Edward J. Farrell

Color Display and Interactive Interpretation of Three-Dimensional Data

Three-dimensional results from engineering and scientific computations often involve the display and interpretation of a large volume of complex data. A method is developed for color display of 3D data with several interactive options to facilitate interpretation. The method is based on representing points whose values fall within a specified range as a single hue. An image is formed by overlaying successive 2D frames with increasing hue lightness towards the front. Interactive options to aid interpretations are viewpoint, contour lines, multiple range display, slicing, veiled surfaces, and stereo image pairs. The display method is successfully applied to several types of data. The overall structure and variations of the 3D data are observable, as well as transients which may be overlooked in a large input data set.

Introduction

Several important engineering and scientific areas require display and interpretation of three-dimensional data or series of two-dimensional data frames, e.g., seismic sections, computed tomographic (CT) images, multi-dimensional optical scans, statistical cluster data, models of three-dimensional fields, and diffusion processes in a plane during a specified time interval. The spacial relationship of various data values is often complex. In most applications, a single data set is at least 2 million bytes and may be as large as 100 million bytes. The display and interactive interpretation of such large volumes of complex data represent a challenging area for research.

Two classes of methods are currently used to display 3D objects from CT scans. One class forms the outline of the objects in each frame [1, 2]. The outlines are connected to form tiled surfaces which are smoothed, shaded, and projected onto the image plane. The second class of methods forms the surface of an object which is defined by a set of 3D volume elements [3, 4]. The set of object volume elements is selected from the total space of volume elements by thresholding the data values. The object surface is formed by tracking connected faces of the 3D volume elements which border the object. The surface is displayed using algorithms to shade and remove hidden surfaces based on the volume element character of the object [5]. Typically, these CT display methods do not use color and have limited interactive features. A notable exception is the interactive graphics programs described by Dwyer et al. for diagnostic imaging in medicine [6]. Atherton describes an interactive method to visualize CAD surface models using a color display [7]. The interactive options include cutaway views, shading, surface transparency, and translucency. The object orientation is fixed. Graphical presentation of general scientific data is surveyed by Graedel and McGill [8]. They describe several two- and three-dimensional methods based on surface line drawings, contour plots, serial color plots, and 3D color diagrams.

The colored-range method developed in this study differs fundamentally from current methods. The basic objective is to interactively interpret the 3D distribution of values, as opposed to displaying a given set of surfaces or to isolating a 3D object and displaying its surface. The method can be applied to a wide range of engineering and scientific data, from continuous 3D potential fields to statistical clusters of disjoint points, as well as conventional objects based on CT scans.

There are four key problems. First, an effective method of representing a three-dimensional field (or series of 2D data frames) and of displaying it on a two-dimensional screen must be developed. To express depth and shape, color is required. Color has inherently greater capability to express complex data since each colored picture element has three attributes (hue, lightness, and saturation), compared to a monochrome picture element which has only one attribute (lightness). Second, a single method of display and color encoding is not sufficient. The data often require interactive interpretation to determine the spacial distribution of various values. Third, the computer software implementation of the display and interactive functions must be capable of responding within a few seconds. Fourth, the large volume of data must be stored and managed to permit interactive interpretation.

The first two problem areas are of primary interest in this paper: color display of 3D data and interactive interpretation. The objective is to develop a basic approach and then specific options that are useful in several engineering and scientific applications. The approach is based on representing all points whose values fall within a specified range as a single hue to produce a 3D colored-range representation. An image is then formed by isometric overlay of successive 2D data frames with increasing color lightness toward the front. Several options are presented for interactive interpretation: viewpoint, contour lines, slicing, veiled surfaces, and stereo pairs. These methods permit the implementation of high-speed software with simplified data management.

Color display of three-dimensional data

A key problem is selecting an effective method of representing a three-dimensional field or a series of 2D data frames. There are several possibilities. For example, small numerals (or symbols) can be placed throughout a 3D rectilinear region to display magnitude, or arrows can be used to represent the direction and magnitude change. A colored-range representation is implemented in this study. For example, if the input data range between 1 and 256, all points in 3D with data values between 50 and 60 are represented as red, points with values between 90 and 100 as orange, 130 to 140 as yellow, and 170 to 180 as green. The choice of ranges depends on the nature of the data. If more than four colors are used, the representation becomes difficult to interpret. The overall result is a 3D rectilinear region with clusters or layers of different colors, i.e., 3D colored ranges.

This 3D colored-range representation of a field must be displayed as a 2D image on a CRT. There are a variety of attributes a human observer uses to interpret the spacial character of objects like the 3D colored ranges: perspective, shading, texture, depth of focus, stereopsis, etc. [9, 10]. Several of these attributes have been successfully implemented for computer display of 3D objects. Perspective was not used in this study because we wanted to ensure a rapid response time. Further, with oblique projection the size of an object in the rear plane is directly comparable to the size of an object in the front plane. Shading/lightness is not used to represent shape based on a light source and viewing position. Lightness is used to indicate z-axis position as described below. This basic approach has two advantages: computation time is less, and the 3D colored range need not be continuous with a surface normal (a disjoint cluster of points can be displayed).

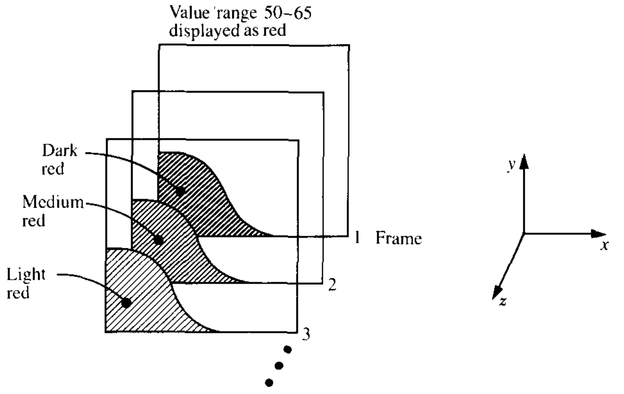

Figure 1 Illustration of colored-range display.

The display method used in this study is oblique projection, with lighter colors for points in the front data frames. This method is illustrated in Fig. 1. All points whose values range between 50 and 65 are colored red. In the rear frame, a dark red is used; in the next frame, a lighter red is used, etc. Lightness is used to indicate the z coordinate of a point; the x, y coordinates can be observed directly since it is an oblique projection. For each range of data values, there is a sequence of colors from dark for the back frames to light for front frames. Each color sequence is predominately one hue, such as red, orange, yellow, or green. Successive frames are offset one image element down and to the left, and replace the color entries of the preceding frame. In this way a 2D image is formed of the 3D colored-range representation of the field. The data can be viewed from eight different "directions" by transposing the data along the x or y axes or by taking the frames in reverse order. Since forming these images does not require arithmetic operations, it is relatively fast. For the variety of applications investigated in this study, it is sufficient to view the data from these eight directions (together with using the interactive functions described later). Viewing the rectilinear data space from any direction would not yield enough additional information to warrant the large increase in computation. In some applications the number of frames may be on the order of 20, while the x and y dimensions are

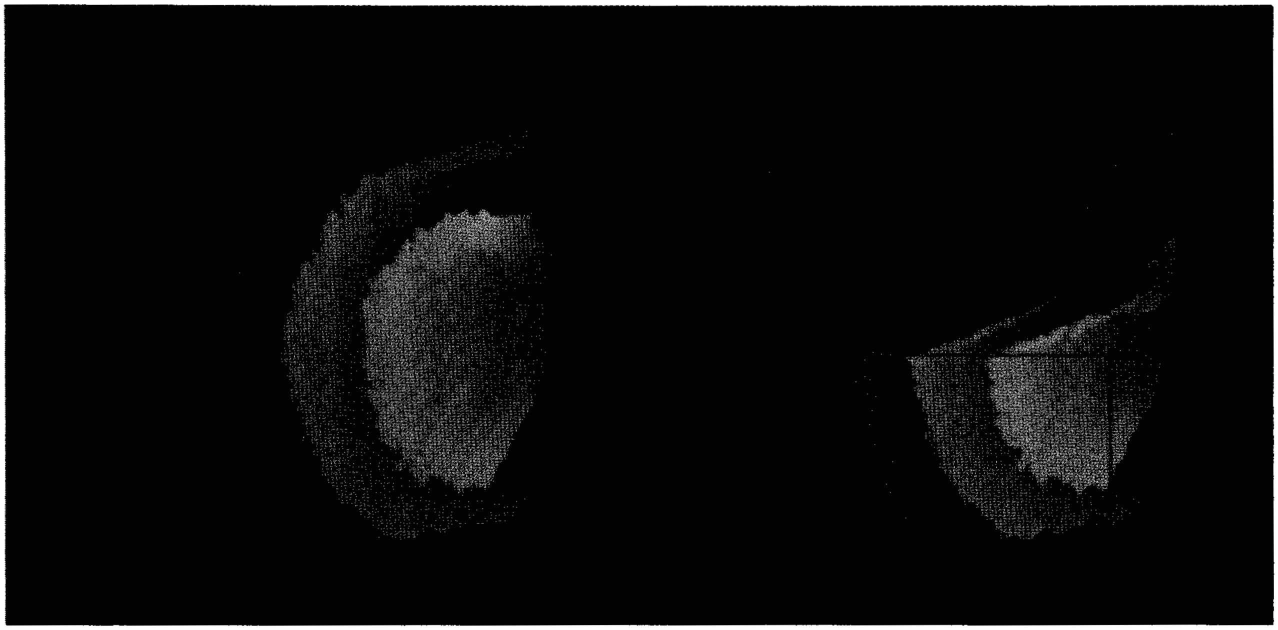

Figure 2 Relative brightness of a test object positioned in an illuminated field of view.

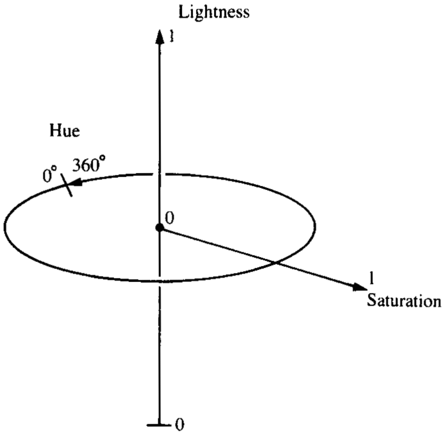

Figure 3 Color space model.

200 to 300. In this case, frames may be repeated once or twice to enhance the three-dimensional appearance of the images.

An example of colored-range display is shown in Fig. 2. (Two imaging systems are used by the author: the IBM 7350 Image Processing System with a screen 1024 × 1024 pixels and a Ramtek Graphic Display System with a screen 1024 × 1280 pixels. A MATRIX Color Graphic Camera is used to photograph the images.) The data are the measured relative brightnesses of a test object as a function of its position in a rectilinear illuminated field of view. The light source is from the observer's right. The points in each colored range are positions of comparable brightness. The yellow points correspond to positions of greatest brightness. Since this method does not require the data fields to be continuous, it can be used to display disjoint points (as illustrated later in Fig. 10) and clusters of statistical data. To correctly interpret the display, it is necessary to visualize dark points as distant and light points as near. The lines representing the coordinate frame are also shaded with the lightest in front (Fig. 2). The lines are handled like data using a gray scale instead of a sequence of colors.

- *Color selection*

The display of a 3D field is based on sequences of colors, such as the green, red, orange, and yellow in Fig. 2. On the other hand, the color display unit requires values of red, green, and blue (RGB) to generate a color image. It is very difficult to define colors based on RGB values since they are not easily related to what is perceived. It is necessary to define color in terms of perceptual parameters, hue, lightness, and saturation (HLS) [10, 11]. They can be interpreted as cylindrical coordinates in the space of colors (see Fig. 3). Hue is periodic and expressed in degrees; lightness and saturation range from 0 to 1. There are several methods of transforming HLS and RGB values [12–15]. From a theoretical point of view, it

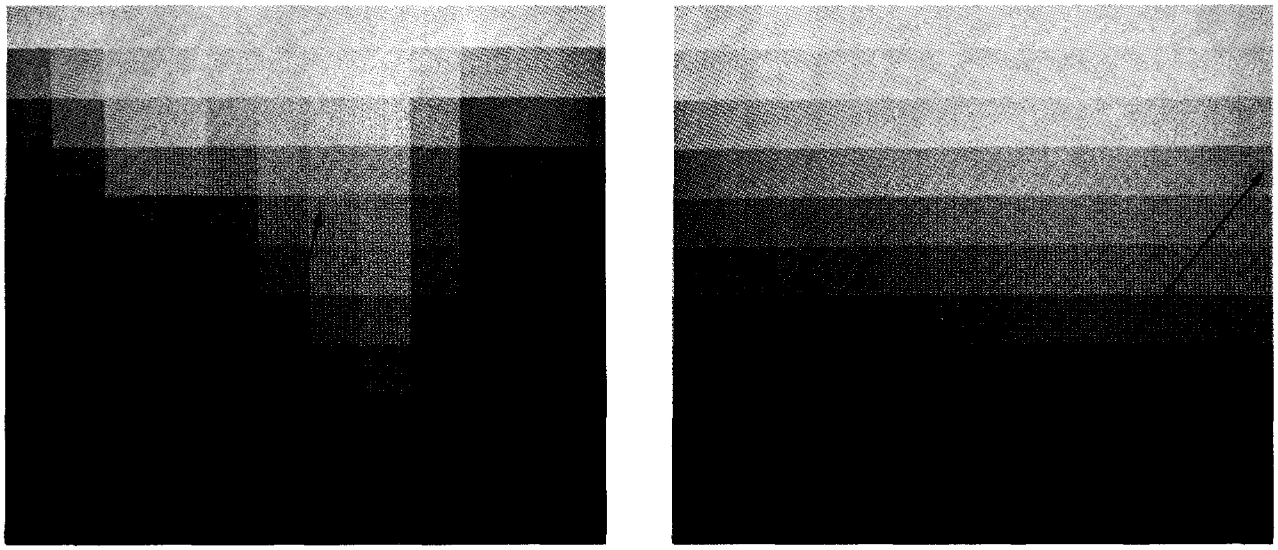

Figure 4 The nominal orange sequence projected onto a surface of saturation 1.0 (left) and constant hue of 120° (right).

would be preferable to have equal increments in HLS parameters result in equal perceived changes on the color display unit and output film.

Munsell defines an equal-perception coordinate system for opaque colors [16]. Based on the tristimulus CIE values for the Munsell colors and the phosphor characteristics of the CRT, a transformation from perceptual parameters HLS to RGB can be defined [15]. As a simple approximation to this exact transformation, the triangle model derived by Smith was implemented [12]. It was modified by removing the intensity normalization. This transformation of HLS to RGB resulted in nearly uniform perceptual increments on the CRT for uniform increments in HLS.

This approach was not satisfied from a practical viewpoint. For a given level of saturation, not all values of lightness are achievable with the CRT. The realizable values of HLS lie in a limited region which cannot be simply described by a mathematical relation [15]. This is characteristic of any method based on a perceptual relation like the Munsell system. Also, the preferable trajectories for the color sequences are not straight lines in such HLS spaces, and they are difficult to define.

The HLS to RGB transformation proposed by the SIGGRAPH Standards Committee has several advantages [14, 17]. (It is presented in the appendix.) All values of HLS result in realizable colors with the CRT, and straight line trajectories in the HLS space result in very useful color sequences for display of 3D data. Also, it is computationally simple. The method does not provide equal perception steps like the Munsell coordinates, but this is not a significant disadvantage. The "best" color to display an image will depend on the character of that particular image. A perceived color of a particular region depends strongly on the surrounding colors [18, 19]. Colors must be selected based on the positions and areas of various elements in the image. Also, it may be necessary to shift the lightness range of the color sequence to emphasize a z-axis interval of special interest. Thus, interactive selection of the color sequences is required. Further, with an interactive selection of colors, it is easy to compensate the color sequences for various output modes (CRT, prints, 35-mm transparencies), display calibration drift, and ambient light. Photographic prints and 35-mm transparencies are sensitive to color selection; an inappropriate color could obscure significant features. The CRT display is less sensitive than film; small structures and color changes can be easily seen.

Six color parameter values must be selected to define each color sequence in the HLS color space: first and last hue, first and last lightness, and first and last saturation. Twenty-four values are needed for the four color sequences. For a practical, interactive selection of color, it is convenient to fix certain parameters. Perceived color is most sensitive to changes in hue and least sensitive to changes in saturation. From trial and error, it was found that the most imaging requirements are met with the first and last saturation values fixed at 0.6 and 1.0. The last value of 1.0 was selected to

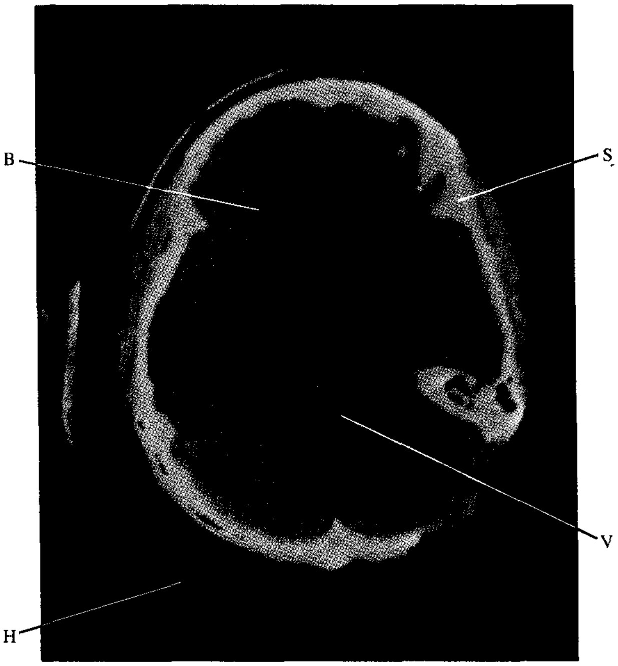

Figure 5 CT scan frame.

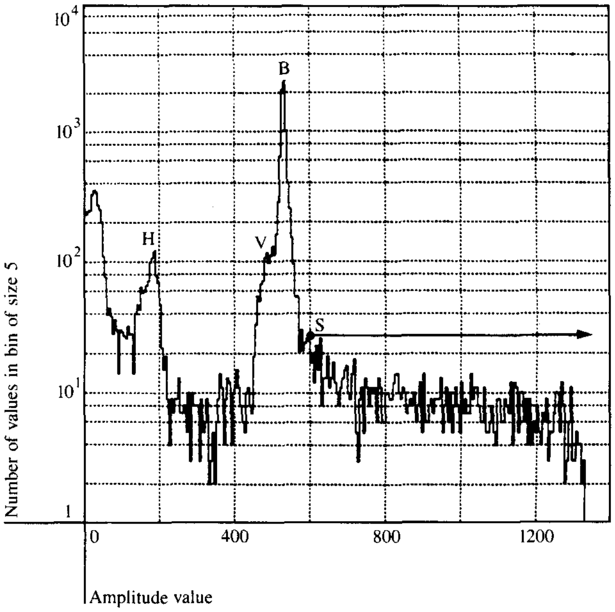

Figure 6 Histogram of CT frame shown in Fig. 5.

obtain maximum perceptual difference between the colors. A value less than 0.6 for the first point reduced the perceptual difference between the sequences; a value greater than 0.6 reduced the "number" of perceived colors in each sequence. Hue overlap between the color sequences must be limited to avoid identifying image points with the wrong sequence. Also, the collective hue range of the four sequences should be distinct from the background hue. Hue ranges of 50° are satisfactory for most cases. The mid-points of the hue ranges are nominally set at 105°, 155°, 190°, and 240°; these values are modified to optimize the perceptual interpretation of the image. The first and last lightness values are nominally set at 0.1 and 0.7 but are also modified interactively. With these conversions, the nominal color sequences are defined, and the user has the option of modifying three values per sequence: hue, first lightness, and last lightness. The nominal orange color sequence is illustrated in Fig. 4 in relation to the total gamut of colors.

The background hue is fixed at 0° blue, since blue is often interpreted as "space." The background lightness is fixed at 0.5, which is between the darkest and lightest parts of the image. A greater sense of depth and shape is obtained with the background lightness between the object lightness extremes [10]. A saturation of 0.5 is used to obtain a distinct blue color, which enhances object colors.

- *Frame preview*

The colored-range representation of a 3D field (or series of 2D data frames) is based on ranges of data values which are displayed as colored regions in the 3D data space. If the numeric character of the data is unknown initially, it is necessary to preview several typical frames before selecting the data value ranges.

Three steps are used in previewing frames. First, a histogram of values is generated. Next, a gray level 2D image is formed, and the values of various regions are determined with cursor pointing. These values can be correlated with features of the histogram. Third, a value range is selected, and a black mask is displayed over the original 2D image to show the areas selected by the value range.

A single frame of a CT scan of a head is presented in Fig. 5, and the histogram of values in Fig. 6. Note the corresponding points S, B, and V. Point H is the head rest, which is not visible with the contrast function used in Fig. 5. To display the ventricle V, a value range of 515 to 525 is selected. Points in this value range are displayed as red in the left side of Fig. 7. There are many points in this value range which are not in the ventricle. In most cases, these points are disjoint, i.e., do not form a cluster or surface. Low-pass filtering would reduce these isolated values, but it would also reduce resolution and increase the computation time. A median filter is effective to eliminate isolated values without reducing resolution, but it also increases computation time [20]. A neighborhood filter is effective in eliminating most isolated points with only a small increase in computation time. It operates

Figure 7 CT scan data. Points for values 515 to 525 without (left) and with (right) 4-neighbor filtering.

on the 2D binary array corresponding to the points selected for the value range. The neighborhood filter retains a point if n of its 8 nearest neighboring points are also in the value range. The choice of n depends on the data. The CT frames are filtered with $n = 4$, and the result is presented in the right side of Fig. 7.

Interactive interpretation

Because the data are complex in most applications, a single image based on colored ranges is not sufficient to interpret the relationship of data values. Different display options must be tried; the option used at a particular step depends on data, prior image options used, and information sought. There are three types of options: shape perception, spacial relationship, and hidden structures.

The shape perception of a colored range is enhanced with three options. First, the 3D data space can be imaged from eight different directions. This is done by changing the indexing order of the rows and columns in each data frame and by changing the order in which the frames are processed. Changing the direction of imaging does not require any computation and is quickly executed. Second, shape perception is improved by forming contour lines on the surface. Figure 8 contains two equi-potential surfaces of three point charges with contour lines. The x and y contours are formed by using the contour color in each tenth column and row of all frames. Third, stereopsis greatly enhances the perception of shape. Since stereopsis is a robust phenomenon, effective stereographic pairs can be formed using the simple method illustrated in Fig. 1 of overlaying successive frames. The horizontal displacement of frames is the integer position of the product of frame number and cotangent of the z-axis elevation angle. Typical elevation angles for the left and right images are 45° and 50°. The angles are selected interactively with larger angle differences for shorter z axes.

The spacial relationship of the value ranges is perceived by interactively selecting different colored ranges for display. The ventricles in a CT scan are shown in Fig. 7. Bone is displayed in the left side of Fig. 9. In the right side of Fig. 9, bone and ventricle are combined to show their spacial relationship. The red points outside of the skull at the lower left result from the head rest. Sulci and nasal sinus have also been interrelated in this way. Any normal or pathologic structure with a radiographic density different from its surroundings can be imaged. These images were formed without manual tracing, surface tiling, smoothing, or surface detection algorithms used in other studies.

Colored ranges hidden behind another range can be seen in two ways. The 3D data representation is sliced to remove a quadrant and expose hidden colored ranges (Figs. 2 and 9). The slice point is interactively varied throughout the data

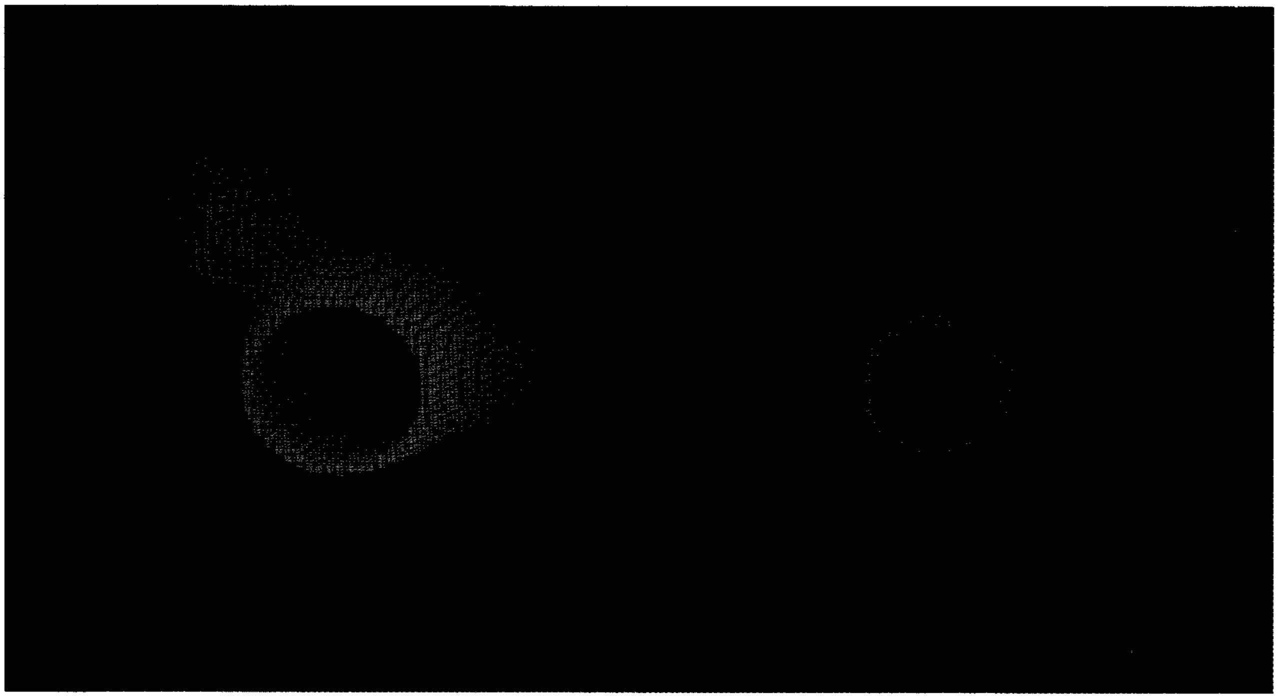

Figure 8 Equi-potential surfaces of three-point charges with contour lines for two potential levels.

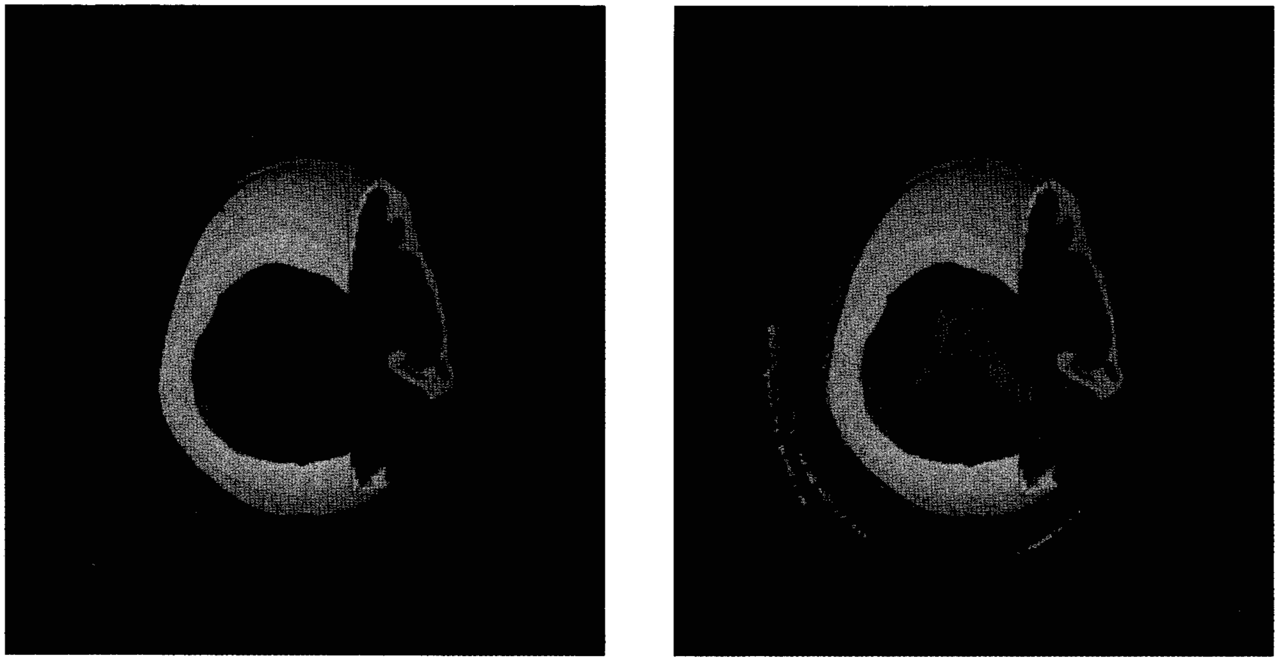

Figure 9 Skull and ventricles constructed from 25 CT scans.

space. A sliced image is formed by excluding a rectangular region while overlaying successive frames for z greater than the slice point. A second way to see hidden ranges is to veil the obstructing range. It is veiled in the process of overlaying successive frames (Fig. 1). Only alternate points in a checkerboard distribution are used from the veiling colored range.

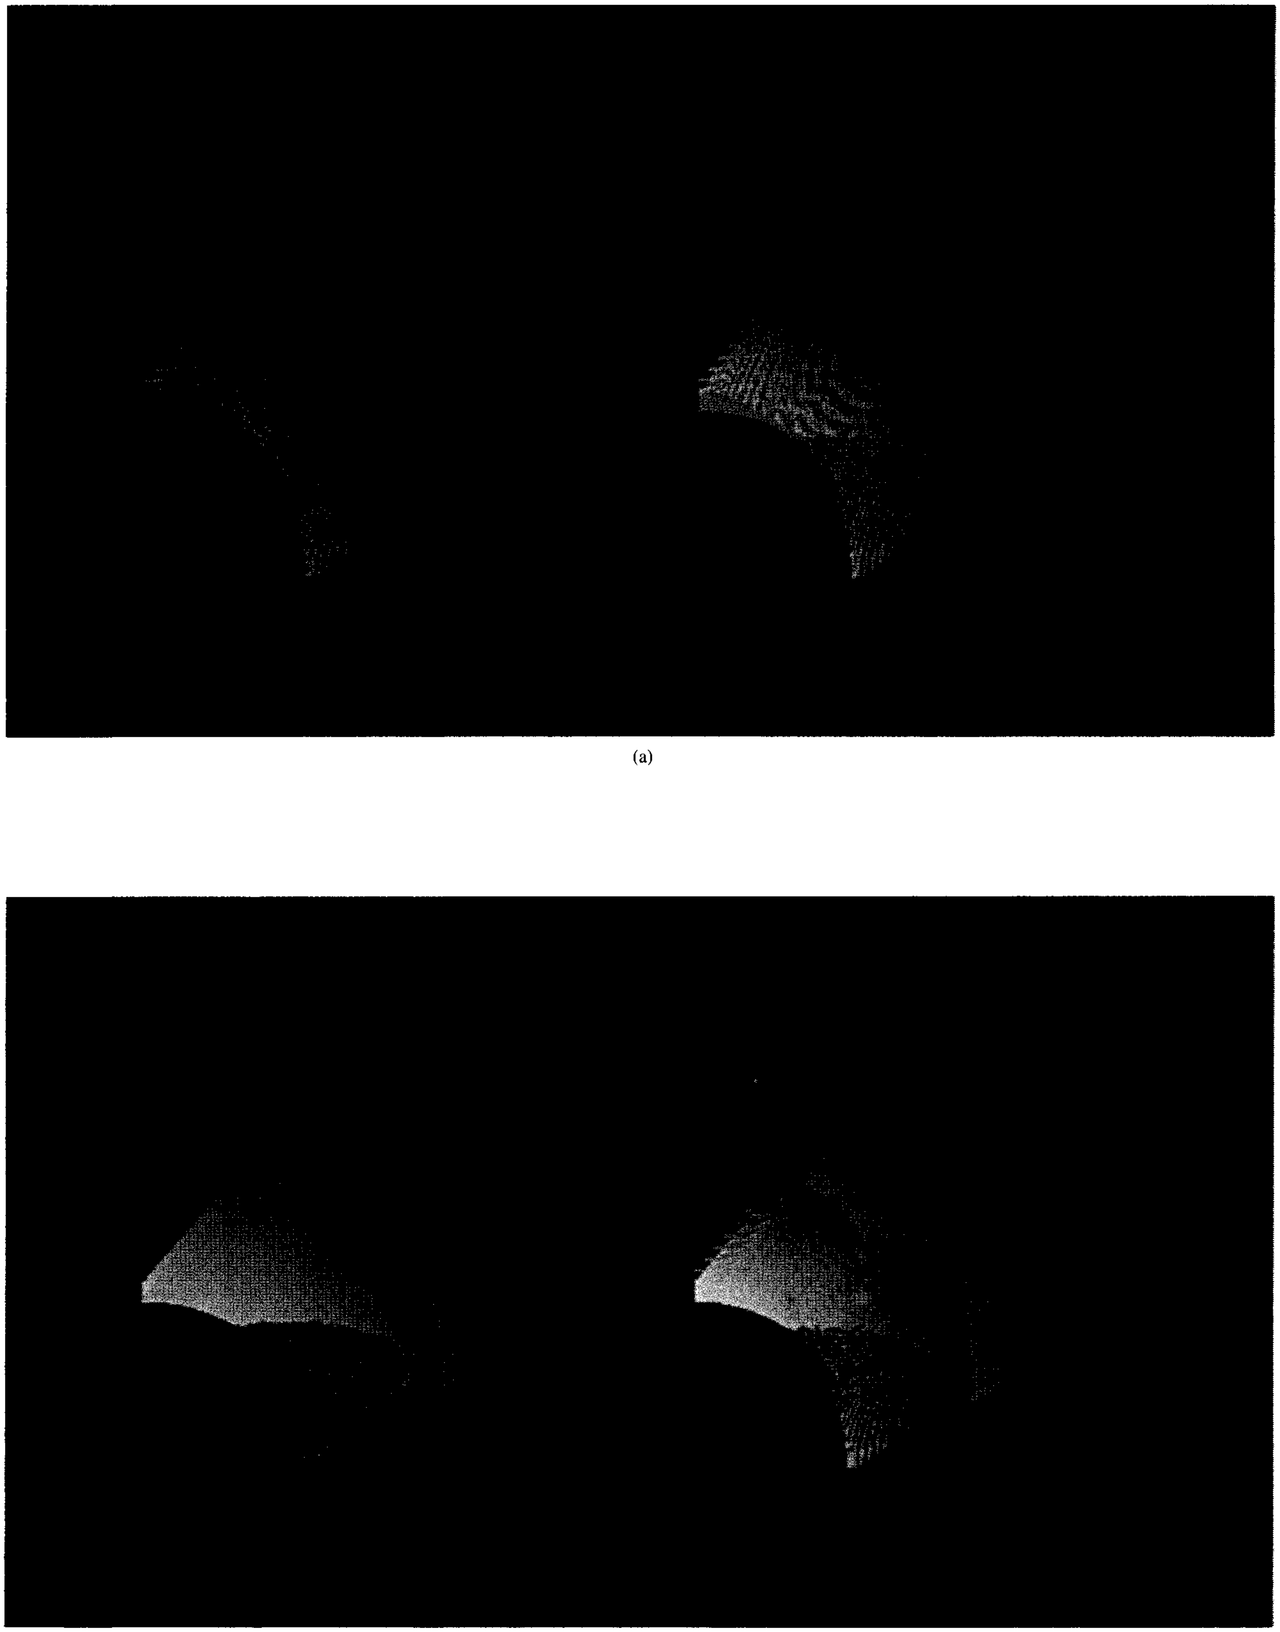

(a)

(b)

Figure 10 Mass density in a plane (x, y axes) as a function of time (z axis). (a) The temporal change of the initial level of mass density (orange) is illustrated in the left image. The right image shows layering with a higher mass density. (b) The space and time relationships of the initial and higher densities are illustrated. The initial density (orange) is veiled in the right image.

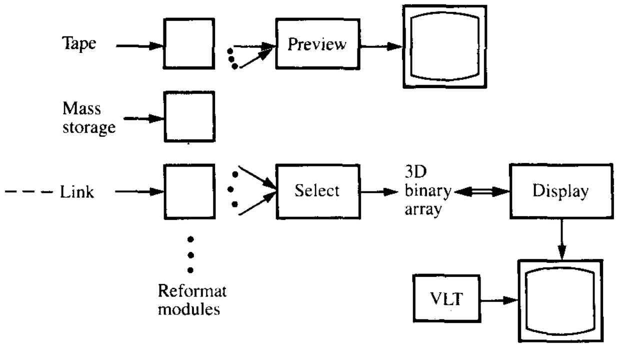

Figure 11 Data and function organization.

This results in a true admixture of the veiling and background colors, as opposed to using a third color to approximate the veiling color. This option is shown in Fig. 10(b), cf. Fig. 10(a). The *x* and *y* axes are position coordinates, and the *z* axis is time. The data are mass density. Initially the density is uniform in an annulus in the first quadrant. The density evolves in an asymmetric way in time, as shown by the orange surface in Fig. 10(a). The red surface is a density range higher than the initial density. Note that this higher density starts after one third of the total time has passed, Fig. 10(b). It is layered with the initial density and has periodic temporal variations. The initial density surface is veiled in the right image of Fig. 10(b) to illustrate the spacial and temporal relationship of the two density levels.

Discussion

Data storage and management were considered in the basic display approach. The colored-range selection results in 3D binary arrays, e.g., one for orange, one for red, and one for yellow. These binary arrays are used for display and interactive interpretation. Various options (slice, viewpoint, grid lines, etc.) can be tried without addressing the original data. Using binary arrays with data compaction can reduce the storage requirements by an order of magnitude. The basic approach is diagrammed in Fig. 11. The input data are stored in their original compacted form on tape, a mass storage system, or a data file accessed by linking. The data are decompressed and reformatted into a 3D array. The PREVIEW function is used to select the value ranges. The SELECT function forms 3D binary arrays based on the chosen value ranges. The DISPLAY function forms images based on the binary arrays; it is not necessary to access the raw data to examine the data with different interactive options. The color sequences are changed by directly loading a new video look-up table (VLT). Figure 11 is an initial approach to data management; other approaches are being considered for this key problem area.

The computation speed of software implementation was considered in the basic display approach. There are very few arithmetic operations; most options involve logical operations or indexing. The display programs were written in APL and executed on an IBM 370/3081. Typically 30 to 90 seconds of CPU time are required to form an image from binary arrays of two to four million bits. This is a key problem area, and several software approaches are being investigated.

In the preceding sections, several methods to display and interpret 3D data were presented. Such methods are essential with large quantities of complex data. For example, the data displayed in Fig. 10 involve almost two million bytes. Without a 3D color display, it is very difficult to perceive how the initial mass density (orange) changes, when and where a higher density (red) occurs, and how the higher density changes with time. The 3D layered structure of Fig. 10 can be clearly seen by reversing the *x* axis and displaying the stereo pair, Fig. 12. (Stereo can be seen with a viewer or directly with a relaxed gaze using the left eye for the left image and the right eye for the right image.) This example illustrates the utility of the interactive, multi-option display method for the user to interpret complex 3D data arrays. He can observe not only global variations, but also transient changes that may be overlooked in a large data array. The display method developed in this study can be applied to a wide range of applications because there are no constraints on the data character, such as value continuity.

Acknowledgments

The interface programs from APL to the display were developed by Paul Sholtz (IBM Research); his valuable work was essential to obtain rapid interaction from APL. Some of the data used in this study were kindly provided by other people: CT scan data by A. Zingher (IBM Research) and Wen C. Yang, M.D. (Department of Radiology, Beth Israel Medical Center, New York, NY); mass density data by K. Winkelmann (IBM Research and GSI, West Germany); and test object brightness data by B. Carrihill (IBM Research). Several colleagues at IBM have given helpful suggestions and assistance: D. Bantz, D. Reich, and A. Appel. This work was supported in part by the National Science Foundation University/Industry Joint Research Grant ECS 80-17683 with the University of Rochester, New York. As principal investigators, Professor Robert Waag of the University of Rochester and C. N. Liu of IBM have been helpful during the course of this study.

Appendix

The following equations are used to transform hue, lightness, and saturation (HLS) to red, green, and blue (RGB) [17]. H ranges from 0° to 360°; L, S, R, G, and B range from 0 to 1.

Set-up equations:

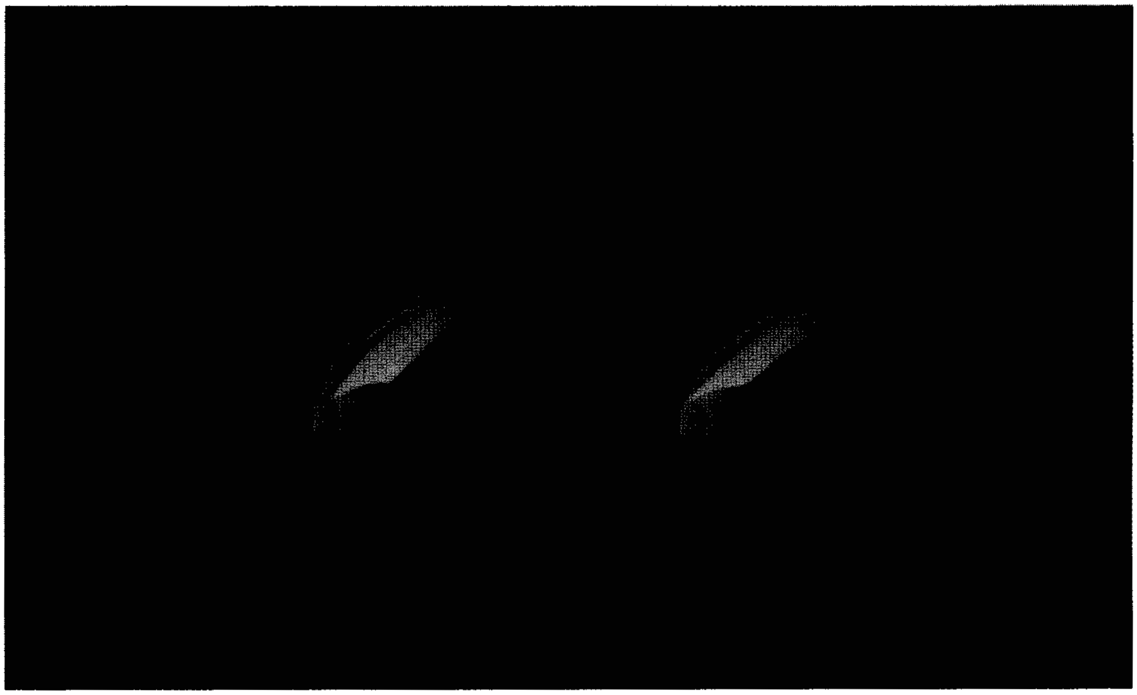

Figure 12 The initial mass density of Fig. 10 is imaged in a stereo pair to clearly show the layered character of the values. The x axis is reversed compared to that of Fig. 10 (a) and (b).

If $L \leq 0.5$, $M = L(1 + S)$.

If $L > 0.5$, $M = L + S - LS$.

$m = 2L - M$.

Equations for calculating R:

If $0 \leq H < 60$, $R = m + (M - m)\left(\frac{H}{60}\right)$.

If $60 \leq H < 180$, $R = M$.

If $180 \leq H < 240$, $R = m + (M - m)\left(\frac{240 - H}{60}\right)$.

If $240 \leq H < 360$, $R = m$.

Equations for calculating G:

If $0 \leq H < 120$, $G = m$.

If $120 \leq H < 180$, $G = m + (M - m)\left(\frac{H - 120}{60}\right)$.

If $180 \leq H < 300$, $G = M$.

If $300 \leq H < 360$, $G = m + (M - m)\left(\frac{360 - H}{60}\right)$.

Equations for calculating B:

If $0 \leq H < 60$, $B = M$.

If $60 \leq H < 120$, $B = m + (M - m)\left(\frac{120 - H}{60}\right)$.

If $120 \leq H < 240$, $B = m$.

If $240 \leq H < 300$, $B = m + (M - m)\left(\frac{H - 240}{60}\right)$.

If $300 \leq H < 360$, $B = M$.

References

1. J. F. Greenleaf, J. S. Tu, and E. H. Wood, "Computer Generated Three-Dimensional Oscilloscopic Images and Associated Techniques for Display and Study of the Spatial Distribution of Pulmonary Blood Flow," *IEEE Trans. Nucl. Sci.* **NS-17,** 353–359 (1970).
2. E. Keppel, "Approximating Complex Surfaces by Triangulation of Contour Lines," *IBM J. Res. Develop.* **19,** 2–11 (1975).
3. E. Artzy, G. Frieder, and G. T. Herman, "The Theory, Design, Implementation and Evaluation of a Three-Dimensional Surface Detection Algorithm," *Computer Graph. & Image Process.* **15,** 1–25 (1981).
4. J. K. Udupa, "Interactive Segmentation and Boundary Surface Formation for 3D Digital Images," *Computer Graph. & Image Process.* **18,** 213–235 (1982).

5. G. T. Herman and H. K. Liu, "Three-Dimensional Display of Human Organs from Computed Tomograms," *Computer Graph. & Image Process.* **9,** 1–21 (1979).
6. S. J. Dwyer, W. H. Anderson, M. A. Tarlton, P. N. Cook, K. R. Lee, L. T. Cook, and S. Batnitsky, "Interactive Computer Graphics for Diagnostic Imaging," *Computer Graph. World* **4,** 46–53 (1981).
7. P. R. Atherton, "A Method of Interactive Visualization of CAD Surface Models on a Color Video Display," *Computer Graph.* **15,** 279–287 (1981).
8. E. T. Graedel and R. McGill, "Graphical Presentation of Results from Scientific Computer Models," *SCIENCE* **215,** 1191–1198 (1982).
9. J. J. Gibson, *The Perception of the Visual World*, Houghton Mifflin Company, Boston, 1950.
10. L. Stroebel, H. Todd, and R. Zakia, *Visual Concepts for Photographers*, Focal Press, Inc., New York, 1980.
11. D. Judd and G. Wyszecki, *Color in Business, Science and Industry*, 3rd Ed., John Wiley & Sons, Inc., New York, 1975.
12. A. R. Smith, "Color Gamut Transform Pairs," *Computer Graph.* **12,** 12–19 (1978).
13. G. H. Joblove and D. Greenberg, "Color Spaces for Computer Graphics," *Computer Graph.* **12,** 20–25 (1978).
14. "Raster Graphics Handbook," Conrac Corporation, Covina, CA, 1980.
15. G. W. Meyer and P. P. Greenberg, "Perceptual Color Spaces for Computer Graphics," *Computer Graph.* **14,** 254–261 (1980).
16. A. E. Munsell, *A Color Notation*, Munsell Color Company, Baltimore, MD, 1946.
17. SIGGRAPH, "Status Report of the Graphic Standards Planning Committee," *Computer Graph.* **13** (1979).
18. J. R. Truckenbord, "Effective Use of Color in Computer Graphics," *Computer Graph.* **15,** 83–90 (1981).
19. J. Albers, *Interaction of Color*, Yale University Press, New Haven, 1963.
20. E. J. Farrell, "Backscatter and Attenuation Imaging from Ultrasonic Scanning in Medicine," *IBM J. Res. Develop.* **26,** 746–758 (1982).

Received January 24, 1983; revised February 22, 1983

Edward J. Farrell *IBM Research Division, P.O. Box 218, Yorktown Heights, New York 10598.* Mr. Farrell joined the Research Division as a research staff member in 1968. His current research is on image reconstruction and multidimensional color presentation of images and data. Prior studies involved the application of simulation and statistical methods to evaluate and predict the likely recovery of critically ill patients. From 1966 to 1968, he was principal mathematician at UNIVAC, St. Paul, Minnesota. Mr. Farrell's work involved processing techniques and simulation related to spacecraft stabilization. From 1964 to 1966, he was a senior research scientist in the Research Division of Control Data Corporation. His research was concerned with feasibility and accuracy of different scanning optical sensors for celestial navigation. From 1959 to 1964, he was a mathematician in the Mathematics and Logic Research Department of UNIVAC, St. Paul. His assignments involved development of signal processing techniques related to sonar target classification and computer system reliability. Mr. Farrell received a B.S. from the University of Minnesota with a major in physics, and subsequently completed his doctoral course work in mathematical statistics. In 1982, he received an IBM Outstanding Technical Achievement Award for his work related to patient recovery.