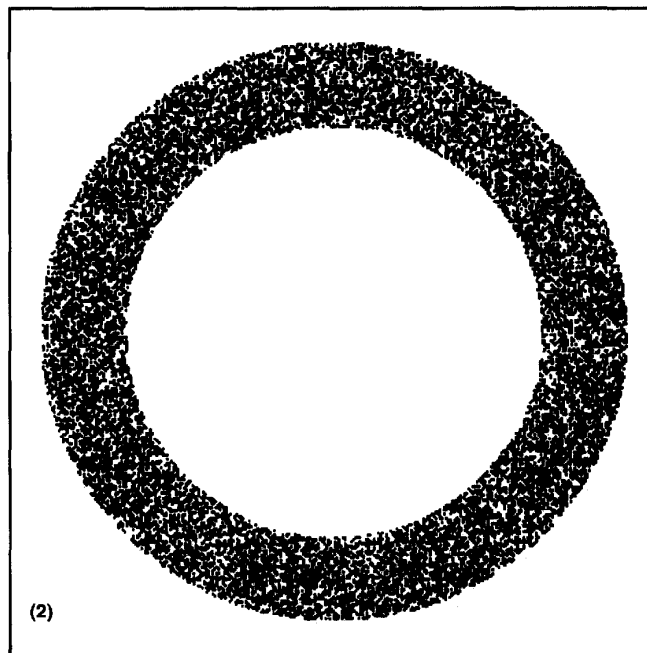
Figure A. (1) Stereogram showing an annulus floating above a surface; (2) the annulus by itself.

Figure 2. The geometry forming the basis for the stereogram-generation algorithm.

sponds to the far plane) to 1 (which corresponds to the near plane). Thus, the point is a distance $\mu z D$ in front of the far plane, or $(1 - \mu z)D$ from the image plane. By similar triangles,

$$s = \frac{1 - \mu z}{2 - \mu z} E$$

which gives the stereo separation s as a function of z. This is the fundamental relationship on which the algorithm is built.

Implementation of the algorithm. The algorithm, programmed in C, is shown in Figure 3. Lines 1-11 set the scene. The screen is $maxX$ by $maxY$ pixels, and the object's z-value is $Z[x][y]$. The depth of field μ is chosen to be 1/3. Neither the eye separation nor the screen resolution is critical; approximate values will do in both cases. The image stereo separation corresponding to the far plane is called *far*.

The program processes one scan line at a time (using the large loop in lines 16-58 of Figure 3). The key to solving the constraint equations is to record what the constraints are in an initial pass (lines 26-52). A second pass allocates a random pixel value (black or white) whenever there is a free choice (line 54); when there is not, the pixel is set to the value that obeys the relevant constraint (line 55).

Pixel constraints are specified by the

same[] array. In general, each pixel may be constrained to be the same color (black or white) as several others. However, it can be arranged that each element of the array need only specify a single constraint on each pixel. The *same*[] array is initialized by setting *same*[x] = x for every pixel, representing the fact that in the absence of any depth information, each pixel is necessarily constrained to be the same as itself (line 24).

Then the picture is scanned, giving, at point x, a separation s between a pair of equal pixels that corresponds to the image stereo separation at that point. Calling these two pixels *left* and *right*, then *left* is at $x - s/2$ and *right* is at $x + s/2$; but just in case s is odd, we can more accurately position *right* at *left* + s (lines 28-29). To avoid a systematic bias in rounding off when s is odd, line 28 adds 1 to s on odd-numbered lines (we might instead have made the offset random rather than in alternate directions on alternate lines, as suggested here). The pair of pixels, *left* and *right*, must be constrained to have the same color. This is accomplished by recording the fact in the *same*[] array. (However, there may be geometric reasons why the corresponding point on the solid object is not visible along both lines of sight; this is checked in lines 35-39, and we discuss it in detail below.)

To ensure that the pixel at *left* is recorded as being the same as the pixel at

right, should we set *same*[left] = *right*, *same*[right] = *left*, or both? Actually, it is unnecessary to set both. When the time comes to draw the pixels, the line will be scanned either left to right or right to left, and in the latter case it is only necessary to ensure that *same*[left] = *right*. It doesn't matter which of the two directions is used. We choose right to left (line 53), set *same*[left] = *right*, and require, as an invariant, that *same*[x] ≤ y whenever x ≤ y.

Now, *same*[left] may have already been set in the course of processing a previous constraint. Recall that *same*[x] records a pixel to the right of x that is constrained to be the same color as x. If *same*[left] is already constrained, that constraint is followed rightward (lines 42-48, using variable k to follow the same links) to find a pixel that is not otherwise constrained. In following the constraints, the variable k may "jump" over *right*. If this happens (line 43), lines 46-47 preserve the assumption that *left* < *right* by swapping them.

Once the constraints have been set for all pixels in a scan line, the line is rescanned in decreasing x order, that is, from right to left as mentioned above, allocating pixel values (lines 53-57). When a pixel is unconstrained (*same*[x] = x), a value is chosen for it randomly and stored in *pix*[x]. Otherwise, its value must be constrained to be the same as some pixel farther to the right, and so *pix*[x] is set to *pix*[same[x]]. The routine *Set_Pixel*()

```
1.  /*    Algorithm for drawing an autostereogram                      */
2.
3.  #define round(X) (int)((X)+0.5) /* Often need to round rather than truncate  */
4.  #define DPI 72           /* Output device has 72 pixels per inch      */
5.  #define E round(2.5*DPI)  /* Eye separation is assumed to be 2.5 inches  */
6.  #define mu (1/3.0)        /* Depth of field (fraction of viewing distance)  */
7.  #define separation(Z) round((1-mu*Z)*E/(2-mu*Z))
8.                           /* Stereo separation corresponding to position Z   */
9.  #define far separation(0) /* ... and corresponding to far plane, Z=0   */
10. #define maxX 256         /* Image and object are both maxX by maxY pixels  */
11. #define maxY 256
12.
13. void DrawAutoStereogram(float Z[][])
14. {                        /* Object's depth is Z[x][y] (between 0 and 1)   */
15.    int x, y;             /* Coordinates of the current point     */
16.    for( y = 0; y < maxY; y++)      /* Convert each scan line independently  */
17.    { int pix[maxX];      /* Color of this pixel        */
18.        int same[maxX];   /* Points to a pixel to the right ...     */
19.                          /*... that is constrained to be this color    */
20.        int s;            /* Stereo separation at this (x,y) point    */
21.        int left, right;  /* X-values corresponding to left and right eyes  */
22.
23.        for( x = 0; x < maxX; x++ )
24.            same[x] = x;  /* Each pixel is initially linked with itself   */
25.
26.        for( x = 0; x < maxX; x++ )
27.        { s = separation(Z[x][y]);
28.            left = x - (s+(s&y&1))/2;        /* Pixels at left and right ...  */
29.            right = left + s;                /* ... must be the same ...    */
30.            if( 0 <= left && right < maxX )  /* ... or must they?      */
31.            { int visible;/* First, perform hidden-surface removal      */
32.                int t = 1;  /* We will check the points (x-t,y) and (x+t,y)  */
33.                float zt;   /* Z-coord of ray at these two points      */
34.
35.                do
36.                { zt = Z[x][y] + 2*(2 - mu*Z[x][y])*t/(mu*E);
37.                    visible = Z[x-t][y]<zt && Z[x+t][y]<zt; /* False if obscured  */
38.                    t++;
39.                } while( visible && zt < 1 );  /* Done hidden-surface removal ...  */
40.                if( visible )  /* ... so record the fact that pixels at    */
41.                { int k;       /* ... left and right are the same      */
42.                    for( k = same[left]; k != left && k != right; k = same[left] )
43.                        if( k < right )/* But first, juggle the pointers ...  */
44.                            left = k; /* ... until either same[left]=left   */
45.                        else          /* ... or same[right]=left       */
46.                        {   left = right;
47.                            right = k;
48.                        }
49.                    same[left] = right;/* This is where we actually record it  */
50.                }
51.            }
52.        }
53.        for( x = maxX-1; x >= 0; x- )  /* Now set the pixels on this scan line  */
54.        { if( same[x] == x ) pix[x] = random()&1; /* Free choice; do it randomly */
55.            else pix[x] = pix[same[x]]; /* Constrained choice; obey constraint   */
56.            Set_Pixel(x, y, pix[x]);
57.        }
58.    }
59.    DrawCircle(maxX/2-far/2, maxY*19/20); /* Draw convergence dots at far plane, */
60.    DrawCircle(maxX/2+far/2, maxY*19/20); /* near the bottom of the screen     */
61. }
```

Figure 3. C-language algorithm for drawing an autostereogram.

then sets the pixel to black or white on the screen as soon as its color is known. (If desired, the number of calls to *Set_Pixel()* can be halved by first drawing a white line across the scan line and then setting only the black pixels.)

The code is followed by drawing two circular marks with their centers near the bottom of the screen, separated by a distance proportional to *far* (lines 59-60). This separation places them on the background of the 3D picture. Some viewers find it easier to see the stereograms by converging their eyes on the near plane, and this is facilitated by separating the marks by *separation*(1) instead of *far* = *separation*(0). Alternatively, if the marks obscure part of the object, it may be better to determine the distance from the depth of the object at that point, to ensure that the eyes' separation is correct for that part of the picture. Initial viewing may be slightly easier if the marks are placed at the center of the screen, though for aesthetic reasons we prefer to place them near the bottom.

Hidden-surface removal. Sometimes, because of a transition in the object (from the near to the far plane, say), a surface in the foreground obscures one eye's view of a more distant point. Hidden-surface removal is a technical detail, and very few objects make it visibly worthwhile. Yet the advantage is that for any part of the object that is strictly hidden, there is one less constraint to process. In turn, this gives the algorithm greater flexibility in allocating pixel colors, which reduces the problem of artifactual "echoes," as discussed in the next section. The hidden-surface-removal code is in lines 35-39 of Figure 3.

Figure 4 shows a point on the object that would cause two pixels of the image to be linked were it not for an obscuring object that interrupts one eye's view. The crucial inequality for hidden-surface removal is that $z_1 > z_t$, where z_1 is the z-coordinate of a point on the obscuring object and z_t is the z-coordinate of a point on the ray from the eye to the original object. The x-coordinate is governed by the distance t, and if such an interruption occurs for any value of t greater than 0 and up to the point where $z_t - 1$, then the original point is no longer visible to this eye. The small shaded triangle is similar

to the much larger one with the same apex but is based on half of the line joining the eyes, and so

$$\frac{t}{(z_t - z_0)\mu D} = \frac{E/2}{(2 - \mu z_0)D}$$

(the relevant distances in the denominators are marked in Figure 4), from which it follows that

$$z_t = z_0 + \frac{2(2 - \mu z_0)t}{\mu E}$$

This is calculated in line 36 of the algorithm in Figure 3. The crucial inequality $z_1 \geq z_t$ appears in line 37, and at the same time a test is made to determine whether the other eye's view is obscured. The whole operation is repeated for different values of t, from 1 up to the maximum possible value, which occurs when $z_t = 1$.

This algorithm checks for obscuring surfaces at each point of the image plane, but if a geometric object were being rendered rather than the $Z[x][y]$ array, a more efficient implementation could be devised using surface intersections.


Figure 4. The geometry for hidden-surface removal, which is performed by the algorithm.

Previous methods of generating autostereograms work by creating a random vertical strip of dots and repeating the pattern independently along each horizontal scan line, modulating the pattern by the depth of the object at that point. We will describe this process more precisely using a quantitative example, but first we need some terminology.

Suppose the object comprises a foreground called the near plane and a background called the far plane; both lie behind the image screen. In Figure A1, for example, the annulus is the near plane and the surface it floats over is the far plane, while the screen corresponds to the piece of paper on which the figure is drawn. Consider the two rays from a point on the far plane to the eyes, shown in Figure 2. These intersect the screen at two pixels; the distance between these pixels is called the stereo separation corresponding to the far plane. If the background is the same distance behind the screen as the screen is from the eyes, then this stereo separation is half the distance between the eyes, say 1.25 inches (half the distance typically separating the eyes).

The width of the initial strip should be the stereo separation corresponding to the far plane. Suppose this is 90 pixels, corresponding to approximately half the distance between the eyes, measured in pixels on a typical display screen with 72 pixels/inch. To give a depth impression of the far plane, the first dot would be repeated at position 91, the second at position 92, and so on. At the point where the foreground object begins, the dots are repeated at a smaller lag, say 72 pixels instead of 90, corresponding to the stereo separation of the near plane. To portray an object that lies between the near and far planes, an appropriate interpolation between 72 and 90 pixels would be used for the lag.

This is a simple but technically incorrect algorithm for generating an autostereogram. Nevertheless, it generally produces reasonable results.

Two problems are obvious. First, when a depth change requires the repetition distance to increase, what should intermediate pixels be set to? Second, copying the initial randomly selected strip of dots gives a directional bias that belies the symmetry of the actual physical situation, and some people have reported difficulty viewing such images.

References

1. C.W. Tyler and M.B. Clarke, "The Autostereogram," *SPIE Stereoscopic Displays and Applications*, Vol. 1,258, 1990, pp. 182-196.

2. G.S. Slinker and R.P. Burton, "The Generation of Random Dot and Line Autostereograms," *J. Imaging Science and Technology*, Vol. 36, No. 3, May/June 1992, pp. 260-267.

Geometrical limitations

Although there are some geometrical limitations to the approach, in practice none actually detract from the depth effect.

Geometrical distortion. As Figure 2 illustrates, the object being viewed will have a slight lateral distortion in the x direction. Suppose a value of x has been reached (Figure 3, line 26) that corresponds with point a on the image plane. To determine the stereo separation, the program will look along the line aA to find the appropriate depth, which in this case locates a point (marked A) on the far plane (line 27). However, the stereo separation should really be governed by point A' instead. There might be something intervening between the eyes and A', but not between the eyes and A. While it would be possible to modify the program to correct the problem, this could only be done for a fixed viewing position.

There is an analogous distortion in the vertical (y-axis) direction, perpendicular

to the plane containing the eyes. Hidden-surface decisions are made by projecting backward from each scan line in a direction orthogonal to the viewing plane, and this is accurate only when the object subtends a small vertical angle at the eyes. Neither the horizontal nor the vertical distortion affects the depth geometry, and neither is noticeable at normal viewing distances. For distant viewing, both errors disappear (the lateral one disappears because points A and A' of Figure 2 coincide).

The amount of available depth resolution is the same as the distance *far* − *near* + 1, in pixels, where *far* and *near* are the stereo separations corresponding to the far and near planes, that is, *separation*(0) and *separation*(1) in the program. This is because points on the far plane are separated by *far* pixels on the screen, points on the near plane are separated by *near* pixels, and intermediate values are interpolated. For the parameters in the program, *far* = 90 and *near* = 72, giving a total of 19 separate depth planes. The number of depth planes increases in proportion with the output device resolution.

The depth planes are not exactly the same distance apart because the stereo

separation is not a linear function of z. Previous schemes for calculating single-image random-dot stereograms (see the sidebar above) interpolate the z-values linearly into the space between the near and far planes, creating a small depth distortion. The algorithm in Figure 3 corrects this problem by using the true nonlinear separation function (line 7).

False fusion: echoes. Suppose there are constraints on pixels such that $a = b$ and $b = c$. It follows that $a = c$. This last constraint, interpreted geometrically, corresponds to a third point in 3D space, and if pixel b lies between pixels a and c, the new point lies farther from the eyes than the points constraining $a = b$ or $b = c$. There may be no such point on the object! This phenomenon, which follows inexorably from the transitivity of equality, can produce visible artifacts that we call echoes.

Most transitive pixel constraints do not cause a visible echo. If two pixels are constrained to be equal but are separated by more than *far*, the echo is conceptually behind the object — indeed, if the pixels are separated by more than E, the echo is im-

Figure 5. Two more views of Figure A, with echoes appearing in the same plane as the annulus: (a) for the algorithm of Figure 3 without hidden-surface removal, (b) for the previous algorithm, described in the sidebar on page 45.


Figure 6. Incorrect convergence of the eyes can lead to a false fusion effect.

possible to see anyway, since it corresponds to a virtual image behind the viewer's eyes. However, for regular objects such as plane figures, echoes may combine in complex but systematic ways right across the image. Periodicities in constraints may result in echoes that appear in the same depth plane where a legitimate part of the object appears. Such echoes can be very distracting, not to mention quite misleading when they occur in unknown or unfamiliar objects.

Fortunately, hidden-surface removal reduces constraints for plane figures to the extent that echoes in the same plane as the figure cease to be a noticeable problem. For example, Figure 5a shows the echoes that arise from the test pattern of Figure A (in the sidebar) if the hidden-surface code is suppressed. Figure 5a was generated as follows. First, a stereogram of Figure A2 was created without hidden-surface removal. Then it was scanned for pairs of pixels with the same value (black or white) that have a separation corresponding to the near plane, that is, *separation*(1). Whenever this occurred, the midpoint was marked black. Of course, it frequently

happens by chance, so the resulting image is mottled, but the visible echoes appear as the large connected black regions visible in Figure 5a. If the procedure is repeated with the hidden-surface code in place, the corresponding picture shows just the single large annulus without any echoes. Visible echoes are also produced by Tyler and Clarke's scheme,[5] described in the sidebar "Previous method for generating single-image random-dot stereograms"; these echoes appear in Figure 5b. The asymmetry of their scheme is clearly evident in this figure.

A visible echo often creates such an implausible distortion that it remains unnoticed — certainly when the object is familiar, such as a sphere or text written in relief. Nevertheless, once an echo has been noticed, it may become hard to ignore. Echoes often persist, although for some people they come and go of their own accord (like depth reversal in the Necker cube illusion). These effects probably occur because both the background plane and the echoes in the foreground are simultaneously present in the autostereogram, and the brain must choose which to see.

False fusion: artifacts. Even though there may be no geometrical constraint that $a = b$, it can happen that a and b are given the same color purely by chance. This will result in a visual artifact that may be distracting to the eye.

With a computer screen display, artifacts can be effectively eliminated by displaying different versions of the picture in rapid succession. This can be done efficiently by storing all the constraints in an array and then repeating lines 57-61, but with different random numbers. This means that artifacts appear only intermittently and are therefore hardly noticed. If the display can be recomputed fast enough, or if several versions are generated in advance and cycled through quickly, a striking "shimmering" effect is obtained. This makes it much easier to see the depth illusion because the eyes cannot latch on to chance patterns or focus on individual dots and the only constant constraints are those that represent the 3D image.

It is also possible to ameliorate artifacts by using gray levels. Rather than color the pixels black and white (half the time any two particular pixels will be the same), we can give them gray levels (or colors) from a larger palette, say eight gray levels. Then chance artifacts become less likely.

False fusion: incorrect convergence. False fusion can also occur simply because the eyes are incorrectly converged. In our experience, false fusion is most commonly caused by having the eyes converged at far too great a distance, thereby seeing the virtual solid object that is imaged by alternate image pixels in the image plane.

Figure 7. Use of the correct viewing technique will reveal a hemisphere bulging out of a plane in this stereogram.

gorithm and produced satisfactory but not ideal results (see sidebar on the next page). Our symmetric algorithm is efficient and in practice as fast as the earlier algorithm. Correct implementation of geometrical constraints, together with the hidden-surface feature, enables higher levels of detail in the solid object to be faithfully recorded in the stereogram. The "shimmering" effect that may occur on a dynamic display renders stereograms much easier to view. These features give artists greater freedom and scope for creative use of this new medium and provide the basis for many interesting experiments with optical illusions. ■

Figure 6 shows the effect: Instead of seeing the annulus in Figure A, one sees two interlocking rings. In fact, if the eyes really are separated by 2.5 inches, as assumed in the program, this effect will be achieved precisely when the gaze is directed at infinity. The risk of perceiving 3D objects with the eyes set at an incorrect convergence makes guide dots essential, especially for more complex objects.

Low-resolution autostereograms. It is possible to draw low-resolution autostereograms using letters (say, 0.5 inch in height), for example, instead of pixels. Some of the fine points of drawing autostereograms, discussed above, were based on an assumption of high resolution. With a low-resolution picture, the pixels themselves become a salient feature, and the geometrical correctness of the picture may have to be compromised to reduce distracting effects.

Hidden-surface removal can cause ambiguous edges around sudden changes in depth. It is generally best, then, not to use hidden-surface removal with low-resolution autostereograms. (With high resolution, the eye doesn't notice the approxi-

mately pixelwide ambiguity.)

Care was taken to avoid the systematic bias created by always rounding in the same direction when the stereo separation is halved (line 28 of Figure 3). However, in low-resolution autostereograms, it is better to round consistently (either up or down) to preserve straight vertical lines, even though they may be systematically misaligned.

A utostereograms are very satisfying to view and are becoming increasingly popular, particularly as recreational posters and in scientific illustrations. A public-domain X Windows demonstration program for displaying and printing autostereograms is available from ftp.cs.waikato.ac.nz. It incorporates the basic algorithm of Figure 3, along with an interface that provides control over all the features discussed in this article. As a final example, and to illustrate an object that requires continuous change in depth, Figure 7 shows a hemisphere bulging out of a plane.

Previous attempts at drawing autostereograms relied on an asymmetric al-

Acknowledgments

References

1. C. Wheatstone, "Contributions to the Physiology of Vision, Part I: On Some Remarkable, and Hitherto Unobserved, Phenomena of Binocular Vision," *Royal Soc. of London Philosophical Trans.*, Vol. 128, 1838, pp. 371-394.

2. C. Wheatstone, "Contributions to the Physiology of Vision, Part II: On Some Remarkable, and Hitherto Unobserved, Phenomena of Binocular Vision (continued)," *The London, Edinburgh, and Dublin Philosophical Magazine and J. of Science*, Series 4, Vol. 3, 1852, pp. 504-523.

3. R.I. Land and I.E. Sutherland, "Real-Time, Color, Stereo Computer Displays," *Applied Optics*, Vol. 8, No. 3, Mar. 1969, pp. 721-723.

The advantage of symmetric eye convergence

For a viewer to perceive stereoscopic depth across a range of depths without seeing double, the angular separation of corresponding parts of the retinas' images must not be too great. Specifically, the angle subtended on a retina — the stereo disparity — must be less than about 0.5 degrees. Under certain circumstances, larger or smaller disparities can be maintained. For example, a rod held on the nose and pointing away from the face subtends too high a disparity to be seen as a single rod near the face, but you know it's a single rod, and this impression suffices to sense, if not patently see, depth uniformly along its length.

The range of angles subtended on a retina depends on which direction the eye itself is pointing. Suppose one eye is converged inward at a fixed angle θ to a perpendicular to the image plane. Setting θ higher for one eye reduces it for the other, given that both eyes converge on the same point in the image plane. Hence, equal, or symmetric, eye convergence permits the largest range of depth change within any given limitation of disparity. Take the eye separation to be 2.5 inches, $\mu = 1/3$, the viewing distance to be 15 inches from the eye to the image plane, and consider a symmetric convergence angle $\theta = \arctan(E/4D)$. Then all depth changes ($0 \le z \le 1$) have a disparity of less than 0.5 degree.

The previous algorithm (described in the sidebar on page 45) assumes that one eye looks straight at the image ($\theta = 0$). For the full range of z and the same viewing conditions, this results in a huge disparity of 5.6 degrees. Put another way, for a disparity of at most 0.5 degree, the effective range of z is halved. Although this assumption makes constructing the image easy, it distorts the image asymmetrically and reduces the complexity of the images that some people can see clearly. Note that a slight rotation of the head has no serious impact on these arguments, since the eyes continue to converge at the same point (though the effective narrowing of the head reduces all angles).

In contrast, the advantages of the symmetric method we describe are twofold: The range of stereo disparity required for viewing is reduced, and parallax distortion is both reduced and symmetric. Further advantages are hidden-surface removal and correct depth interpolation, though these refinements could also be applied to the previous algorithm. Although the symmetric convergence approach places multiple constraints for both eyes on each image pixel, these constraints are readily solved.

4. B. Julesz and J.E. Miller, "Automatic Stereoscopic Presentation of Functions of Two Variables," *Bell System Technical J.*, Vol. 41, Mar. 1962, pp. 663-676.

5. C.W. Tyler and M.B. Clarke, "The Auto-stereogram," *SPIE Stereoscopic Displays and Applications*, Vol. 1,258, 1990, pp. 182-196.

6. H.W. Thimbleby and C. Neesham, "How to Play Tricks with Dots," *New Scientist*, Vol. 140, Oct. 9, Issue No. 1,894, 1993, pp. 26-29.

7. D. Marr and T. Poggio, "Cooperative Computation of Stereo Disparity," *Science*, Vol. 194, Oct. 15, 1976, pp. 283-287.

8. D. Marr and T. Poggio, "A Computational Theory of Human Stereo Vision," *Proc. Royal Soc. of London*, Vol. B204, Royal Society, London, 1979, pp. 304-328.

9. W.L. Gulick and R.B. Lawson, *Human Stereopsis: A Psychological Analysis*, Oxford University Press, Oxford, UK, 1976.

Harold W. Thimbleby is professor of computing research at Middlesex University, England. Formerly a lecturer at Queen Mary College (now Queen Mary and Westfield College) of London University and York University, England, he joined Stirling University, Scotland, in 1988 as professor of information technology and moved to Middlesex in 1994. He has written more than 200 publications, including three books, and sits on several editorial boards.

Thimbleby received a PhD from London University in 1981. He is a chartered engineer and a fellow of the Royal Society of Art. In 1981 he was awarded the British Computer Society's Wilkes Medal.

Stuart Inglis is a doctoral student at the University of Waikato, New Zealand, where he graduated with a degree in computer science. He was recently a research assistant at the University of Calgary, Canada. He is interested in all aspects of computer graphics and is pursuing research on compression-based methods for template matching and the compression of textual images.

Ian H. Witten is professor of computer science at the University of Waikato, New Zealand, having been previously affiliated with the University of Calgary, Canada, and Essex University, England. His current interests include prediction and modeling, machine learning, and programming by example. He has published widely on machine learning, speech synthesis and signal processing, text compression, hypertext, and computer typography. He has authored or coauthored five books.

Witten received degrees in mathematics, computer science, and electrical engineering from Cambridge University, the University of Calgary, and Essex University, respectively. He is a member of the IEEE Computer Society.

Thimbleby can be contacted at the Department of Psychology, University of Stirling, Stirling, Scotland, FK9 4LA, e-mail hwt@compsci.stirling.ac.uk. Inglis and Witten are at the Department of Computer Science, University of Waikato, Hamilton, New Zealand, e-mail {singlis, ihw}@waikato.ac.nz.