CPG MEMO NO. 11

DISCUSSION OF THREE DIMENSIONAL VECTOR DISPLAY SYSTEMS

Written by Jim Torson
Version 3
April 13, 1982

1. INTRODUCTION

Various types of display devices are being considered for use with the "pipeline" mapping system. The basic purpose of the display system is to provide capability that will allow users to make effective use of the "pipeline" system. This will consist of displaying the maps that are produced. However, it will also include interactive display and editing of the visibility data.

There clearly is a need for an image display device for displaying the maps. Image display devices have been extensively investigated and a discussion of their features is contained is a separate document. We have chosen the IIS as the image display device. The current plan is to attach this display to a PDP-11/44 as the host computer. In the "final" configuration, this system will have shared disk access with the GRIDDER disk(s) on which finished maps are stored.

This document will discuss the possible use in the pipeline display system of refresh vector displays which can display perspective (and non-perspective) views of three dimensional data plots. (These display systems can of course also display the usual types of two dimensional data plots.) These displays are not being considered as a substitute for an image display. Instead, they are being considered for use in conjunction with the image display.

2. THE MAIN DIFFERENCE BETWEEN AN IMAGE DISPLAY AND A REFRESH VECTOR DISPLAY

In an image display such as the IIS system, the CRT beam draws the image on the screen by using a raster scan pattern. That is, the beam starts at the top left and scans along a horizontal line. It then flys back to the left edge and scans across the second line. It continues in a similar manner until the bottom horizontal line is drawn on the screen. This raster scan pattern is repeated at a regular rate such as 30 or 60 times per second. An image display system contains an image refresh memory which can be thought of as storing a two dimensional array of numbers. Each stored number determines the gray scale intensity or color of the corresponding dot or "pixel" on the displayed image. When the computer changes values stored in the refresh memory, corresponding changes occur in the displayed image. In most image displays, the stored pixel value does not directly determine the displayed intensity or color. Instead, the stored value is used as an index into a "look-up table," and the corresponding value in the table determines the displayed intensity or color for that pixel. This is extremely useful because reloading the relatively small look-up table allows dynamically changing the scaling from stored pixel values to displayed intensity or color.

An image display can of course display a picture that is composed of lines. All you have to do is store appropriate values into all the refresh memory pixels that make up the lines. However, showing a picture that is composed of many lines will involve a large number of pixels. Thus, it is not possible to rapidly change such an image.

In a refresh vector display, the CRT beam "randomly" moves around the screen to draw each line, dot or character of the picture in sequence. A piece of hardware called the display processor reads a "display list" from memory and positions the CRT beam accordingly. When it finishes the display list, it just starts reading it again so that the image on the screen is continually refreshed. The display list contains some control information to tell the display processor what to do, but most of it consists of data which specifies the coordinates of the endpoints of the lines to be drawn. Moving a line on the screen can be accomplished by simply changing the endpoint coordinates in the display list. Additionally, a picture composed of many lines can be described in terms of coordinates relative to an initial positioning command in the display list. Moving a complex picture on the screen can thus be accomplished by changing the coordinates in the initial positioning command. Thus, pictures can be moved or changed much faster than could corresponding pictures that are displayed on an image display system.

The refresh vector displays to be considered below have an additional important feature. The display list can contain a description of a three dimensional data plot by storing three dimensional line endpoint coordinates. A small list of parameters tells the display processor what view of this three dimensional object to display on the screen. Thus, changes to the parameters can easily result in showing a picture that is dynamically scaled, rotated, etc.

The above description is somewhat simplified and is a little inaccurate in some of the details. However, in general terms it correctly shows why a refresh vector display can do things which are quite impossible on an image display. Image displays have a number of additional features, some of which are described in another document. The additional features of refresh vector displays are described below.

3. GENERAL COMMENTS ON IMAGE DISPLAYS

When we proposed purchasing the first image display system for use at NRAO, many people questioned whether such a system would be useful enough to justify its cost. The users were accustomed to displaying their maps by making contour plots and they felt that that was sufficient for what needed to be done. However, most people would probably now agree that an image display system is well worth the cost. Experience has shown that an image display allows the user to easily see features in a map which can be difficult to see in a contour plot. An example of this is the stripes that can result from using bad visibility data points in mapping. Of course contour plots are still a very useful way to display maps. The important point is that contour plots and gray scale or color images are different ways to display the same data. One is most useful for some purposes and the other is most useful for other purposes.

When we display a radio map as a gray scale image, we are of course encoding different data values as different intensities on the screen. Interpretation of the displayed image makes use of the intensity resolution of the human visual system. However, it is important to recognize that the visual system cannot measure absolute intensities. For example, consider a room with white walls and a window in one wall as the only light source. When we look at the wall that is perpendicular to the wall that has the window, we are receiving a relatively high light intensity from a section of the wall near the window. A section of the same wall far from the window will be giving us a much lower intensity of light. But we perceive the entire wall as being white. It would be very difficult to estimate the relative intensities that we are receiving from different parts of the wall. And, an accurate estimate of the absolute values of the intensities would probably be totally impossible. In fact, a black object near the window can be giving us a higher light intensity than a section of the white wall that is far from the window. Of course it is very useful for our visual system to work this way. A "black" object should always be perceived as being black. Things would be very confusing indeed if an object sitting in a dimly lit corner of a room looked black and an identical object sitting near the window looked white.

This ability of the visual system to give a constant perception over a range of absolute intensities can be very useful in interpreting gray scale displays of radio maps. For example, the radio image of NGC6543 is easily perceived to be a number of overlapping filaments even though different filaments have different absolute intensities and a given single filament can have a varying intensity over its spatial extent. It can be much more difficult to see this pattern in a contour plot of this object. Incidentally, some of these differences between image displays and contour plots were not fully appreciated when we argued for the initial image display.

Our perception of color is certainly different from our perception of gray scale intensity. In many cases, different colors can be distinguished more easily than can different gray scale intensities. Thus, more detail is often visible in a pseudocolor display in which different map values are being encoded as different colors on the screen. Of course this may also make it more difficult to see a feature which is composed of a relatively large range of map values since its different perceived colors may get lost among the surrounding colors. In any event, our ability to perceive absolute colors is not nearly as accurate as may be supposed. The perceived color of a point is affected by what surrounds it, and a gradual and/or small change in color across an area may be difficult to detect.

Although changes in color or gray scale intensity over a certain range are difficult to perceive, very large changes can of course be easily seen. One of the main uses of the look-up table in an image display is to exaggerate the range of intensities or colors that is encoding a given range of map values. However, it is important to realize that when this is done, changes within other ranges of map values usually become more difficult to perceive. Since you don't know ahead of time which range of data values will be of interest, it is very important to provide the user with dynamic and interactive control over the look-up table contents. Also, there may be more than one range of data values that is of interest. In the case of radio maps, it is often adequate to examine these one at a time.

4. DISPLAYING VISIBILITY DATA ON AN IMAGE DISPLAY

We are planning to display visibility data on the image display system and to allow the user to interactively edit the data by pointing to displayed values that are to be flagged. One way that the data can be displayed is to load the amplitude (or phase) for a given correlator as a function of time into successive pixels values across a horizontal (or vertical) line in the image refresh memory. Different correlators will be stored in different lines (or columns). The data can then be displayed with the usual gray scale or pseudocolor encoding of stored data values. Another way to display the data would be to store amplitude in one image refresh memory and to store phase in another. A color display in which phase is encoded as hue and amplitude is encoded as intensity could then be produced. Although we do not have any experience with this type of display, it seems fairly certain that a bad correlator would show up as a horizontal line with an abnormally low or high value. In order for this to work effectively, it may be necessary to sort the baselines according to UV distance in order to minimize the normally expected fluctuations between adjacent lines on the display. Interference would probably show up as a vertical pattern on the screen since there would be a disturbance in a number of correlators at the same time.

Large problems with the data are expected to show up fairly clearly with this type of display. Of course it may be reasonably easy to write a program which would scan the data and automatically detect these large problems. The main motivation for displaying the data would be to make use of the enormous power of the human visual system and the human brain in detecting more subtle problems with the data. However, there may be several difficulties encountered when trying to see subtle problems when the visibility data are displayed in this way on an image display.

As described in the previous section, the human visual system cannot easily detect a small and/or gradual shift in gray scale intensity or color. Thus, a small change in the output of a correlator may not be visible on the screen.

Changing the look-up table contents so that a small range of data values is encoded with a larger range of gray scale intensities or colors can make it easier to see a small change in stored data values for a given correlator. However, it would then become more difficult to see similar changes in other correlators which are occurring over a different range of values. It would thus be difficult to see a pattern of changes that occurs in a group of correlators.

As described previously, displaying a map with pseudocolor encoding can increase the amount of detail that is visible. However, the human visual system has less spatial resolution for changes in hue than for changes in intensity. (That is why the standard color television broadcast signal has a larger bandwidth for the intensity part of the signal than for the hue part of the signal.) Pseudocolor often increases the detail visible in a displayed map because most map features are fairly broad in spatial extent and the lower spatial resolution for changes in hue is not important. However, the visibility data display is a totally different kind of image. The features of interest can have a much smaller spatial extent than the features in a map display. Some visibility data features may even involve only a single pixel. Thus, it may be necessary to use a zoom factor to make some features in the visibility display visible. Enabling the zoom factor of course reduces one of the advantages of using an image display system to display the visibility data, i.e., the ability to see a large amount of data at one time. This may make it harder to see overall patterns in the data.

As discussed previously, our visual system does not give us an absolute measure of gray scale intensity or color. Even relative measurements are difficult. Thus, we cannot tell much about the quantitative aspects of a change that is visible. It would not even be possible to tell whether the change was linear or non-linear. If two changes at different levels are visible on the screen, we cannot tell whether the rates of change or the amounts of change are the same. Changing the look-up table contents can make small changes in the stored data values more visible, but this can also make it more difficult to perceive quantitative aspects. It may then be more difficult to distinguish a normal level of fluctuation from an unacceptable level of fluctuation.

The human visual system is sensitive to changes in the spatial derivitive of the intensity variation. At places where there are changes in the derivitive, we perceive spikes in the intensity distribution which are not actually in the intensity image. This effect is call Mach banding. For map displays, this is not an important effect since the intensity varies rather smoothly in a map. However, this might cause significant distortions in our perception of visibility data. It could make quantitative estimates or comparisions more difficult, especially when zoom is enabled and each data value is displayed as a discreet block consisting of an array of pixels on the display screen. This effect might also cause us to "see" bad data points which aren't really bad. It might also cause our perception of what is happening with one correlator to be affected by what is happening on the adjacent correlators.

Our visual system is constantly trying to find "edges" in any image that we look at. This may help in detecting a change in a correlator's data value. However, such an "edge" would be very small compared to the "edge" between adjacent correlators. Our visual system may be so busy with the large sharp edges that it is harder than might be supposed to see the possibly fuzzy "edge" formed by a change in correlator value.

In summary, an image display system is very useful for displaying and interpreting the two dimensional data arrays which we call radio maps. Visibility data can certainly be put into a two dimensional data array and displayed as an image on an image display system. It is highly likely that this will be a useful way for the user to see some features of interest in visibility data. However, the resulting image differs from map images in some important ways. It may be a mistake to assume that such a method of display will allow the user to easily see all of the features in the visibility data that are of interest.

5. GENERAL COMMENTS ON VECTOR DISPLAYS

When we display data on a vector display device, we are generally encoding different data values as different spatial locations on the plot, e.g., different dot or symbol locations or different line endpoint coordinates. Interpretation of such a plot makes use of the spatial resolution of the human visual system. Unlike our intensity perception, our spatial coordinate perception is quite good at giving us an absolute measurement. It is even better at making relative measurements. When we reach for an object, it is relatively unimportant to know the absolute or relative intensity of the light comming from it, but it is very important to know where it is and to know where our hand is relative to it.

For a plot such as amplitude as a function of time, the range of vertical spatial coordinates in the plot encodes the range of data values being displayed and corresponds to the range of displayed gray scale intensities or colors in an image display. Changes in a sequence of plotted values of course need to be greater than some threshold value in order to be visible. Variations which are just above the threshold of detectibility can be expressed as a percentage of the full range of spatial coordinates in the plot. Variations of this same percentage in the intensity or color of an image display would not be visible. This is a direct result of the higher spatial resolution of the visual system versus its intensity or color resolution.

As was mentioned above, the visual system is constantly trying to find edges. That is why lines in a data plot can be seen so clearly. And, the visual system is very good at detecting the straightness of a line. A slight deviation from straightness can be easily detected. Also, the visual system is very good at comparing different lines. A slight deviation from parallelism between two lines can be easily detected. Patterns between lines which aren't straight can also be clearly seen. If several different things are plotted together, it will be easy to detect if one of them varies in a slightly different way.

When dots or other symbols are plotted for individual data points, our high spatial resolution allows us to easily see even a single data point which differs significantly from other points in the plot.

In general, a data plot on a vector display device can give us a much better understanding of quantitative aspects of the data. We can easily tell whether a change is linear or non-linear. We can easily tell if two different changes represent the same rates of change or the same magnitudes of change. We can get a good feel for the range of fluctuations in a group of data points. Points which lie outside of the normal range of fluctuations can be identified even if they exceed the normal range by a small amount. Adjacent data points have little effect on our ability to see quantitative aspects of a particular point.

There are of course limits in our ability to interpret data plots on a vector display. If changes in the data values are scaled into very small changes in spatial coordinates on the plot, then we won't be able to see the changes. Also, if the plotted symbols or lines are spread out too much, then we will not be able to see trends in the data. In any given situation there will be some range of scaling which allows us to see the features of interest. As with the scaling of data values into displayed intensities on an image display, we won't be able to tell ahead of time which scaling of data values into spatial coordinates will be appropriate. Thus, having dynamic and interactive control over the scaling can be very useful.

When we change the image display look-up table contents to emphasize a range of data values, it becomes harder to see changes in other ranges of data values. This is much less of a problem with data plots on a vector display since we can see small changes across a much wider range of the thing that is encoding the data values, i.e. spatial coordinate. This would make it much easier to see such things as similar patterns of small changes in correlator outputs for a set of correlators that are at different levels. We would be able to clearly see patterns which could not be seen on an image display with any amount of look-up table modification. On a vector display plot, we could even independently and dynamically change the scaling for each different correlator in order to more clearly see patterns at different levels for different correlators.

6. SOME TYPES OF THREE DIMENSIONAL DATA PLOTS THAT COULD BE USED FOR DISPLAYING VLA DATA

The following sections describe several possible types of three dimensional data plots. Some uses of these types of plots for various types of VLA data are suggested. Of course some of these may turn out to be more useful than others for displaying our data. Also, there are probably various other types of plots that could be tried and could turn out to be more useful for certain types of data.

6.1. THREE DIMENSIONAL SURFACE PLOT

Any two dimensional array of data could be displayed as a three dimensional surface plot. The X and Y dimensions in the plot correspond the the rows and columns in the data array. The height of the surface (Z coordinate) at each X-Y point is determined by the corresponding value in the data array. This type of plot is illustrated in Figure 1. In some situations, it may be more useful to draw lines only along the X direction or the Y direction. Another variation would be to just plot a dot at the location of each data point rather than connecting the data points with lines. This would probably make it a bit more difficult to see the "surface." However, different intensity dots could be drawn to add another variable to the display. Another variation would be to draw a line parallel to the Z axis for each data point as illustrated in Figure 2. This type of display might be useful for some purposes since a data value which differs greatly from its neighbors might be less distracting but still visible.

The three dimensional surface plot could of course be used for displaying VLA map data. This might be useful for simultaneously showing small map value variations at different locations within the total range of values in the map. This type of display might also be useful for displaying gridded visibility data. Another application might be for displaying spectral line map data in which one coordinate is velocity and the other is a spatial map coordinate.

6.2. MULTIPLE TWO DIMENSIONAL DATA PLOTS

Multiple two dimensional data plots can be drawn as a three dimensional data plot. This is illustrated in Figure 3. Each different plot lies on a plane of constant Y value. By looking at this plot along a line parallel to the Y axis and removing the perspective effect, we see a normal two dimensional plot in which multiple things are plotted on top of each other (see Figure 4). This could be useful for comparing the absolute levels of the various things being plotted. By moving the viewing location so that the view direction makes an angle with the Y axis, we could see the various plots stacked along the vertical dimension of the screen as illustrated in Figure 5. By looking parallel to the X axis as shown in Figure 6, we can compare the levels and ranges of variations in the different plots.

This type of plot can of course be used for displaying amplitude or phase as a function of time for a number of different correlators. Another use would be for displaying spectral line data. The different plots could be the bandpass for a given correlator at different points in time. Or, the bandpass for a number of different correlators at the same time could be plotted. Monitor data as a function of time could be plotted. Each different plot could correspond to a different antenna. This type of plot could also be used for displaying the antenna gains produced by ANTSOL.

6.3. COMPLEX VALUES AS A FUNCTION OF TIME

Figure 7 illustrates a technique that could be used for displaying a complex value as a function of time. At each point in time a line is drawn from the origin line to a point whose Y and Z coordinates correspond to the real and imaginary values. Successive points in time are also connected by lines. The display then shows a "ladder" which twists around the origin line. The length of a "rung" of the ladder corresponds to the amplitude of the data value and the angle of the "rung" corresponds to the phase. A set of complex values could be plotted by giving each one an origin line at a different Y-Z coordinate.

This type of display could of course be used for displaying visibility as a function of time for a number of different baselines. The locations of the different origin lines could be arranged in a matrix form as is done in LISTER so that all the plots in a given line or column correspond to baselines associated with a single given antenna. This is illustrated in Figure 8.

This type of display could also be used for displaying the complex antenna gains produced by ANTSOL.

6.4. UV TRACKS

Visibility data can be displayed as a three dimensional plot in which X and Y correspond to U and V. The Z value at which a point is drawn would correspond to the visibility amplitude or phase. Looking straight down parallel to the Z axis would then show what UV values have been sampled. Looking at some other angle would show the variation of amplitude or phase as a function of U and V. A variation on this display would be to connect the points for each correlator with a line to show the time sequence.

6.5. THREE DIMENSIONAL SCATTER-GRAM

The three dimensional location at which a point is plotted could be determined by the values of three different variables at the same point in time. Plotting a point for each different time sample would result in a plot such as that illustrated in Figure 9. Viewing the plot along lines parallel to the three axes would show relationships between the three possible pairs of values. Viewing from other angles would show relations among the three variables together. A variation on this type of plot would be to connect the dots with a line to show the time sequence as illustrated in Figure 10.

A possible use for this type of plot would of course be to see relationships among three different monitor point values.

6.6. ADDITIONAL VARIATIONS

In addition to the types of plots described above, a number of additional variations could be imagined that might prove useful. For showing multiple two dimensional plots, we could draw lines down to the Z = 0 plane. This might help to visualize which plane a particular plot lies in and thus help to distinguish the different plots. This technique might also be useful for the UV track display.

If we are plotting amplitude as a function of time for several correlators, we could draw phase information as small vectors that start at the plotted amplitude value and extend parallel to the Y axis. The size of the vector could correspond to its value and the direction could correspond to its sign.

If we are plotting averaged data, it might be useful to draw RMS or range bars.

Perhaps it would be useful to show both vertical RMS bars and horizontal phase vectors on a set of amplitude plots.

7. WORKING EXAMPLES OF THREE DIMENSIONAL DATA PLOTS

About four and a half years ago our first array processor was briefly connected to the PDP-11/40. I took this opportunity to quickly implement some test programs which use the array processor to show dynamically changing views of some three dimensional data plots on the VT-11 refresh vector display that is used by IMPS. An array processor is once again temporarily connected to the PDP-11/40, and some sample plots can be illustrated using various types of VLA data. These simple test programs can give you some feeling for what a three dimensional vector display could do. These programs illustrate one of the main advantages of such a system, i.e., the ability to show dynamically rotated and scaled views of a three dimensional data plot. However, it should be recognized that a real three dimensional vector display system would have significant additional advantages. The important capabilities of these systems are discussed in the following section.

8. CAPABILITIES OF THREE DIMENSIONAL VECTOR DISPLAY SYSTEMS

The following discussion will describe some of the important capabilities that would be provided by one of the new three dimensional refresh vector display systems that are currently on the market. This will be compared with the much more limited capabilities that can be demonstrated with the test programs which use the array processor and VT-11 display.

8.1. IMAGE COMPLEXITY

The VT-11 is limited to displaying about 2,000 points or lines at one time. A new system would be able to display images that contain 20,000 or 40,000 lines or points. For two dimensional data plots, up to 95,000 lines could be displayed on the screen. In addition, a more complex plot could be stored in the display system with the capability to dynamically roam around with the display "window."

8.2. SCREEN RESOLUTION

A new system would have a CRT spot size (line width) which is the same or smaller than the spot size on the VT-11. Also, the size of the CRT itself would be bigger. Thus, there would be an effectively greater screen resolution.

8.3. SCREEN PHOSPHOR

The VT-11 has a relatively long persistence phosphor. Thus, when the image moves on the screen, "streaks" are left behind which interfere with the three dimensional effect, especially when dots are being plotted. A new system would have a short persistence phosphor to eliminate the "streaks."

8.4. DEPTH QUEUING

The test programs draw all lines and dots at the same intensity. This makes it difficult to tell which line is which when the viewing orientation makes many lines overlap on the screen. A new system would have a feature called depth queuing in which parts of the image farther from the viewer are drawn with a lower intensity. This enhances the three dimensional effect and makes it easier to distinguish different lines.

8.5. THREE DIMENSIONAL CLIPPING

The test programs do not allow any clipping of the displayed image. However, a new system would allow dynamic three dimensional clipping. Conceptually, there would be a three dimensional box (or truncated pyramid for perspective viewing) that surrounds the data plot. Anything that lies outside the box is clipped, i.e. not displayed. The boundaries of the box can be dynamically and interactively varied so that parts of the plot that are not of current interest can be eliminated from the display. This can be very useful for temporarily eliminating parts of the plot that lie in front of or in back of the region of current interest.

8.6. IMAGE SCALING AND ROTATION

The test programs allow rotation of the data plot about only two axes. There are thus some viewing orientations which cannot be shown. A new system would provide complete flexibility in choosing viewing orientations.

The test programs consider the entire data plot to be a single three dimensional object and only allow rotation and scaling of the entire object as a whole. Scaling of the object is limited to scale factors which keep the entire object visible on the screen. A new system would allow a much greater range of scaling. In conjunction with the three dimensional clipping, a small part of a plot could be rapidly and dynamically magnified to fill the entire screen for close examination. The plot coordinates could be stored as 16-bit values. This would allow dynamic scaling throughout the full dynamic range of the original data.

A new system would allow independent scaling, rotation and translation of different parts of a plot. For example, the test programs allow scaling the entire plot of complex values as a function of time. However, a new system would allow you to do such things as changing the scaling of each plot relative to its local origin. This would effectively change the amplitude scaling while leaving the origin line of each plot in the same location. This type of capability might also be of use with multiple two dimensional plots. One of the plots could be dynamically translated to the plane of one of the other plots so that the two could be compared more easily.

8.7. MULTIPLE VIEWS OF THE SAME PLOT

The test programs allow displaying only a single data plot at a time. However, a new system would allow displaying multiple different plots or different views of the same plot. For example, part of the screen could display an overall plot of a large amount of data. Another part of the screen could then show some magnified view of a small part of the overall plot. A box could be drawn in the overall plot to show what part is being displayed in the magnified plot.

8.8. BLINKING AND HIGHLIGHTING

The test programs allow blinking of parts of the plot in a very limited fashion. A new system could allow much more flexible control over this. For example, in a plot of amplitude as a function of time for multiple baselines, the user could dynamically select an antenna number and then have all the baselines which use that antenna blink on the screen. This could help in seeing patterns that are associated with a particular antenna. Or, it may be more useful to change the displayed intensity of a part of the plot so that it is still visible ~~be~~ but can be distinguished from other parts of the plot. Using different line textures, e.g., dashed or dotted, may also be useful for distinguishing parts of a plot.

Another use of blinking would be to blink compare different things, e.g., amplitude and phase, different time periods, different baselines, etc.

8.9. AXES AND LABELLING

The test programs can only draw simple axes without labelling. Obviously any serious implementation on a new system could have flexible axes, labelling, and grid line options.

8.10. INTERACTIVE POINTING AND DATA EDITING

A new system could allow the user to interactively point to parts of the plot and get a numerical read-out of data values, e.g., which correlator is being pointed to, what amplitude, what time, etc. Interactive pointing could also be used for interactive editing of the visibility data. The user could point to bad data points that are to be flagged. Or, he could point to ranges of data points to be flagged.

8.11. CONTROL OF INTERACTIVE OPERATIONS

The test programs use the data tablet for all interactive control of the image. The different scaling and rotation functions are selected one at a time by typing commands on the keyboard. A new system could have a control panel consisting of an array of knobs and function switches and buttons. Turning two knobs at the same time could then do such things as simultaneously changing the rotation and scaling. These things would be standard system options that are supported by the standard software package.

8.12. SELECTION OF DATA FOR DISPLAY

The test programs were meant to test the visual aspect of dynamic three dimensional data plots and thus do not include a flexible data selection capability. Any serious implementation on a new system would of course provide this capability.

8.13. COLOR

Some of the new systems that are available can display different colored lines. This could be useful for distinguishing different parts of a data plot or for adding another variable to the displayed information.

9. A FEW TECHNICAL IMPLEMENTATION DETAILS

The display list for the VT-11 is stored in the PDP-11 main memory. This uses up program address space and also real memory. It also uses memory cycles to constantly refresh the display. A new system would have its own memory for storing the display list.

Some new systems can locally handle inputs from the interaction devices and produce corresponding changes in the displayed image. This would reduce the amount of CPU processing needed for standard interactive functions.

Some new systems have a display processor which can handle two independent display screens. Thus, adding a second user capability would not involve a doubling of the price.

10. SUMMARY AND CONCLUSIONS

When data is displayed as a gray scale or color image, the intensity or color resolution of the human visual system is used in interpreting the display. When data is displayed as an image composed of lines, dots or characters, the spatial resolution of the visual system is used in interpreting the display. Both of these types of displays are very useful. However, since they make use of different parts of the visual system, they will allow the user to see different aspects of the data. A refresh vector display which can display three dimensional data plots would thus be a highly desirable tool to use in conjunction with an image display.

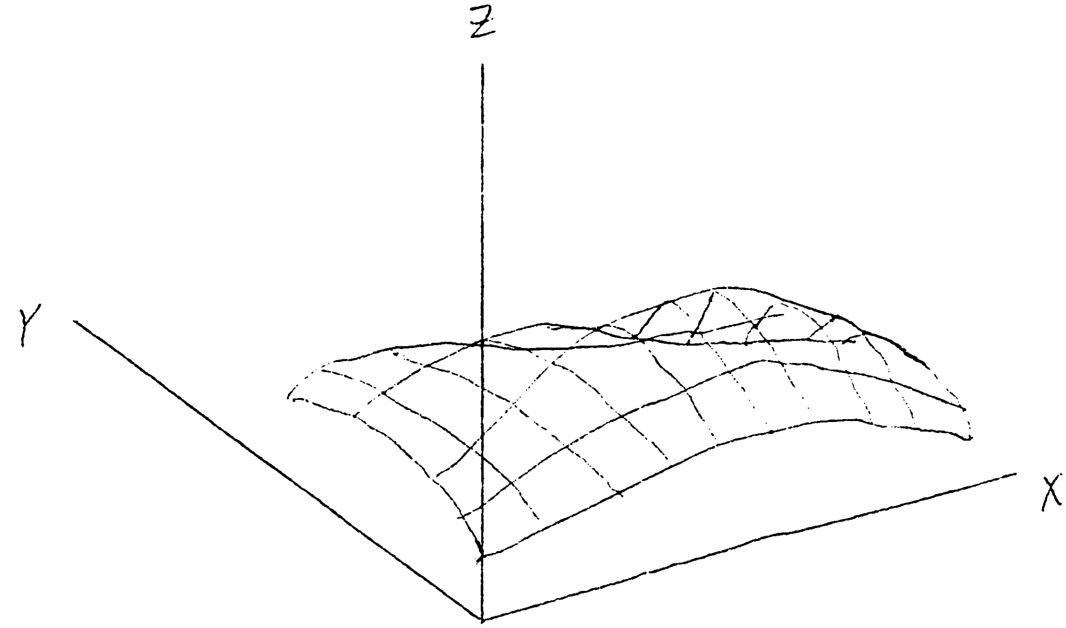

Figure 1. 3-D Surface Plot.

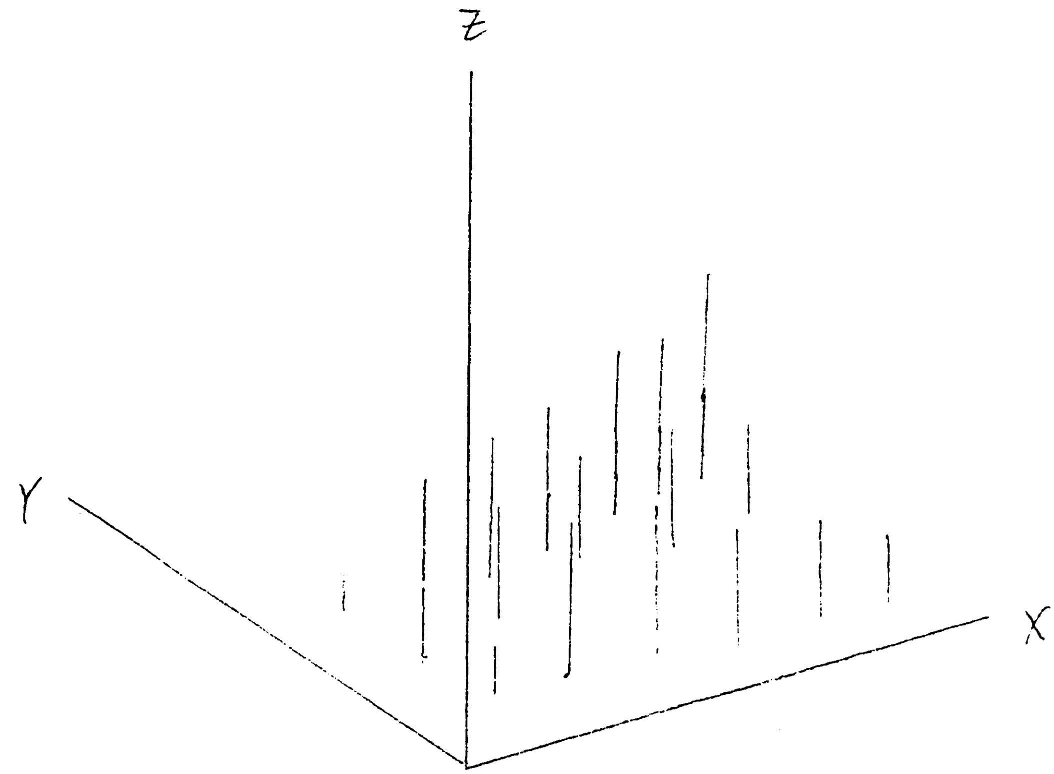

Figure 2. Variation on 3-D surface plot.

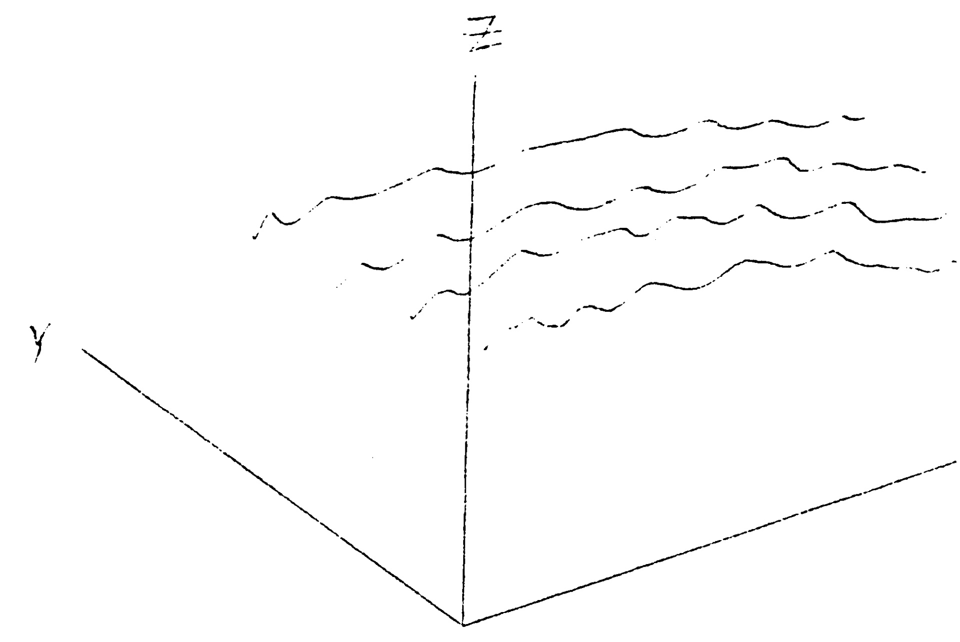

Figure 3. Multiple 2-D Data Plots.

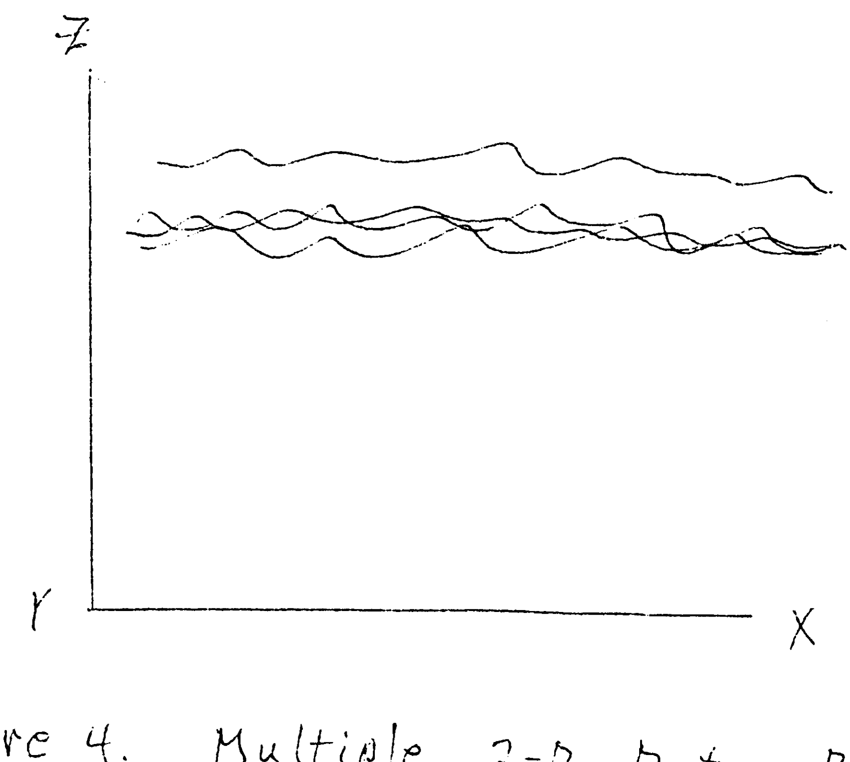

Figure 4. Multiple 2-D Data Plot on Top of Each Other.

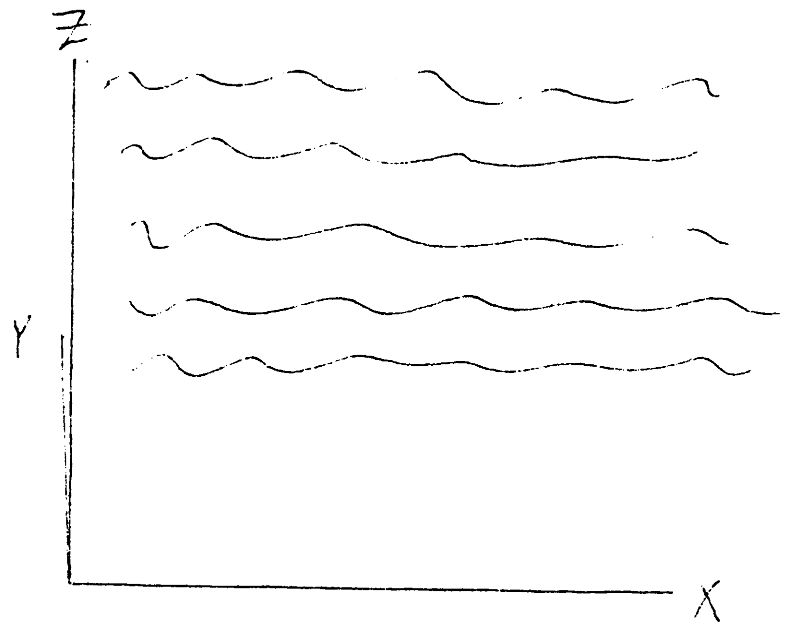

Figure 5. Multiple 2-D Data Plots Stacked Vertically Across the Screen

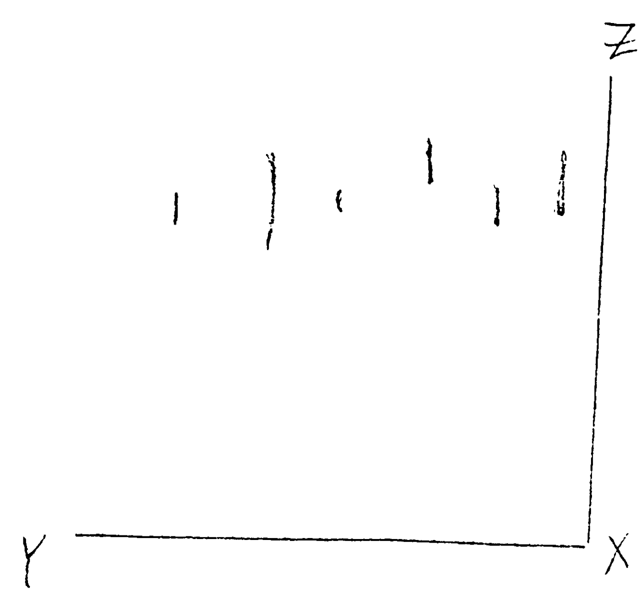

Figure 6. Multiple 2-D Data Plots Viewed to Compare Ranges and Leve

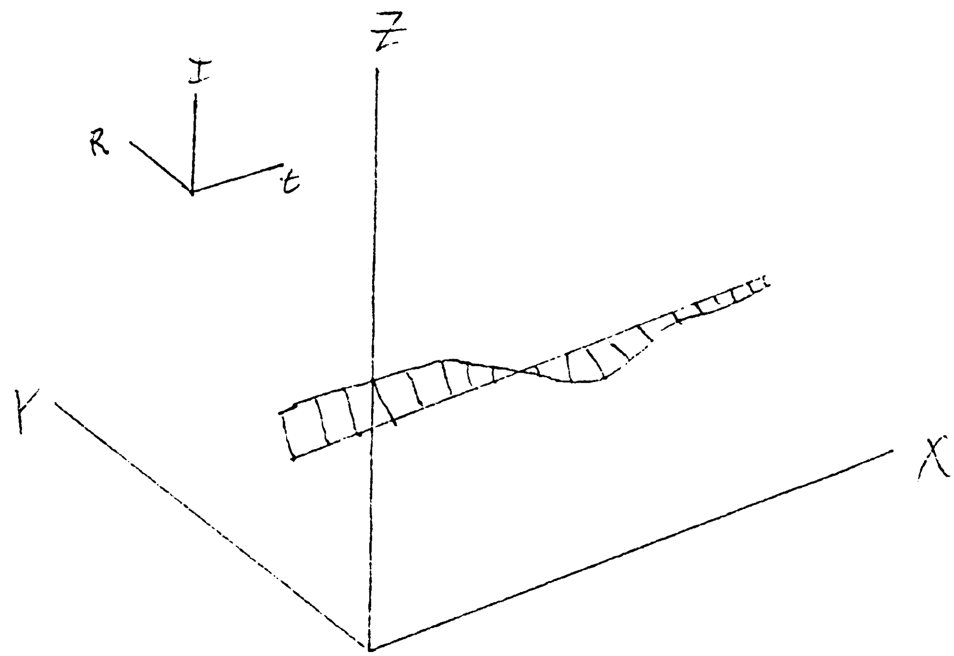

Figure 7. Complex Value as a Functio
of Time

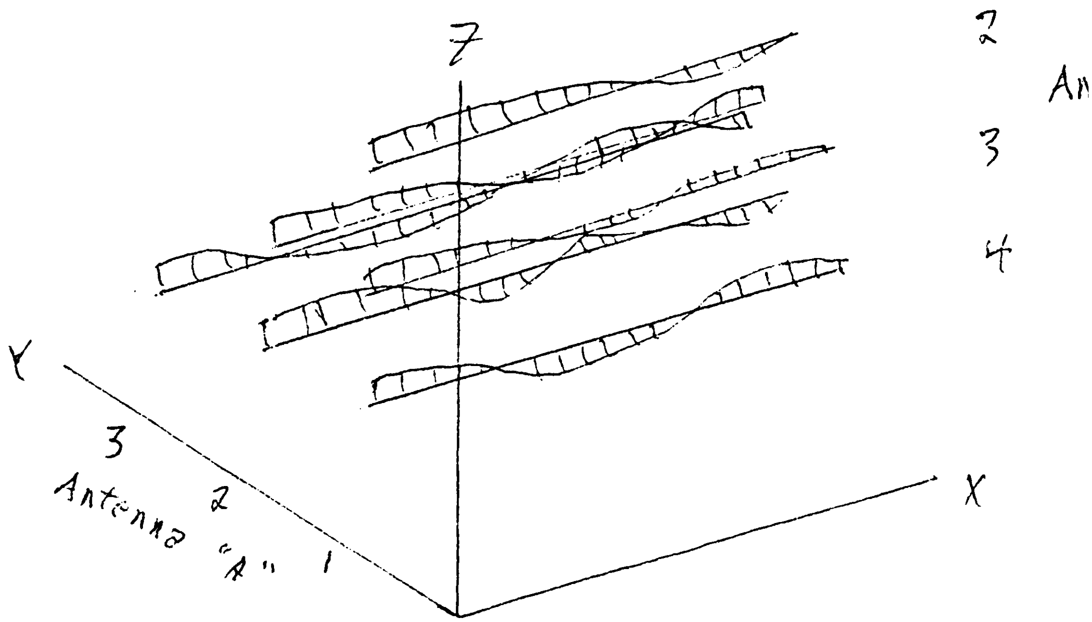

Figure 8. Multiple Complex ~~Vales~~ Values
Functions of Time

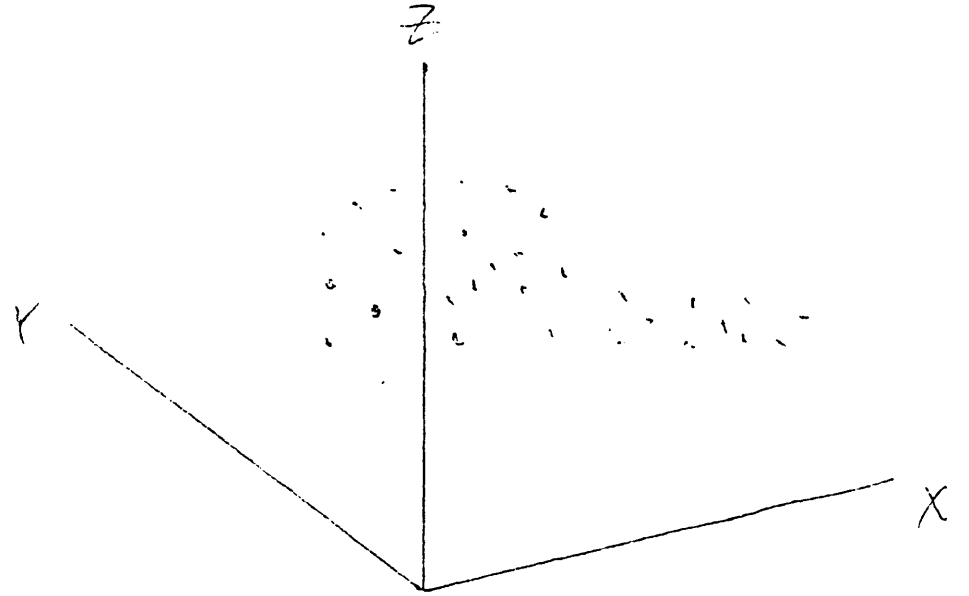

Figure 9. Comparison of Three Variables.

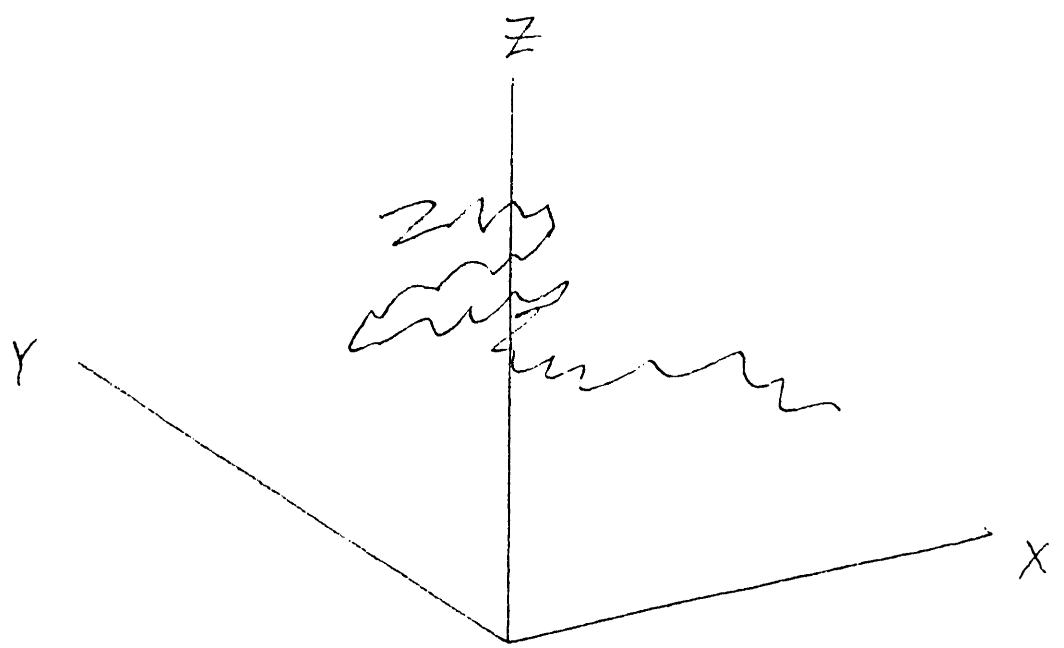

Figure 10. Comparison of Three Variables Plus Time Sequence.