TR 85-007

Systems for 3D Display in Medical Imaging

Stephen M. Pizer

Department of Computer Science
University of North Carolina
Chapel Hill, North Carolina 27514

March, 1985

To appear in Proc. NATO Advanced Studies Institute on Pictorial Information Systems in Medicine, 1985.

Systems for 3D Display in Medical Imaging

Stephen M. Pizer, Ph.D.
Departments of Computer Science and Radiology
University of North Carolina
Chapel Hill, North Carolina, U.S.A.

Abstract

Three-dimensional display can be accomplished by simulating either of the two means of three-dimensional presentation in nature: reflection from object surfaces and self-luminous objects. Both forms of display benefit significantly from real-time interaction.

Reflective displays are produced by shaded graphics, which require raster graphics hardware: a frame buffer, scan-out components, and often a processor. Hardware to speed display transformations is presently becoming available. The structures of both the standard and fast systems will be described. Also briefly presented will be the design of interactive head or hand trackers for moving and pointing within a 3D image.

Self-luminous display can be produced for surface representations made by dots or lines or for space-filling grey-scale distributions. In both cases one may calculate one or more projections which may be presented at various times or to different eyes to achieve the percept of depth. Alternatively one may place all of the individual points in 3-space by appropriate optical maneuvers. Display systems for both projective and optically based 3D display will be described. The former category includes vector, point, and raster graphics systems. The latter includes holography and systems with rotating or vibrating screens or mirrors, most interestingly the varifocal mirror and rotating LED-panel systems.

1. Introduction

There are two ways in which the three-dimensional world presents itself: by reflection and self-luminosity. Similarly, there are in essence only two forms of three-dimensional display: reflective and self-luminous. Shaded graphics is the method by which reflective display is normally presented (see figure 1). Self-luminous display is provided by many approaches, including holography, vibrating or rotating mirrors and screens, and vector graphics of wire frames (see figure 2) using, for example, the kinetic depth effect.

Figure 1. Shaded graphics display of a tumor in a brain in a head, from CT scans.

Figure 2. Wire-frame display of a brain, from CT scans.

It appears [Fuchs and Pizer, 1984] that reflective display is superior to self-luminous display for the presentation of known surfaces, e.g. the surface of an organ to a surgeon or radiotherapist who will use the result for planning. However, the need to know the surfaces before the display is produced makes reflective display weak for the explorational needs of a radiologist. In contrast, self-luminous display seems to have some real strengths for explorational purposes. Therefore, in the following first reflective and then self-luminous display systems will be covered.

2. Reflective Display Systems

The hardware currently supporting reflective display is the raster graphics system, the major component of which is the *frame buffer* (see figure 3). The frame buffer is a memory holding values for each of red, green, and blue intensities for each picture element (pixel) on the display screen. Alternatively, each pixel may hold an intensity value only, or values representing hue, saturation, and intensity. Provided with the

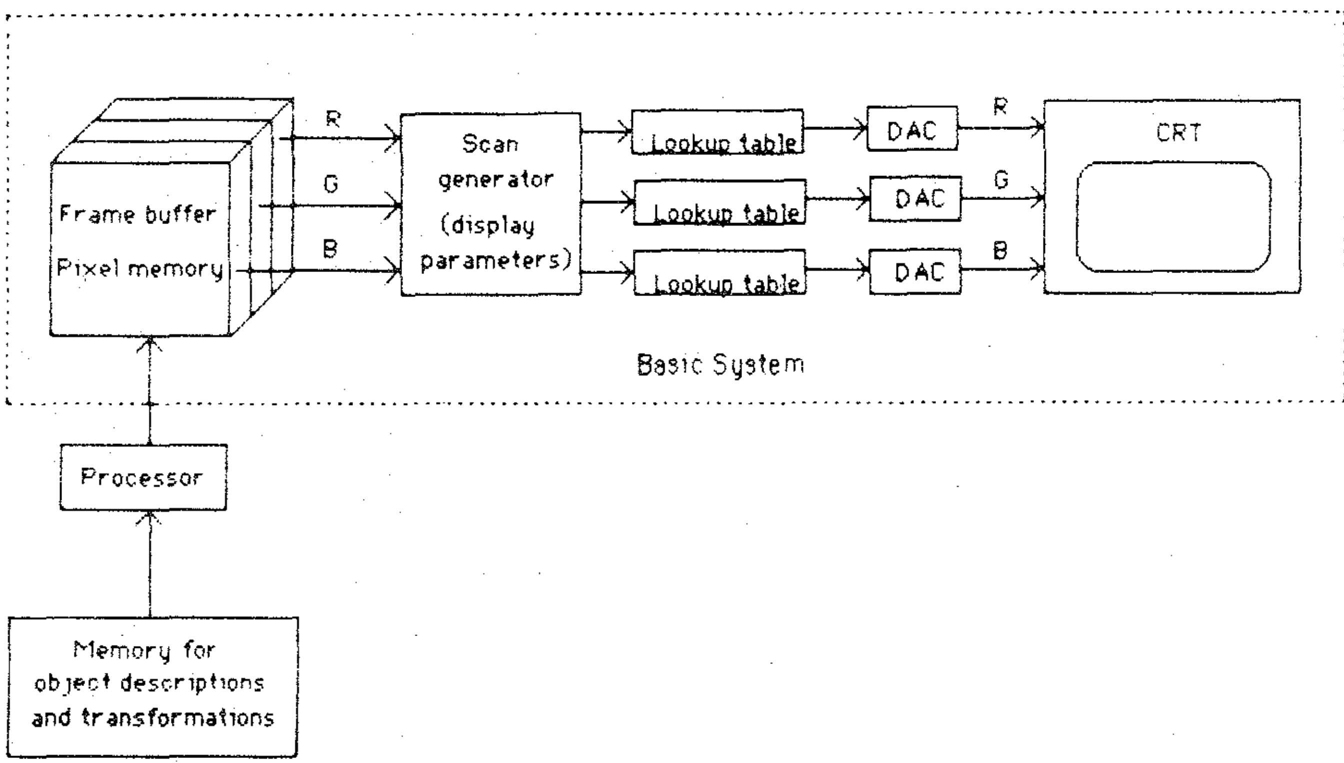

Figure 3. The structure of a raster graphics system.

frame buffer is a mechanism for scanning out these values to control the brightness of three color guns as the image is swept out onto the screen as a raster at 25-60 frames/sec. This mechanism, called the *scan generator,* includes parameters that cause the appropriate portion of memory to be read out, thus allowing pan and zoom. The three brightnesses can be modified, and thus the resulting color controlled, by passing the digital intensities through lookup tables before they are made analog.

Both effective display and image manipulation depend on interactive modification of the contents of the frame buffer and thus of the displayed image. For example, one may wish to rotate an object or move a cursor. This is accomplished by keeping descriptions of image objects in a separate memory and then filling the frame buffer by applying transformations on these objects [Newman and Sproull, 1979]. The object descriptions are normally in the form of lists of polygonal tiles. These tiles are frequently triangles fit between contours of the surface on successive slices [Fuchs, 1977] or rectangles representing the faces of object voxels (3D pixels) [Herman, 1985]. The transformations, represented by matrices and function subroutines, allow such operations as rotation, translation, scaling, perspective division, and lighting calculations. A processor, frequently a special one internal to the display system, applies the transformations, does *visibility* calculations to determine at each pixel which is the frontmost polygon at that position, and *shading* calculations to determine the light reflected from that polygon at the pixel in question, possibly smoothing the polygonal surface as it shades [Sutherland, 1973]. The result is loaded into the frame buffer.

The speed of calculation and loading of pixel values can be slow, since calculations are required at every pixel or for every tile. Speeding up this process is important in that it allows interactive control of the display parameters, and this improves the 3D perception as well as making the system convenient to use. Two hardware solutions, both involving parallel computation, are presently being developed to provide this speedup. In the first solution, processors are pipelined so that, for example, of a sequence of operations to be carried out on each tile, operation n is being carried out on tile k while operation n-1 is being carried out on tile k+1, etc. In the second solution, some processing capability is placed at each pixel so that in effect each pixel can do its own hiding and shading calculations. The latter approach is taken by the elegant design of Fuchs.

Fuchs's *Pixel-planes* structure is based on the realization that the calculations for determining 1) whether a pixel is in a particular tile, 2) which tile is hidden by which at each pixel, and 3) what shade each pixel should take, involves or can be well approximated by the calculation of a linear function Ax+By+C. Fuchs et al [Poulton, 1985] therefore have created a structure in VLSI (see figures 4 and 5) in which a tree connection of the pixels allows every pixel in the frame buffer (each with a position

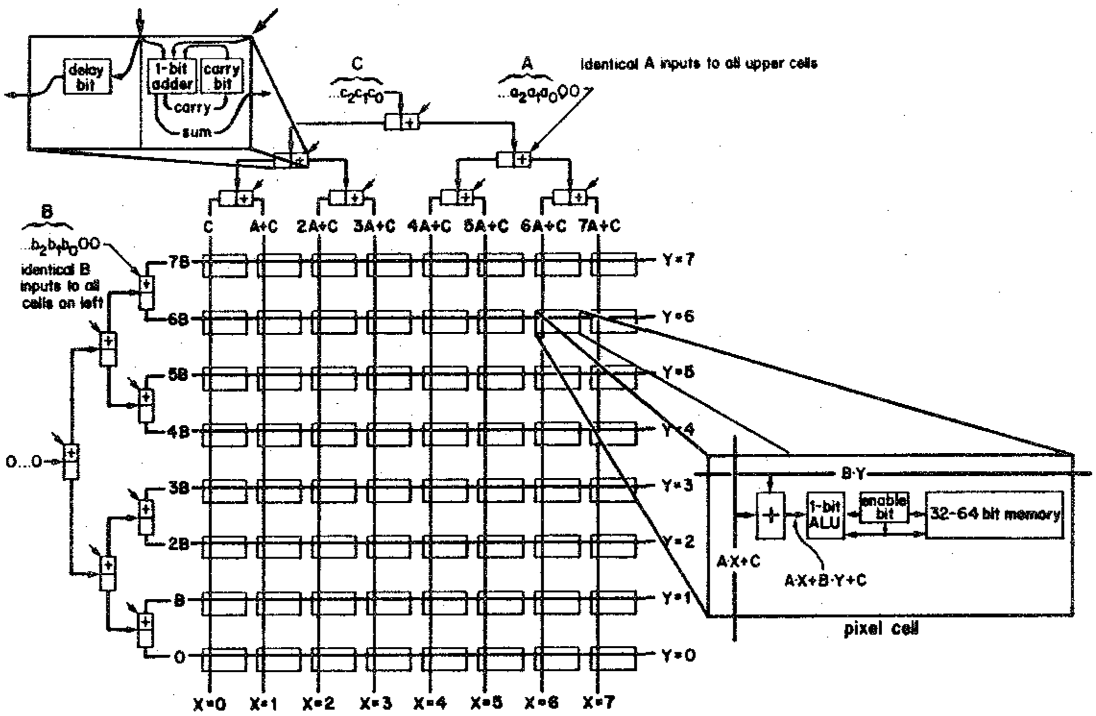

Conceptual design of an 8 x 8 pixel PIXEL-PLANES image-buffer memory chip.
Scan conversion, hidden-surface elimination and color-rendering commands are translated outside the memory system into A,B,C coefficients and associated PIXEL-PLANES commands.

Figure 4. The structure of the Pixel-planes memory.

x,y) to calculate this linear function for its own value of x,y such that this operation is completed for all pixels in 30 clock cycles (presently under 3 microseconds). This machine, which on the basis of prototypes promises to take only about twice the chip space as an ordinary frame buffer, will be able to do the hiding and shading calculations for over a thousand tiles in the 1/30 second frame time of a video display. Thus the displayed image for a complex medical 3D object should be calculated in a fraction of a second, allowing interactive control of translation, orientation, lighting, transparency, etc.

2. 3D Trackers

Interaction in 3D, for example to control location and orientation in 3-space, requires special interactive devices. These may be hand-held or, even better, head-mounted [Sutherland, 1968] so that the display can be controlled by the viewing po-

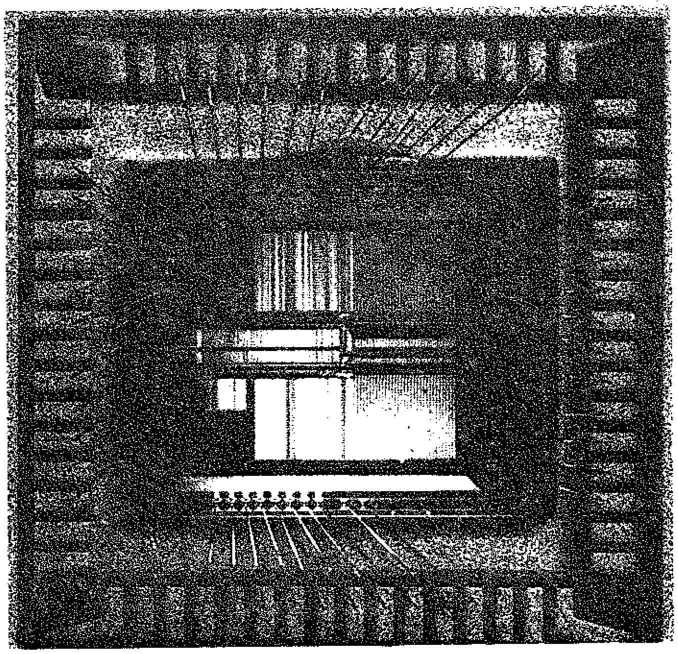

Figure 5. A prototype Pixel-planes chip (stage 3).

sition and direction of the viewer. With an adequately fast display system and a head-mounted display device, the viewer can feel that he is in the image space, such that ordinary movement of his body makes him feel that he is moving through the image (this has been called a *walkaround* display).

Among the 3D trackers available are mechanical devices such as trackballs and joysticks, pens that use the delays of sound to determine position, devices based on a magnetic field, those following special lights, either mounted on the observer and followed from special room locations or vice versa, and a newly developed prototype that follows natural features of the room. The 3D mechanical devices are in common use but are not very portable and allow the specification only of orientation, or only translation for some joysticks. Furthermore, 3D joysticks frequently do not provide continuous rotation or the specification of all possible angles, and 3D trackballs provide only relative orientation. Similarly, the sonic pen allows the specification only of location. The magnetic devices and those based on fixed or user-mounted lights overcome this restriction. However, the former have a somewhat limited range, and the latter provide somewhat limited accuracy, at least in a room-sized environment.

The "self-tracker" of Bishop [1984a,b] is designed to avoid all these difficulties without requiring special instrumentation of the viewing environment with lights or sensors, other than the necessary computing equipment.

Figure 6. A mockup of the Bishop self-tracker (photograph by Jerry Markatos).

The Bishop self-tracker (see figures 6 and 7) consists of many 1D sensors looking out in various directions from the viewer to the surrounding environment. It operates by correlating pictures at successive intervals of one or more milliseconds and calculating movement from the shifts required to achieve maximal correlation. The basic chip acts both as the 1D sensor and as a the correlating computer, and a central system combines the reports from the various chips into a single report of 3D location and orientation. Prototypes of this chip have been built and tested, and simulation based on their properties suggests that system development, now ongoing at Bell Laboratories, will produce a successful system by perhaps 1987.

Besides determining rotation and translation, 3D interactive devices are useful for selecting slices from 3D images for 2D display, for spatial clipping to limit obscuration or hiding problems, and for contrast enhancement in self-luminous displays. The devices mentioned above for determining rotation and translation can also be used for

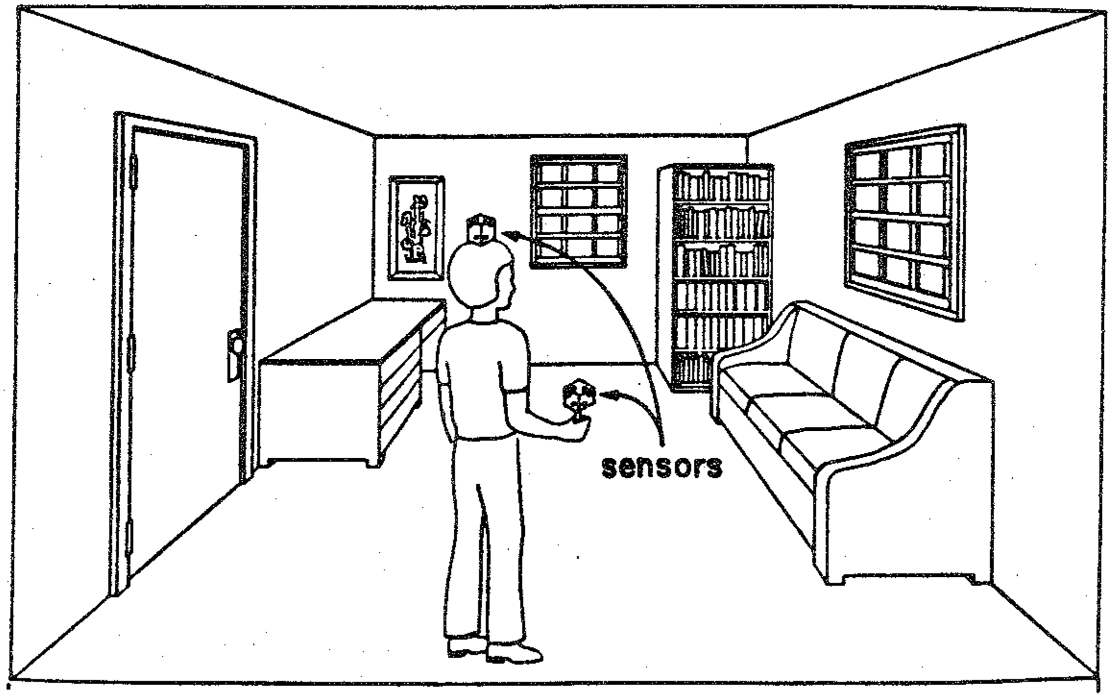

Figure 7. Use of the self-tracker in a "walkaround" display. From [Bishop, 1984a].

slice selection and spatial clipping. Contrast enhancement for self-luminous display is normally controlled by sliders, knobs, or mice.

3. Self-luminous Display Systems

Self-luminous display can be achieved via a directly viewed CRT or by doing the optical equivalent of putting points in 3-space by means such as holography or vibrating or rotating screens or mirrors. With any of these we may display 3D distributions of points, each with its own intensities, or of structures made from lines, such as wire frames. With space-filling self-luminous display, obscuration is such a problem that one never wants to display a full 3D raster of points, and interaction is frequently important to select the region to be displayed or the orientation of view.

3.1 3D Self-luminous Display by CRT

With CRT's the three-dimensional effect is given by providing stereopsis, by producing the kinetic depth effect, in which the object appears to move continuously in 3-space, or by using head-motion parallax via a head-mounted sensor as discussed above. Combining many of these cues is also possible and increases the strength of the 3D percept. All of these CRT-based methods require that the 3D image, consisting of a set of points or lines in 3-space (with co-ordinates x,y,z), be projected onto the 2D CRT (with co-ordinates x',y') from one or more orientations (along z').

Projection involves computing x', y', and z' from the recorded x, y, and z values for each point to be displayed, discarding z', and displaying at screen location x',y' a point with the recorded intensity of the point to be displayed. The transformation to the primed (*viewer*) co-ordinate system from the unprimed (*world*) co-ordinate system involves matrix multiplication, as well as division if perspective projection is used. This calculation must be done for every point.

For the display of lines the endpoints of each line must undergo a projection calculation as above, and one of the following three steps then must be undertaken with these endpoints together with any individual points to be displayed.

(1) Interpolate points intermediate to the line's endpoints, and put the result into a frame buffer, summing intensities in pixels that each picture element (pixel) on the display screen. Alternatively, each receive more than one result. One may possibly remove the stairstep artifact caused by discrete pixels by an approach called *anti-aliasing*, which involves a slight blurring of the line [e.g., Crow, 1981].

(2) Interpolate the points on the line as in method 1 and place the result on a point-plotting CRT, or deliver the endpoints of the line to a vector graphics display system, which will draw the line using analog means. Both point-plotting and vector graphics displays are given x,y,intensity information, rather than a 2D array of pixel intensities, as with a raster graphics system (frame buffer).

(3) Interpolate the points on the line as in method 1, and use a color frame buffer in the following nonstandard way [Fuchs, Pizer, et al, 1982] to make it behave like a point plotter (see figure 8). Load the locations in the frame buffer not with red, green, and blue values as is normally expected, but rather with x,y, and intensity values. Replace the video CRT at the output with a point-plotting CRT, with the three wires that are outputs of the frame buffer (labeled red, green, and blue) plugged into the x, y, and intensity terminals of that point-plotting CRT. With this scheme the intensity lookup tables normally provided for each color in a frame buffer system can be used to transform x, y, and intensity respectively, thus allowing, clipping, scale change, translation, and contrast enhancement.

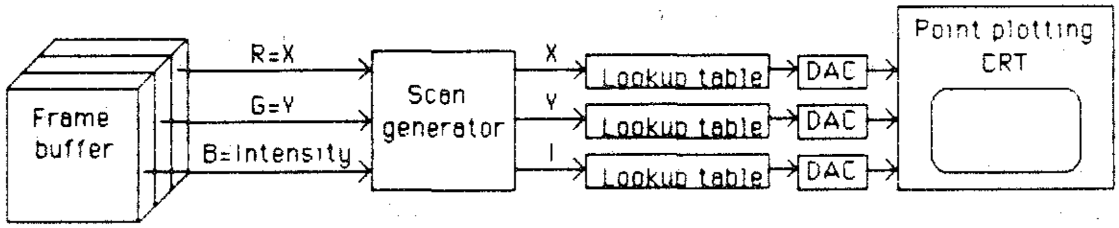

Figure 8. Use of a frame buffer to produce a point-plotting display.

Stereo display simply requires the calculation of two projections, and either the simultaneous display of both images or the placement of the results into two display refresh buffers and the alternation between these buffers. With simultaneous display, mirrors or lenses are necessary to direct the pictures separately to the respective eyes. If the two projections are alternated, one must have a means of alternately blocking the two eyes. Common means for this blocking are via a rotating mechanical device, via goggles with a birefringent material that can by electrical means be made to change between opaque and transparent, and by superimposing a polarizing plate on the screen that can under electrical control change polarization [Bos, 1983], while the viewer wears a differently polarized lens on each eye. The latter method allows many viewers to view the same display and for the goggles to be very inexpensive.

Display using the kinetic depth effect can be achieved by precomputing the projections corresponding to a rocking or rotating motion, but interactively controlled motion provides a more convincing 3D effect [Lipscomb, 1983]. Furthermore, such interactive control is useful for choosing viewing directions that avoid obscuring objects. But this requires calculation of the required projection(s) at on-line speeds. Even fast processors that operate serially on the points cannot provide adequate speed for more than on the order of 10,000 points. What is needed is a system that can transform many points in parallel, if necessary interpolate lines from endpoints in parallel, and if using a raster graphics system add the results in parallel into the frame buffer.

3.2 3D Self-luminous Display by Optical Means

Failing such new fast projection hardware, one can avoid the need to do projection calculations and add the (probably most) important depth cue of head-motion parallax by putting points in 3-space by optical means. The first means that comes to mind is holography, but it is orders of magnitude too slow in producing a new result to allow the critical step of interaction, and it is technically quite demanding. An attractive alternative is to write planar images while moving the writing plane cyclically, fast enough for the planes to fuse into a 3D image. Either screens or mirrors can be moved to provide the required motion of the writing plane.

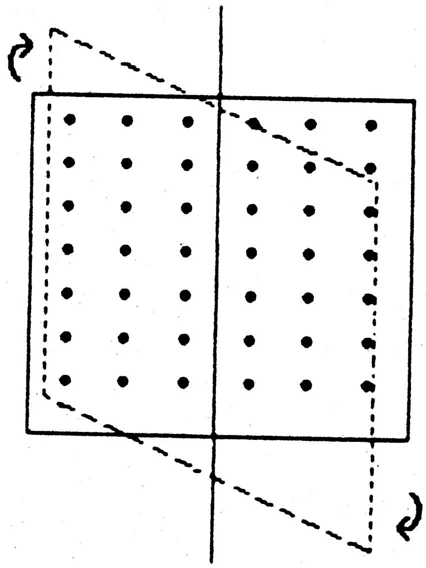

Figure 9. The rotating LED-panel display [Jansson, 1984].

(1) Among the many systems using rotating or translating screens, the most attractive seems to be the rotating LED panel [Jansson, 1984]. In this device (see figure 9) a closely packed array of LED's is rotated about an axis centered within the plane of the array, so that at any time the array forms a one-sided radial slice of a cylinder. The intensities within that slice are formed by controlling the fraction of the time that each LED is on while the array is near some angle. The strength of the device is that it allows all viewing angles and a viewing position close to the 3D display region. Its weaknesses are the need to store and transmit the zero intensities that must dominate a self-luminous display if obscuration is not to be debilitating, and limitations in the spacing of the LED's and the size of the panel for it to be rotated with adequate speed.

(2) Among the many systems using rotating or translating mirrors, the most attractive

is the varifocal mirror because of the small mechanical motion that is required to obtain adequate depth [Traub, 1967]. The principle of the device is shown in figure 10. The viewer looks at a screen reflecting in a mirror that can be made to vibrate between convex and concave, normally by placing a loudspeaker behind the mirror membrane or plate. Because the apparent depth of the reflected screen behind the mirror varies with the convexity or concavity of the mirror, the writing plane is swept over 20-30cm for just a few millimeters of motion of the center of the mirror. Thus, points written on the screen within a short time interval appear on a thin slab parallel to and behind the mirror face at a depth dependent on the time during the display cycle at which the points were written.

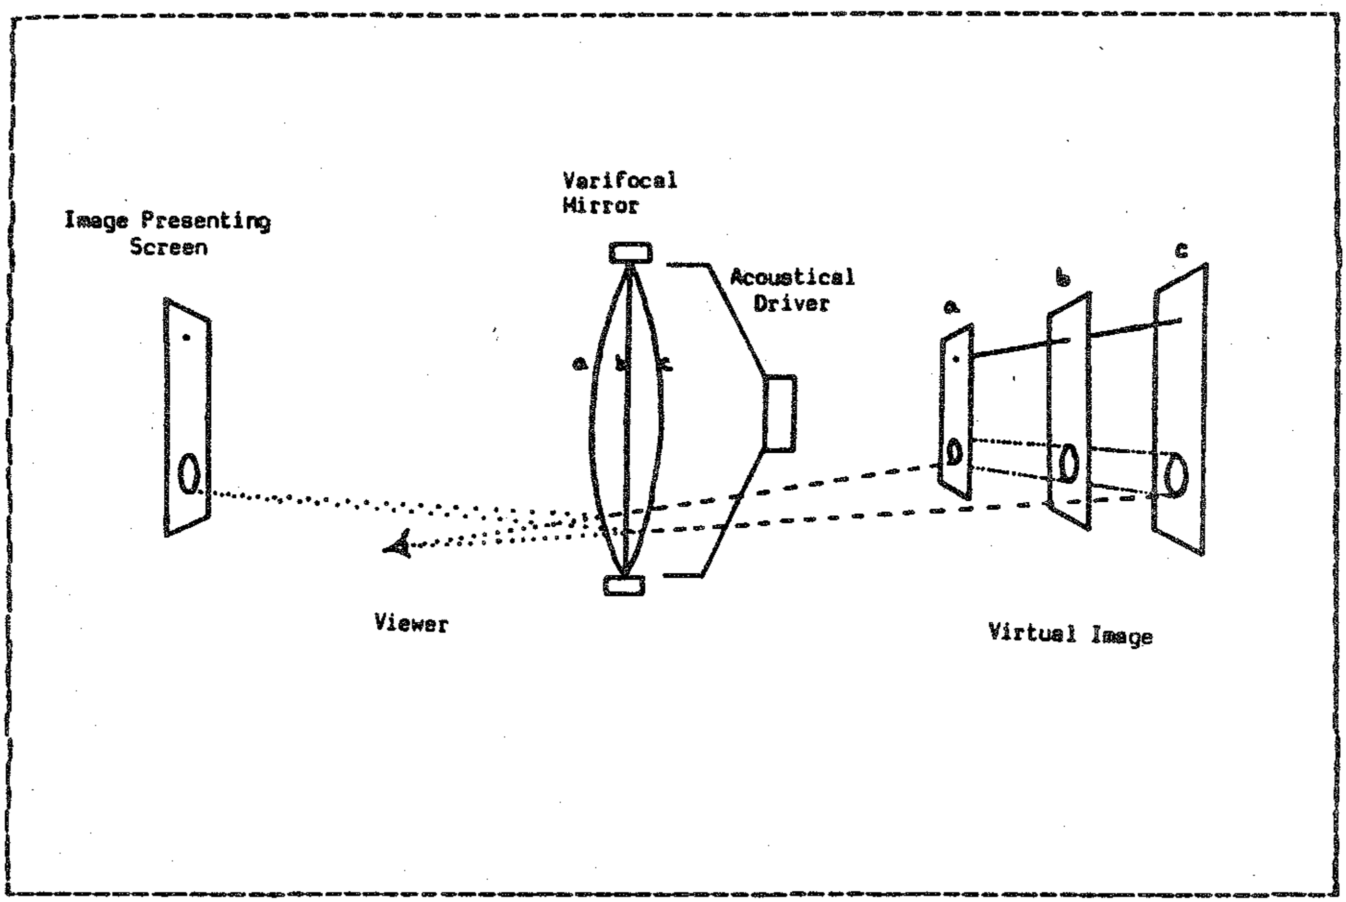

Figure 10. Principle of varifocal mirror display.

Only points with nonzero intensities need to be written onto the screen. These can be written in the form of a 3D raster or as 3D line segments [Baxter, 1980;

Sher, 1980]. Alternatively, if the screen is a serial point-plotting CRT, one can use the method given earlier of turning a raster graphics system into a point plotting system (see figure 8) with the additional restriction that the location in the frame buffer (time in the display cycle) at which a point's co-ordinates are stored must be related to the depth at which it is to appear [Fuchs, Pizer, et al, 1982]. The lookup tables can be used for clipping in the x and y dimensions and for contrast enhancement, and clipping in the z dimension, very important for avoiding obscuration, can be achieved using the registers normally used to specify the y-viewport (which region in y appears on the video screen) when the system is used for raster display.

The fact that the mirror's change from convex to concave causes a variable magnification with depth (so-called anomalous perspective) can be counteracted by varying the x and y scale of the image over time as the writing plane (the screen's reflection) moves in depth. Specifically, the x and y analog signals, taken relative to the screen center, need to be multiplied (gain controlled) by the speaker's sinusoidal input, offset by a constant.

Interactive modification of rotation, objects selected, blinking, scale, and other display parameters can be achieved by the standard graphics approach of having basic object descriptions as points in the world space and transformations from these to viewer space by a processor that frequently is part of the display system (see figure 3). If this processor is not fast enough to transform all of the points in one display cycle, as is commonly the case, a useful approach is that of *successive refinement*, in which a selected group of points are first transformed and displayed, forming a coarse display, with the remainder being transformed and added to the display in later cycles if no change in the interactive parameters controlling the transformation is made in the meanwhile. This approach appears to have some merit as well with reflective display.

4. Conclusion

The preceding has been a brief overview of 3D display devices. While 3D display has great promise in medical imaging, further work is needed in developing devices that are adequately speedy and provide adequate interactive control. New systems will almost certainly be invented, and developments and evaluation of applications in medical imaging are being carried out in large numbers.

5. Acknowledgements

I thank Henry Fuchs and Edward Chaney for useful discussions. I am indebted to Joan Savrock for manuscript preparation and to Bo Strain and Karen Curran for photography. The preparation of this paper was prepared partially with the support of the National Institutes of Health grant number 1-R01-CA39060.

6. References

B. Baxter, "3-D Viewing Device for Interpretation of Multiple Sections", *National Computer Conference, NCC-80*: 437-440, 1980.

G. Bishop, *Self-Tracker: A Smart Optical Sensor on Silicon*, Ph.D. dissertation, University of North Carolina, Department of Computer Science, 1984a.

G. Bishop and H. Fuchs, "The Self-Tracker", *Proc MIT VLSI Conf*, Artech House, Dedham, MA, 1984b.

P. J. Bos, P. A. Johnson, Jr., K. Rickey Koehler/Beran, Tektronix, Inc., "A Liquid-Crystal Optical-Switching Device (π cell)", *Proc. Soc. for Info. Disp.*: 30, 1983.

F. C. Crow, "The Aliasing Problem in Computer-generated Shaded Images", *Comm. ACM, 20* (11): 799, 1977.

H. Fuchs, Z. M. Kedem, and S. P. Uselton, "Optimal Surface Reconstruction from Planar Contours", *Comm. ACM, 20* (10): 693-702, 1977.

H. Fuchs, S. M. Pizer, E. R. Heinz, L. C. Tsai, and S. H. Bloomberg, "Adding a True 3D Display to a Raster Graphic System", *IEEE Computer Graphics and Applications, 2* (7): 73-78, 1982.

H. Fuchs and S. M. Pizer, "Systems for Three-Dimensional Display of Medical Images", *Proc Int Joint Alpine Symp on Med Comp Graphics and Image Commun and Clin Adv in Neuro CT/NMR*, IEEE Comp Soc [IEEE Catalog No 84CH2006-5]: 1-6, 1984.

G. T. Herman, "Computer Graphics in Medicine", in this volume.

D. G. Jansson and R. P. Kosowsky, "Display of Moving Volumetric Images", *Processing and Display of Three-Dimensional Data II, 507*: 82-92, 1984.

J. S. Lipscomb, "Motion Decomposition, Orthogonality, and Stereo Display in a Molecular Computer Graphics System", *ACM Transactions on Graphics*, 1983.

W. M. Newman and Robert F. Sproull, *Principles of Interactive Computer Graphics*, McGraw-Hill, New York, 1979.

J. Poulton, H. Fuchs, J. D.Austin, J. G. Eyles, J. H. Heinecke, C. Hsieh, J. Goldfeather, J. Hultquist, and S. Spach, "Implementation of a Full Scale Pixel-Planes System", to appear in *Proceedings of Conference on Advanced Research in VLSI*, UNC-Chapel Hill, 1985.

L. Sher, *The SpaceGraph Display: The Utility of One More Dimension*, internal paper, Bolt, Beranek, and Newman, 1980.

I. E. Sutherland, "A Head-Mounted Three Dimensional Display", *Proc. Fall Joint Comp. Conf.*: 757, 1968.

I. E. Sutherland, R. F. Sproull, and R. A. Schumacker, "A Characterization of Ten Hidden-Surface Algorithms", *Comput. Surv. 6*(1): 1, 1974.

A. C. Traub, "Stereoscopic Display Using Rapid Varifocal Mirror Oscillations", *Applied Optics, 8*(6): 1085-1087, 1967.