8125600

LIPSCOMB, JAMES SARGENT

THREE-DIMENSIONAL CUES FOR A MOLECULAR COMPUTER GRAPHICS SYSTEM

The University of North Carolina at Chapel Hill PH.D. 1981

University Microfilms International 300 N. Zeeb Road, Ann Arbor, MI 48106

PLEASE NOTE:

In all cases this material has been filmed in the best possible way from the available copy. Problems encountered with this document have been identified here with a check mark __√__.

1. Glossy photographs or pages __✓__
2. Colored illustrations, paper or print ______
3. Photographs with dark background __✓__
4. Illustrations are poor copy ______
5. Pages with black marks, not original copy ______
6. Print shows through as there is text on both sides of page ______
7. Indistinct, broken or small print on several pages __✓__
8. Print exceeds margin requirements ______
9. Tightly bound copy with print lost in spine ______
10. Computer printout pages with indistinct print ______
11. Page(s) __________ lacking when material received, and not available from school or author.
12. Page(s) __________ seem to be missing in numbering only as text follows.
13. Two pages numbered __________. Text follows.
14. Curling and wrinkled pages ______
15. Other ____

THREE-DIMENSIONAL CUES FOR A MOLECULAR COMPUTER GRAPHICS SYSTEM

by

James Sargent Lipscomb

A Dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Computer Science.

Chapel Hill
1981

Approved by:

Adviser: F. P. Brooks, Jr.

Reader: H. Fuchs

Reader: W. V. Wright

P. Calingaert

D. L. Parnas

JAMES SARGENT LIPSCOMB. Three-dimensional Cues for a Molecular Computer Graphics System. (Under the direction of DR. FREDERICK P. BROOKS, JR.)

ABSTRACT

The perception of depth in three-dimensional objects presented on a two-dimensional display can be enhanced by many techniques. This thesis develops programs for implementing a number of such techniques flexibly and efficiently, including rotation, intensity modulation, and stereo.

These techniques are applied to the GRIP-75 molecular computer graphics system with which crystallographers fit proteins or nucleic acids to their electron density maps. A fragment of a molecule is superimposed on a contoured electron density map and presented to the user for manual manipulation.

My thesis is that motion decomposition pervades surprisingly many aspects of our computer-graphics-guided manual manipulation, that smooth rotation appears to be the single best depth cue currently available, that a variety of multiple image effects, due to disparity between refresh and update rates, can plague moving pictures, that the good 3-D perception provided by stereo can hinder manipulation, and that achieving orthogonality among functions is surprisingly important and difficult.

The programs run on a satellite graphics system consisting of a DEC PDP-11/45 computer driving a Vector General display system, and a high-speed selector channel which provides communication with a time-shared host computer, an IBM System/360 Model 75. The programs developed are written in PL/I and reside mainly in the satellite.

To My Family

ACKNOWLEDGEMENTS

I wish to thank Dr. Frederick P. Brooks, Jr., for creating the GRIP-75 project of which this thesis is a part. He and Dr. Victor L. Wallace advised my work. Other members of my committee, Dr. James D. Foley and Dr. William V. Wright, coordinated software development. Dr. Wright also built an early version of GRIP on which the current version was based.

I must also thank the many student members of the GRIP-75 team. I particularly thank E. G. Britton and M. E. Pique. Just as much of their work overlaps mine, so do many of their ideas.

I thank P. J. Kilpatrick and M. Hollins for several stimulating discussions, and W. Siddall for maintaining and improving my programs. I thank the many computer graphics and psychology professionals who guided me through the literature. I thank E. G. Britton, M. L. Hamilton, S. M. Pizer, and the members of my committee for their criticism of this thesis in draft. The text was edited on CALL-OS (IBM) and formatted by SCRIPT (University of Waterloo).

The GRIP-75 molecular graphics system was built by the University of North Carolina Department of Computer Science with the help of many users. Major system contributions were made by E. G. Britton, J. S. Lipscomb, and M. E. Pique, under the direction of W. V. Wright; and J. E. McQueen, Jr., under the direction of J. Hermans. F. P. Brooks, Jr., is Principal Investigator. The system development has been supported by NIH Biotechnology Research Resource grant #RR00898, NSF grant #GJ-34697, AEC contract #AT-(40-1)-3817, and the IBM Corporation.

Copies of this thesis are available from:

Department of Computer Science
New West Hall, 035A
University of North Carolina
Chapel Hill, NC 27514

University Microfilms Inc.
300 N. Zeeb Road
Ann Arbor, MI 48106

CONTENTS

LIST OF FIGURES

LIST OF TABLES

Chapter 1

INTRODUCTION

1.1 OVERVIEW

Computer graphics can be useful in many areas of design. One common theme is "construction". The user wants to build a picture that represents some characteristics of a system or physical object.

The large-molecule crystallographer wishes to construct a three-dimensional (3-D) atomic-detail model of the structure he is studying. He has as given the structural elements, units strung in a chain each with a few atoms, which he must position in space guided by a rough three-dimensional map of the electron density. This requires good three-dimensional perception.

Many techniques exist for the presentation of depth in three-dimensional objects drawn on a two-dimensional (2-D) display. This thesis analyzes the most promising of these techniques traditional in molecular applications, and describes their implementation on the GRIP-75 system. This software is designed to facilitate productive work by chemists, to provide a base for the investigation of the human factors requirements of each technique, and to provide a comparative evaluation of these depth cues.

Effective tools for three-dimensional perception and manipulation using computer graphics are not straightforward to build. Many of the issues are subtle but important, and sometimes solutions that were intuitively obvious turned out not to be the best.

1.2 STATEMENT OF THESIS

My thesis is that the following statements about the design of three-dimensional cues for a molecular computer graphics system can be supported by experimental observation:

1. motion decomposition [Kilpatrick, 1976] pervades surprisingly many aspects of our computer-graphics-guided manual manipulation;

2. smooth rotation appears to be the single best depth cue currently available;

3. a variety of multiple image effects, due to disparity between refresh and update rates, can plague moving pictures;

4. the good 3-D perception provided by stereo can hinder manipulation;

5. achieving orthogonality among functions is surprisingly important and difficult.

The chapter on results gives these and other conclusions in an expanded form with page number references to more detailed discussion. However, orthogonality, the last point mentioned above, permeates this entire document. Its influence in my work is so widespread and so important that I shall introduce it now.

1.3 ORTHOGONALITY

It appears to be both more important and more difficult than heretofore believed, to build related three-dimensional display functions in molecular computer graphics so that they appear to be mutually orthogonal.

I use the term _orthogonal_ to describe a relationship between three-dimensional display techniques that have three properties:

1. the user may invoke either independently by independent controls;

2. invocation of one technique does not change the perceived effect of the other;

3. user-controlled variation in parameters for one tehnique does not change the perceptual effect of the other.

This concept can be generalized from pairwise orthogonality to a system of mutually orthogonal features. The detailed discussions of orthogonality in this thesis, however, will generally be cast in terms of pairwise orthogonal features.

The value of orthogonality will be demonstrated by two parallel discussions. One addresses the topic directly, and the other addresses the implied assumption that related 3-D display techniques are needed at all.

What was heretofore believed about three-dimensional display techniques in molecular computer graphics is summarized in a chapter on related work. The reader will see that rather little has been concluded on the issue of orthogonality. Then, after presenting a history of the GRIP-75 system, several chapters of discussion are given to the human factors of trial configurations of three-dimensional display techniques, the difficulties (mainly arising from lack of orthogonality) that were encountered in smoothing the human factors, the work that was required to overcome these difficulties, and the correlation between the orthogonality and the utility of the three-dimensional display techniques as ultimately implemented. This correlation is then illustrated in a table of users in the chapter on results.

The subject of orthogonality is hardly worth discussing if just one or two three-dimensional display techniques will do. I originally intended the wide variety of techniques implemented to be a tool for exploring their relative utility, with an eye towards selection of a small, good set. I was surprised to find that most users are, indeed, best served by offering them a rich menu of redundant techniques. I demonstrate this by presenting the literature on related work in which the reader will note unresolved disagreement on the intrinsic value of several three-dimensional display techniques. Then in the chapters that cover individual techniques, I show how sharp differences in personal preference, which are made clear by a large user community, may account for a great deal of disagreement about the utility of different three-dimensional display techniques. Finally, the chapter on results shows, in a table of users, these differences in personal preference as perturbations on (or as noise in) the correlation between the orthogonality and the utility of the three-dimensional display techniques as ultimately implemented.

Achieving orthogonality can surprisingly be hard. There being no literature thus far on orthogonality in molecular computer graphics, I must restrict the argument to showing instances where the reader will be surprised that achieving orthogonality is not straightforward. Subtle complications arise in the interaction between rocking rate and amplitude, in the manual control of relative fading, in the "rotation" needed for stereo disparity, and in the use of the light pen in stereo.

1.4 SCOPE OF THESIS

The GRIP-75 system's goal was "early" biochemical usefulness. This confined the range of three-dimensional techniques implemented for this thesis, to those traditional in molecular computer graphics, and it necessitated close attention to human factors.

I made no attempt to try new or experimental three-dimensional techniques. Further, I avoided those traditional molecular graphics display techniques that are unsuitable for our display unit.

As in GRIP-71 [Wright, 1972a], I represent a molecule's bonds as sticks connecting atoms.

Wright [1972a] found perspective to be unhelpful with molecular structures, a result consistent with the psychology literature reviewed by Braunstein [1976]. Surface representation of molecules and hidden-line removal combined with smooth rotation on our display were too demanding for early implementation.

1.5 VARIETY OF CUES

1.5.1 Three Classes of Depth Cues

The techniques traditionally employed for presenting depth in molecular computer graphics fall into three classes:

1. smooth motion;

2. intensity;

3. stereo.

1.5.2 Variety Within Each Class

GRIP-75 provides a number of cues in each class, giving the user a variety of tools for three-dimensional perception.

Our users can manually move the image smoothly with three devices. An articulated linkage sets the overall viewing direction with two nested rotations. The automatic rotation cues are sinusoidal rocking and two forms of spinning about the vertical axis. One option is a continuous spin at a smoothly-controlled rate, the other is a smooth 90-degree turn. Several manipulation techniques are provided that give some depth perception during their use. A 3-D translation joystick moves a "jack"-shaped position-indicator (cursor) or a molecular substructure about in three dimensions. A 3-D spring-return device allows three degrees of rotational freedom for the substructure it controls. Together with the 3-D translation, it allows six degrees of freedom of motion for the substructure. Manipulation internal to the substructure is described by Pique [draft].

Three uses of intensity provide structural cues to the third dimension. All are controlled by adjacent sliders. The first is a relative fading control which fades out the molecule in one half of its range and the map or a second molecule in the other half. The second is intensity depth-cueing which fades the more distant objects in the picture. Finally, a "curtain of invisibility" can be swept through the picture. This simulates walking into the molecule. The curtain is parallel to the screen and only objects behind it are visible.

Stereo is a good depth cue for people able to see it. GRIP-75 has three forms of stereo: rotating-shutter stereo using a Bausch and Lomb tubular viewer to which the display is synchronized, side-by-side stereo which is familiar to chemists, and up-down mirror stereo employing a half-silvered mirror and some polarizing material.

Rocking and spinning are mutually exclusive, as are the three forms of stereo, but in all other respects the three-dimensional features can be used in any combination. They are compatible with all manual manipulations and, except for intensity control, the light pen.

1.6 <u>WHAT THE USER SEES AND FEELS</u>

The three-dimensional features covered in this document are part of a system for determining the structure of large molecules. An outline of a typical session follows, with typical values for all parameters.

The user sits at a workstation (Figure 1.1), signs-on, and asks for his personal library of molecules and electron density maps, which appears as a list of his molecules. He chooses one of them with the light pen, and presses the button for three-dimensional viewing. The selected molecule and the set of electron density maps associated with his library are now ready for display.

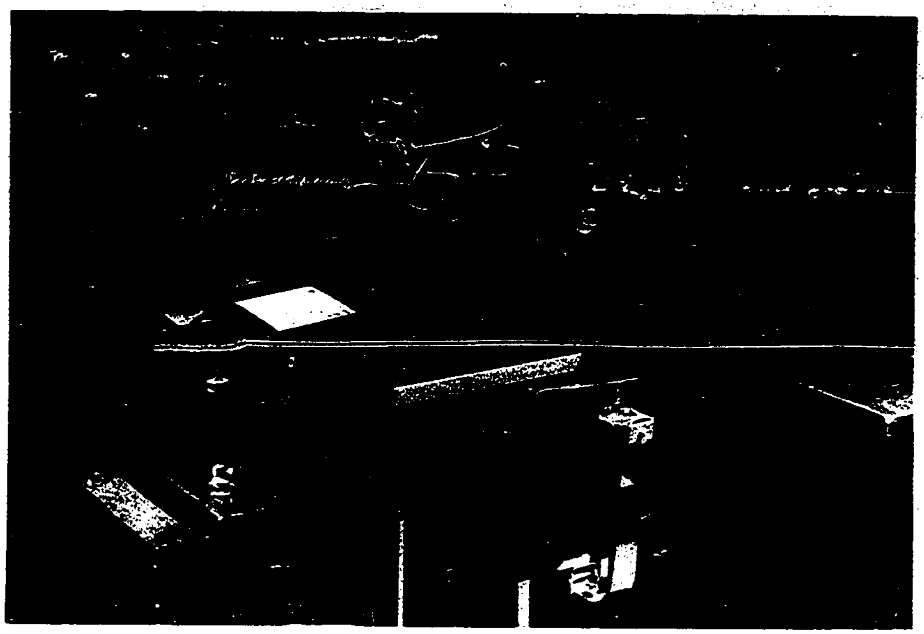

Figure 1.1: Workstation

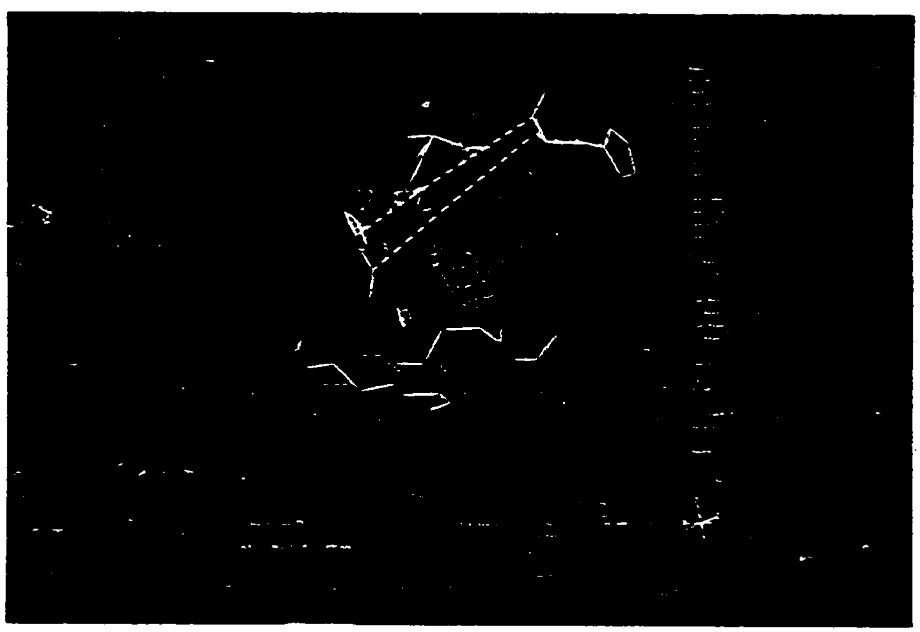

Figure 1.2: Fitting Closeup

After interactively centering and scaling for the desired view in a cubical region of space, usually selected to be about 18 angstroms on a side, within the molecule, he selects the electron density maps he wishes to see and the electron density level to be contoured.

When the map appears, he increases the scale to about a 12 angstrom cube, selects a substructure of the molecule with the light pen, and then selects a manipulation command. The system "breaks" the substructure loose as a unit from the surrounding molecule and sets it to translate and rotate by joystick control with its interior bonds twistable by knobs (Figure 1.2).

The user positions the substructure with joysticks and then selects the "accept" verb or the "cancel" verb to accept or to reject the new fit. When he is finished fitting substructures, he presses the button for display of his molecule library and types a "save" command.

1.7 TERMINOLOGY

Interactive computer graphics and x-ray crystallography both have their own particular set of terms. This section explains some of the more important terms dealing with both fields of endeavor.

The type of person for whom GRIP-75 was built is called the user, the client, the chemist, or the biochemist. Crystallographers form the largest group of these biochemists.

Rotation of the entire 3-D portion of the picture, simply to see it better, is called a viewing motion. This motion changes the viewpoint from which the structure is seen. It is therefore sometimes called viewpoint-rotation.

The small portion of the molecular structure which one is free to move about relative to the rest of the structure is called the substructure. Since this is usually an amino acid residue of a protein, or a nucleic acid residue (nucleotide) of a nucleic acid, it is often called simply the residue.

The experimentally determined data which guide the placement of the molecular structure represent a three-dimensional scalar field of electron density. The contoured representation of this data presented to the user is called the contoured electron density map. It is also more loosely described as an electron density map, a density map, or simply a map. The act of placement is called fitting the substructure to the map.

The classical modeling device for this work employing brass rods for the molecule and transparent sheets for the map is called a Richards box [Richards, 1968]. A computer graphics system, such as GRIP-75, that simulates this device is called an electronic Richards box [Collins et al., 1975].

Pique [draft] and Britton [1977] have established the following hierarchy of dynamic, interactive, and batch operations:

1. Dynamic operations are those that appear to be flowing and continuous, even if not quite smooth. Examples of dynamic operations in a car are the steering, brake, and gas controls. Examples in this molecular graphics system include viewpoint-rotation, substructure rotation and translation, intensity control, etc.

2. Interactive operations are in the form of discrete responses to discrete commands. In a car these include a manual gearshift, the lights, and the horn. Our molecular graphics system includes changes in stereo mode, contour level, position within the model, and scale of display.

3. Batch operations are not changeable during an interactive session. With the car example these include the minimum turning radius, the amount of water in the cooling system, and the number of tires. In the GRIP-75 system these include the maximum size of molecule and density map that can be loaded, the connectivity of the molecule, the interactive commands available, and the allocation of tasks to, and the limits of action of, most of the dynamic controls.

A distributed processing computer system can be organized as a network with a time-shared mainframe computer and an attached dedicated mini-computer. When the mini-computer is then coupled to a graphic display, we have what is called a satellite graphics system. The mainframe is the host, and the mini-computer is the satellite. In our system the host computer is an IBM System/360 Model 75 and the satellite is a DEC PDP-11/45.

The graphic display is called a display unit or simply a display. It is issued control information and a series of drawing instructions which it uses to draw a picture. A collection of these data and instructions is called a _display list_ or a _display program_. In GRIP-75 the display lists are constructed from the molecule and density data by the host in response to interactive commands.

When the display program is executed once, the picture is drawn once on the screen of the display. This is called a _refresh_. Because the image fades with time, the refresh must be repeated regularly. In GRIP-75 the refresh happens about 30 times per second depending on the complexity of the picture.

Dynamic motion is achieved by making a sequence of small changes in the position or orientation of an image. Each of these changes is called an _update_. GRIP-75 can execute about 10 updates per second depending upon the number of independently movable objects.

The coordinate system fixed in the laboratory is called the _user's frame of reference_ or the _user space_ (the screen coordinate system in Newman and Sproull [1973]). The coordinate system that rotates with the viewpoint-rotation is the _viewing frame of reference_ or the _viewing space_ (the object coordinate system [Newman and Sproull, 1973] of the outermost nested rotatable object). A rotation or translation is said to take place in one of these spaces when the axes of the joystick controlling the motion maps onto axes fixed in that space (e.g. viewing-space translation).

When an axis defined in the viewing space is set parallel to an axis defined in the user space, the definition of the viewing-space axis has _compensated_ for having been defined in a rotated space. We call the calculation relating the two, _axis compensation_.

The stereo effect is achieved by delivering a slightly different picture to each eye. The images differ by a rotation about the vertical axis. The angle of rotation is called the _stereo disparity_. It is also known as the _stereo angle_, the _stereo separation_, or the _amount of parallax_.

The interactive act of designating an item on the screen with the light pen is called _selecting_, _picking_, or _hitting_.

The vocabulary and grammar of meaningful interactions is the _interaction language_, or the _command language_, of a system. An _action language_ includes the dynamic motions as well. Where the language is limited to sequences of selections of objects and commands from a menu, it is called a _button language_.

When the deflection of a joystick causes a corresponding deflection of an image, the joystick is a position device. A velocity device is one whose deflection is proportional to the rate of motion of the image.

Manual motions (those controlled by a joystick) are dynamic. The image may also be made to move continuously without human intervention. I call such dynamic motion automatic motion, to contrast with manual motion. For example, the computer may rock the image back and forth or spin it. Manual motion using a velocity device has a strong automatic-motion character that makes such motion hard to classify.

Chapter 2

RELATED WORK

2.1 INTRODUCTION

Much of the work described in this thesis either builds upon earlier work done elsewhere or can be compared to later efforts by other builders. Only work pertinent to molecule or electron density map display and to dynamic motion is described. A comparison of these features in several electronic Richards box systems is also presented.

2.2 PERCEPTION

2.2.1 Background

The standard literature on perception includes many books and articles that directly address topics in this thesis. These are referenced in the body of this work.

However, some of the standard works provide valuable background information in the field of visual perception. They are well worth reading, although for one reason or another I do not reference them elsewhere. I commend these books to the reader's attention (see bibliography):

1. Boring [1942];
2. Cornsweet [1970];
3. Gibson [1953];
4. Graham [1965];
5. Helmholtz [1910];
6. Johansson [1950a];
7. Kling and Riggs, et al. [1971]
8. Kolers [1972];

9. Michotte [1946];

10. Monty and Senders [1976];

11. Spigel [1965];

12. Wallach [1976];

2.2.2 Wallach and O'Connell [1953]

Wallach and O'Connell [1953] coined the term *kinetic depth effect* (KDE) to describe the phenomenon by which a 2-D projection of a rigidly rotating 3-D object is perceived as 3-D. The wire frame models used in a number of their experiments resemble strikingly, though inadvertently, a variety of simple stick-figure molecular models. The orthographic projection of these wire frame models onto a screen shows black lines on a white background, but when a luminous rod is substituted with no projection (white lines on a black background) nothing important to the kinetic depth effect is changed.

They conclude that the kinetic depth effect arises from simultaneous change in both length and direction of the projected lines.

> Both these changes must be given together. A change in length alone is not sufficient to produce a reliable kinetic depth effect.

Wallach and O'Connell show that only the relative motion of the endpoints of the lines really matters.

> It is apparent that the kinetic depth effect will readily yield a perception of a rigid spatial arrangement of unconnected objects.

It is not important that these points are connected by lines, but if they are, loss of the kinetic depth effect occurs where the lines are clipped by the edges of the display, giving anomalous motion to the perceived endpoints. It is also not important that the lines touch each other to form vertices.

Subjects adapt to prolonged exposure to the kinetic depth effect. This facilitates the ease with which difficult structures may be seen as three-dimensional by use of the kinetic depth effect. However, fatigue also sets in:

> Where the kinetic depth effect takes place, a rigid three-dimensional form in rotation is seen instead of the distorted two-dimensional figure given on the shadow screen. The distorting two-dimensional shape may occasionally be _seen_ after prolonged exposure of the same kinetic presentation, or when one looks from a sharply oblique direction at the screen. The [subject's] experience of a continuously flowing form seems to him abnormal and unusual although this is exactly what is given on the retina.

2.2.3 _Wallach, O'Connell, and Neisser_ [_1953_]

Wallach, O'Connell, and Neisser [1953] continued the kinetic depth effect (KDE) experiments just described. They found that after motion had ceased, subjects continued to see the stationary projections as three-dimensional. Wallach, O'Connell, and Neisser call this the _memory effect of the KDE_.

> . . . a retinal pattern, which at the outset is seen as a plane figure, gives rise under identical external conditions to the perception of a three-dimensional form after an intervening exposure of the same pattern given under conditions which cause it to be seen as three-dimensional.
>
> . . . .
>
> It appears . . . that the aftereffect can be obtained virtually undiminished after longer time intervals [days], and so should be termed a memory effect.
>
> . . . .
>
> The influence which causes the perception of three-dimensional form in the stationary test exposure appears to come from a previous three-dimensional perception of the _same_ figure only.

The only "exception" to this last point is that subjects tend to interpret ambiguous stationary structures as identical to three-dimensional structures seen before by the kinetic depth effect.

2.3 HUMAN FACTORS

2.3.1 Woodson and Conover [1964]

Woodson and Conover [1964] discuss device layout and labeling. They suggest:

1. Controls should be within easy reach.
2. When reaching for a device the user's arm should not obscure the display.
3. Clothing should not catch on near devices when reaching for far devices.
4. Similar controls should be grouped together.
5. Control motion should be compatible with the expected response of the image.
6. "On" is up or to the right, "off" is down or to the left.
7. Labels should be placed along the direction of motion of the control.
8. The control should never cover a label.
9. Labels should be words rather than symbols.
10. Labels should be lettered horizontally.
11. Labels should be in all capital letters.
12. The type font for the labels should be simple.
13. Labels should be well separated from each other.
14. Labels should be close to their own devices.
15. Labels should be at a uniform distance from their devices.
16. Labels should be easily read from a normal operating position.
17. The label should say what the device does, not what the device is (ALTITUDE not ALTIMETER).

Our device layout follows this scheme with some necessary compromises. Labels had to be placed below some buttons to be visible from a normal viewing position. For z-clipping "on" (picture visible) is "down" (towards the user) to present the expected response (no clipping). "Off" (no picture) requires pushing back the "curtain of invisibility" away from the user until it covers all of the picture. Relative fading and intensity depth-cueing then had to follow suit for consistent action.

One aspect of device labeling that I did not consider is that at low light levels, white letters on a black background may be superior to black letters on a white background [Shurtleff, 1967]. A summary of Shurtleff's results may be found in Kriloff [1976].

2.3.2 Hansen [1971]

Hansen [1971] states some important ideas for the design of graphic systems.

The system should minimize the memorization required of the user; for example, the user should select functions from a menu rather than enter them from memory. The items selected should be labeled with names not numbers, and the resulting behavior should be predictable. In fact, if the hand-eye coordination is sufficiently close and natural, no conscious recall is required. The interaction takes on a lexical nature, and one can almost say that one's "muscles remember" how to perform it.

Engineering for user errors demands that all actions be reversable. Engineering for system errors demands that more than one means to any given end should be built in.

2.3.3 Foley and Wallace [1974]

Foley and Wallace [1974] discuss the role of the system designer who specifies the interaction language:

> One of the principal goals of interactive systems, graphic or otherwise, is the symbiosis between man and machine. When a terminal user is able to interact with a computer so that he is unaware either of the computer or of the medium of communication, this interaction can be said to be conversational. Then the capabilities of the two partners, man and computer, become as one working together on a single task.

> While this ideal is rarely completely attained, it represents the goal which, through skillful design of the machine and the medium of communication, can be met in large measure today.
>
> . . . .
>
> The language in computer graphics communication is not one of spoken or even written words, but rather one of pictures, and actions such as button pushes, lightpen intractions, and joystick movements which serve as words.
>
> The language and context of the conversation must be the language of the man and must be natural to him. . . . In our case, either the computer graphics system must "learn" the user's language, or the user must learn the computer's. Clearly the former should be the designer's goal. In Sackman's [1967] words, "the need (is) to adapt the machine to the man rather than the man to the machine". A biochemist studying protein conformations wants to speak to the computer in terms of atoms, bonds, dihedral angles, residues, and bonding forces, not in terms of linked lists, iterations, and subpictures. The former are words in his vocabulary. The latter are words in a foreign language which he has neither the time nor the inclination to master.

They discuss some psychological principles of system design.

> . . . the system should avoid psychological blocks that often prevent full user involvement in a interaction. The most typical of these blocks are <u>boredom</u>, <u>panic</u>, <u>frustration</u>, <u>confusion</u>, and <u>discomfort</u>.

They define a <u>sentence-structured</u> language as desirable.

> An action language is sentence structured if, within a given phase or subdomain of discourse, each complete user thought can be expressed in a continuous sequence of input device manipulations with standard patterns of beginning and termination. Upon termination, the machine returns to a state from which similar action sequences, other sentences, can begin.
>
> The essential features of this sentence structure are indivisible, complete thought; unbroken action; a well-defined "home state"; regularity of

> pattern. Obviously, these properties are modeled after spoken discourse.
>
> . . . .
>
> The notion of sentence structure does not preclude the possibility of paragraphing, whereby a particular local domain of thought can be delimited, as evidenced by multiple "phrases" in many systems, such as construction, solution, labeling, etc.

Recovery from incorrect or unwanted action (backup) and prompting are recommended. The issues of continuity are then presented in the context of a sentence-structured interaction language.

> The idea of visual continuity is that, within a given sentence, the eye should focus on a single area which moves in a continuous manner throughout the expression of the sentence. For example, a lightpen or cursor often provides the visual cue to action, or a sequence of flashing objects may direct attention from one area of the screen to another. Contrarily, a <u>need</u> to scan a displaced prompt field during an action is undesirable and unnatural, as is the provision of error messages on a separate device.
>
> Tactile continuity refers to the need to avoid groping or searching with hands or feet once a sentence has begun. None but a touch-typist should be required to use a typewriter keyboard as an input device within the expression of a sentence. When programmed function keys are used by the left hand, it should be possible to keep the hand in a standard position for striking keys "by feel" throughout a sentence. It should not be necessary to move the hand to an independent device during the sequence, most especially not one like the lightpen which might not be found in a fixed position. Input languages which adopt a one-device philosophy, using a lightpen or tablet as the sole hand-activated device for action sequences, emphasize tactile continuity and have great attractiveness when they can be applied.
>
> . . . .
>
> Contextual continuity refers to the absence of unrecognized side effects resulting from a user's actions. It is best achieved by providing only immediately perceivable responses to reinforce the effect of every step in the action sentence. If the reinforcement, or feedback, is visual, it

> should be within the primary field of vision (focus) for that part of the sentence. Nevertheless, at the expense of decreased visual continuity, contextual continuity can be preserved if the user is trained to look at some standard place for special reinforcements, such as diagnostics.

Finally, the need for flexibility is expounded.

> In man-man communication we usually have several ways of expressing a thought or asking a question. Yet how many graphics systems allow the user several paths to a given end? Precious few. Drafting systems offer an exception, because they typically offer several ways to construct a line such as two endpoints, or point, angle, and length. The underlying question for system designers is whether one's attitude toward the user will be "This is how you _must_ express yourself to the system," or "here are several ways to express yourself, choose whichever you prefer."

2.4 RELATED SYSTEMS

These are not electronic Richards box systems, but their builders had to face some of the same issues as I did. Table 2.1 summarizes their display techniques. Publications (if any) are indicated by year.

Publications often do not mention every single feature in a given system. The reader should understand that this chart is in some cases incomplete.

Windowing indicates the ability to scale up the picture without lines wrapping around the screen. Z-clipping here indicates the ability to change the shape of the viewing region, which contains the picture. Sequential brightening is a picking technique whereby the computer rapidly brightens each line in the picture in turn. The user simply stops this process at the line he wants to pick, backing up if necessary.

Obviously there exist many more "related" systems. Since this is not a history of three-dimensional computer graphics, I include only articles in which I feel the author has made particularly interesting observations pertinent to my work.

2.4.1 Johnson [1963]

Johnson [1963] built SKETCHPAD III, a three-dimensional extension of SKETCHPAD [Sutherland, 1963], a system for mechanical drafting.

Johnson added the third dimension to Sutherland's system, as well as dynamic scaling (picture size) and dynamic rotation, by means of a knob, in the user's frame of reference.

Three orthogonal non-perspective views and one perspective view are presented together. The user draws with a light pen in the plane of the screen. The orientation of the plane of the screen with respect to the object being modeled defines the three-dimensional direction of the line.

Our 3-D positioning motions are also in the user space and nested within the viewing space.

TABLE 2.1

Features in Related Systems

Key

- D: Dynamic
- I: Inter-active
- P: Present
- *: Absent

Viewpoint-Rotation

- M: Manual
- R: Auto. Rocking
- S: Auto. Spinning

Structural Cues

- F: Relative Fading
- D: Intens. Depth-cue
- Z: Z-clipping
- W: Windowing
- P: Perspective
- C: Color

Two-Image Display

- O: Orthogonal views
- R: Rotating shutter
- S: Side-by-side ster.
- P: Polarized mirror
- X: Other

Picking (selecting)

- L: Light pen--tablet
- S: Sequential bright.

			Rotation			Structural						Two-Image					Pick	
	Systems		M	R	S	F	D	Z	W	P	C	O	R	S	P	X	L	S
1	Johnson	1963	D	*	*	*	*	*	*	P	*	P	*	*	*	*	P	*
2	Levinthal	1966	D	*	*	*	*	*	*	*	*	*	*	*	*	*	P	*
3	Joyce	1967	*	*	*	*	*	*	*	*	*	*	*	*	*	*	P	*
4	Strauss	1968	I	*	*	*	*	*	I	I	*	P	*	P	*	P	P	*
5	Sutherland	1968	D	*	*	*	*	*	D	P	*	*	*	P	*	*	*	*
6	Ophir	1969	I	*	*	*	*	*	*	P	*	*	*	*	*	P	*	*
7	Barry	1969	D	*	D	*	D	*	*	P	*	*	*	P	*	*	*	*
8	Hobgood	1969	D	*	*	*	*	*	*	*	*	*	*	*	*	P	*	*
9	Ortony	1971a	D	*	*	*	*	*	*	*	*	*	*	P	P	*	P	*
10	Ortony	1971b	D	*	D	*	D	*	*	I	*	P	*	P	*	*	P	P
11	Barry	1971	D	*	*	*	*	*	*	*	*	*	*	P	*	*	*	*
12	Tountas	1971	D	D	D	*	*	*	*	*	*	*	*	*	*	*	*	*
13	Perkins	1971	D	*	*	*	*	*	*	D	*	*	*	P	*	*	*	*
14	Wright	1972	I	*	*	*	*	*	I	*	*	*	*	*	*	*	P	*
15	Levinthal	1974	D	D	D	I	*	I	I	*	*	*	*	*	*	*	*	*
16	Plott	1975	*	*	D	*	*	*	*	*	*	*	*	P	*	*	*	*
17	Kilpatrick	1976	D	*	*	I	*	*	*	D	*	*	P	*	P	*	P	*
18	Feldmann		D	*	*	*	*	*	D	I	I	*	*	*	P	*	*	*

2.4.2 Levinthal [1966]

Levinthal [1966] and coworkers built a system for interactively building protein structures. He used a two-degree of freedom, spring-return velocity device for manually turning the image on the screen.

> If this is done, the brain of the human viewer readily constructs a three-dimensional image from the sequential display of two-dimensional images. Such sequential projections seem to be just as useful to the brain as simultaneous stereoscopic projections viewed by two eyes. The effect of rotation obtained from the continuously changing projection nonetheless has an inherent ambiguity. An observer cannot determine the direction of rotation from observation of the changing picture alone. In the Project MAC display system, designed by John Ward and Robert Stotz, the rate of rotation of the picture is controlled by the position of a globe on which the observer rests his hand; with a little practice the coupling between the hand and the brain becomes so familiar that any ambiguity in the picture can easily be resolved.

2.4.3 Joyce [1967]

Joyce [1967] argues that a light pen should indicate to the user what it sees _before_ the user commits himself to select that object. Joyce chose to intensify the line seen by the light pen. However, if an object is not selectable, it does not intensify or otherwise react when the light pen sees it. This gently steers the user away from meaningless actions. He calls this technique a _reactive display_.

We have a similar technique -- a small circle shows what the light pen will pick if tapped.

2.4.4 Strauss and Poley [1968]

Strauss and Poley [1968] built a three-dimensional graphic system for piping layout. This system indicates locations in space with a "tracking asterisk" or "jack" which has three orthogonal intersecting arms each parallel to an axis in the viewing space. It moves in one of these six directions selected by push buttons and at a dynamically controlled rate. An arrowhead appears on the arm pointing in the direction of motion.

It is interesting that Strauss and Poley have no difficulty with viewing-space translation while we find user-space translation necessary. This probably arises from the difference in applications. Their pipes run mostly parallel with the viewing axes and motion along them is meaningful. In molecular work the application axes have little bearing on molecular structure, though they are related to crystal structure.

Top, front, and side orthographic projection and sudden 90-degree turn are provided. Two stereo techniques are also available:

> Two side-by-side images, each of the piping system as it would be seen by one of the viewer's eyes looking through a window are displayed on the scope face and can be viewed through a stereoscope. The stereoscope fuses the images and produces the illusion of depth for the viewer. The use of true stereo, as contrasted with perspective or isometric representation, allows a realistically complex piping system to be displayed with adequate clarity and detail, in spite of the small size of the viewing window.
>
> Another method of stereoscopic display suitable for group viewing, is for two closed-circuit television cameras, one focused on the left eye and one on the right eye image, to be connected directly to a color television receiver. If the image for one eye is displayed in red and the image for the other eye is displayed in blue, the fused 3-D image can be observed by viewing the television screen through a pair of glasses with one blue and one red lens.

Plotted output of the screen is available. The display unit is an IBM 2250.

2.4.5 <u>Sutherland</u> [<u>1968</u>]

Sutherland [1968] has constructed a head-mounted display that delivers a stereoscopic image to the subject by means of two small CRT's alongside the head and a pair of prisms in front of the eyes. A passive mechanical arm connects the headgear to the ceiling and senses the position of the subject. This allows the subject to "walk around" the stick figure object he is viewing. It appears to float in the middle of the room. I have worn this device and can pronounce the sensation realistic.

Vickers [1974] installed a companion position-indicating device (handgrip or wand) suspended from the ceiling by three spooled lines. It is free to move about a substantial portion of the room. In one experiment subjects were presented a cube made of twelve edges floating in the center of the room. With the wand they could grab a corner of the cube and stretch the three lines that form it, distorting the figure. Vickers reports that with the same image presented to each eye (no stereo) users walking around the cube reported feeling they saw stereo even though they actually were not. However, when they began to distort the cube with the wand, they had great difficulty perceiving what was happening.

The display tubes physically move with viewing motions keeping the user space and the viewing space in register. This obviates the need to consider compensating one into the other.

2.4.6 Ophir et al. [1969]

Ophir et al. [1969] built an inexpensive stereo display employing a standard raster scan television display. The picture is viewed through red and green glasses (the anaglyph technique). They displayed stick figure objects with perspective.

> An interesting improvement in three dimensional presentation was realized when the bright and dark parts of the display were reversed (accomplished in the program) so that the figures were presented as dark lines on a bright field. This substantially enhances the subjective reconstruction of the object in space -- possibly because the lines to be superimposed in the brain are seen in the same "color" (black) by both eyes rather than as red to one eye and green to the other. . . . The bright field technique provides a more restful and less fatiguing display.

These comments steered me away from anaglyph stereo for a vector display.

2.4.7 Barry et al. [1969]

Barry et al. [1969] give an extensive treatment of dynamic rotation, intensity depth-cueing, and side-by-side stereo. The consequences of the approximation

a = sin a

1 = cos a

for small "a" are discussed, and some general conclusions are drawn.

> The rotation method has an ambiguity with regard to the direction of rotation which is difficult to resolve, even with some sort of perceptual feedback, when the object is unfamiliar to the viewer. Of course, all three-dimensional perception halts when the rotation stops. Leisurely study and contemplation of a rotating figure is difficult, at best. In addition, continuous rotation puts large computational demands on the computing hardware which may seriously degrade the quality of the display.
> . . . .
> . . . removal of hidden lines is probably too time consuming for a dynamic display without special hardware.

Their hardware is a Link 300 computer with a dual trace oscilloscope display.

2.4.8 Hobgood [1969]

Hobgood [1969] built and tested a varifocal mirror viewer for an IBM 2250 display unit following the work of Traub [1968]. It displays stick figure objects.

> [The] images . . . are reflected in a flexible mirror whose focal point changes continuously as the mirror shape is deformed from concave to convex and back. . . . The displayed material appears as a virtual image in a well-defined space behind the mirror at a depth determined by the curvature of the mirror surface.
>
> The mirror is a sheet of silvered Mylar. When stretched over the front of a large loudspeaker which is driven at a low frequency sinusoid, the

> mylar bulges inward and outward, and the locations of the virtual images move three to six inches in the space behind the mirror.

This technique, as it now stands, holds promise for simple 3-D pictures and so-called 2 1/2-D applications such as tomography where objects lie mainly in the plane of the screen. Henceforth, the term plane of the screen will not distinguish between the physical plane of the screen and those planes parallel to it.

Because picture complexity is limited by computation speed, the varifocal mirror may not be suitable for electron density fitting without special-purpose hardware to display many vectors and to perform 3-D dynamic rotation. This hardware, unlike the special-purpose hardware in current vector displays, is not, at present, commercially available.

2.4.9 Ortony [1971a]

Ortony [1971a] designed a stereo viewer (British patent appl. 13844/70) for a molecular computer graphic system. He uses a horizontal two-way (half-silvered) mirror to combine images drawn on the upper and lower halves of a CRT. The lower image is seen directly through the glass, the upper image is seen reflected. Crossed polarizers placed flat against these images allow exactly one of them to be seen by each eye of a user wearing polarizing glasses.

> The most important advantage which the transmission/reflection method has over comparable rivals such as the mirror stereoscope or fusion by prisms is based on the way in which the images are discriminated. Since the images are already overlaid, image routing is a question of isolation by filters rather than by geometry. This means that the isolation is position-independent, and this in turn allows freedom of movement without loss of perception.

He goes on to mention that this allows simultaneous viewing by several people at little added expense.

Based on these favorable comments and the general popularity of this display, we built a viewing device for it.

2.4.10 Ortony [1971b]

Ortony [1971b] describes his experiences with his molecular graphic system.

Items may be selected directly with a light pen or by sequentially brightening each bond in the molecule, stopping, and backing up with a twist of a potentiometer until the desired bond is illuminated steadily.

> . . . if a truly stereoscopic image is perceived it is difficult, to say the least, to point the light-pen at a line seen to be in front of the plane of the screen, and impossible to point it at a line seen to be behind this plane.
> . . . .
> There is a problem with identification by sequential brightening if one wishes to use interactive stereographics, together with intensity modulation, in that the brighter lines at the front are difficult to perceive with the increased intensity which takes place during the sequential line brightening. The use of some other identification mark, such as a cross, would solve the problem.

Points in space are specified with a three-dimensional "bug" under dynamic control.

Continuous spinning is available. Ortony comments on what he calls a structural after-image (the memory effect of the KDE [Wallach, O'Connell, and Neisser, 1953], page 13 above):

> A common objection to dynamic rotation as a clue to the three-dimensional structure of an object is that with the cessation of rotation, there is associated a loss of insight into the structure. However, our experience shows that in very many cases there is what one might call a "structural after-image" associated with dynamic rotation. The associated objection concerning the difficulty of interaction has more force. There seem to be two objections to dynamic rotation which are more severe. With depth-ambiguous objects there is the problem of inversion. During rotation the effect of this, as mentioned earlier, is to reverse the apparent sense of rotation. In many cases, this is confusing, and in all cases it is disturbing. The second objection is that it is often not possible to rotate a complex object sufficiently

> quickly to produce the desired effect, without special-purpose hardware.

He states that this "structural after-image" effect carries over into stereo.

> It is worth noting that a user may gain so much insight into the nature of his material from one brief glance at it displayed stereoscopically, that his interpretation of a monocular view is subsequently improved to the extent that he is able to do a great deal without further use of stereo.

Two double-image display modes are available in his system, side-by-side orthogonal views as in a mechanical drawing, and polarized-mirror stereo as described in Ortony [1971a] (page 25 above). He divides objects to be viewed into several categories. There are the familiar objects with parallel lines, like buildings, which he calls <u>visually-predictable</u>, and objects with unusual lengths and angles, molecules for example, which he terms <u>visually-unpredictable</u>. For each of these it is not always possible to inherently deduce which of the image's parts are nearer the observer than others, as in the Necker Cube. In this case the scene is called <u>depth-ambiguous</u>.

> Stereo seems to be of no assistance with visually-predictable objects if they are not depth ambiguous, thus no additional insight is gained by the use of stereo in viewing, for example, a building. With visually-unpredictable objects, however, stereo comes into its own; for with such objects, the traditional cues used in computer graphics fail to adequately clarify the object. Furthermore, there is little doubt that stereo greatly increases the ease and comfort with which depth-ambiguous visually-predictable objects are viewed.

Intensity depth-cueing he calls <u>intensity modulation</u>.

> Intensity modulation was found to be of greatest help in cases where groups of associated lines lying in the same or nearly the same plane could be identified as groups in particular views where they overlapped by virtue of the differing intensities. In certain views of molecules, for example, which are more or less unintelligible without intensity modulation, their intelligibility was greatly increased because it was possible to

> recognise the bonds belonging to particular rings in this way.
> . . . .
> For visually unpredictable objects, stereo alone or together with intensity modulation with depth, provides an impressive and clear representation of an object. Our experience so far suggests that interactive stereographics will be most important in the modeling of complex molecular structures where stereo is already being widely used for the purpose of visual representation.

He closes on a note of caution.

> One should spend available computer time and space to provide the most appropriate useful facility, rather than the most true to life facility, which frequently does not give the greatest insight.
> . . . .
> In conclusion then, the use of stereo displays and interactive stereographics should be restricted to cases where it is really needed.

2.4.11 Barry et al. [1971]

Barry et al. [1971] describe a system for displaying side-by-side stereo images of molecular structures on a LINK 300 computer's oscilliscope.

These can be dynamically rotated by the user with a two-axis mechanical linkage that the GRIP-75 system's manual viewpoint control resembles.

The mathematics of stereo is discussed.

2.4.12 Tountas and Katz [1971]

Tountas and Katz [1971] give a good discussion of schemes for nested rotation and for user-space dynamic rotation.

These algorithms are implemented on an Adage AGT/50 display unit.

2.4.13 Perkins et al. [1971]

Perkins et al. [1971] have a system for viewing molecular structures. Their system features side-by-side stereo and dynamic user control of viewpoint orientation (three Euler angles), picture size, perspective, stereo disparity, and image separation by means of seven linear potentiometers.

It runs on a Honeywell DDP 516 digital computer with an EAL TR48 analog computer and a Lanelec "19" oscilloscope.

They conclude:

> The [stereo map reader] mirror system . . . was found to be adequate for viewing the display tube by a single observer. For additional observers, wedge lenses were tried but were found to be less effective than small hand held versions of the mirror units. . . . During general demonstration, adaptation to stereo viewing was found to be quite varied, some persons finding it quite impossible whilst others disdained the use of external devices since they experienced little difficulty in discomforting their vision to accommodate the stereo view.
> . . . .
> The addition of the third dimension to displays has been shown to be of immediate value in the examination of molecular structures. When locating specific molecular positions the 3-D effect obtained by rotation of a single image is sometimes adequate, but the stereo display is certainly needed for observing stationary displays, . . .

2.4.14 Wright [1972]

Wright [1972a], [1972b], built a graphic system called GRIP-71 for discrete manipulation and interactive energy relaxation of proteins. It used a light pen, a box of buttons, and a keyboard. Interactive viewpoint-rotation was possible. Interactive 90-degree turn, as well as perspective, were tried but were not found helpful.

It used an IBM 2250 display unit and a time-shared IBM System/360 Model 40 computer. This system is the direct ancestor of the GRIP-75 system now running on different hardware.

Some of Wright's results have been summarized by Brooks [1977].

2.4.15 Levinthal et al. [1974]

Levinthal et al. [1974] have a system called CARTOS for tracing nerve fibers through tissue. His pictures are similar to those used by crystallographers fitting molecules to electron density maps. It runs on an Adage AGT/50 display unit connected to an IBM System/360 Model 91 computer with a 100 kilobaud communication link.

Dynamic position-indicating in x and y is done with a small hand held "mouse". Continuous rocking and spinning are available. Interactive windowing and interactive relative fading of two superimposed images are provided.

2.4.16 Plott et al. [1975]

Plott et al. [1975] built a system for perceptual studies using a small digital computer, a digital-to-analog converter, and a Tektronix 555 oscilloscope. The pictures they generate are mathematical functions which very much resemble a molecule inside an electron density map.

They performed an experiment with ten subjects who saw a side-by-side stereo display through both a stereoscope and some sort of viewing device with polarized lenses. In each case it took about five minutes for noticeable eyestrain to develop. When the images were increased in size the strain was reduced.

They also tried continuous spinning, and conclude:

> The rotational motion feature, incorporated into the computer program, provided an excellent illusion of the image slowly revolving around a fixed point. It is felt this aspect should certainly be embodied in systems of this nature.

2.4.17 Kilpatrick [1976]

Kilpatrick [1976] has studied

> . . . how force cues can be used to augment visual cues in an interactive computer graphics system.

He did this by building a system for manipulating visually predictable objects (blocks) in a perspective display using an Atomic Energy Commission Argonne Remote Manipulator (ARM). It has force feedback servos in all its seven degrees of freedom (rotation, translation, and pinch).

His work was carried out in parallel with the GRIP-75 project using much of the same hardware in the same graphics laboratory. Although working on a different problem, many of the conclusions he draws seem relevent to molecular computer graphics.

Just as in the GRIP-75 project, which does not employ perspective, he found dynamic control of disparity desirable.

> The user can change various things other than the viewpoint. For any given viewpoint and user, there is one theoretically correct amount of stereo separation. In practice, however, users prefer to adjust the amount of separation. The maximum amount of separation a user can merge is a function of several psychological factors plus the mechanical behavior of the Lorgnette [spinning disk stereo]. A larger separation than the theoretically correct amount enhances the stereo effect and makes relative distance discriminations more precise. As the Lorgnette slips out of adjustment, the user is not able to work with as much separation. Furthermore, a large separation tends to accelerate eye fatigue. Moreover, one can merge a wider separation reached by a continuous increase from zero than a separation discontinuously switched on. It is desirable to be able to adjust the magnitude of this cue to the task at hand and to personal preference. The stereo separation is adjusted with a slide potentiometer.

He comments on the desirability of exaggerating the reach of a translation control with respect to the motions it produces on the screen.

> A seemingly natural choice would be to have the device the same size as a typical CRT screen, approximately 10 inches on a side. This choice would give absolute correspondence between the physical and virtual movements. However, as noted earlier, the average user can normally detect

movements of 0.04 inch. Not only is this much coarser than the visual modality, it is substantially coarser than commercially available CRTs which routinely provide 0.02 to 0.01 inch resolution. If the user cannot input information as accurately as the screen can display it, some of the output resolution is effectively lost.

This suggests that the kinesthetic display be two to four times the size of the CRT screen. Experience with our ARM has shown that users adapt very quickly to an "enlarged" kinesthetic space as long as the relative movements are consistent.

He concludes,

> Several subjects explicitly divided the cues into these two groups.
>
> <u>Decomposition</u>: While watching subjects, I became convinced that most subjects immediately started looking for a way to decompose the three-constraint problem of placing the center of the tongs close to the center of a block into a 2-D fit and a 1-D fit. In fact, the one subject who initially tried to pick up the blocks without doing this had a very difficult time. Finally, he too developed a strategy of decomposition.
>
> . . . .
>
> The movable viewpoint allows the user to present a series of nearly orthogonal views. This is a method of decomposition, but it is awkward in that it requires two hands, it requires a shift of visual attention, and it often takes several iterations. Each iteration requires that the user move the joystick. This operation typically takes almost a second.
>
> Movable viewpoint would probably be better for this task if only two viewpoints were available. If the viewpoint were changed with a button on the handgrip, the need to use two hands and the shifts in visual attention would be eliminated. Even in this configuration, however, explicit iteration would be necessary. With both the markers [on the screen] and force [feedback], explicit iteration is not necessary because those cues allow the user to switch between aligning along the various axes at mental speed.

The stereo cue [with perspective] provides no method of decompostion along the familiar [viewing-space] axes. The best technique is to align the tongs with the block on the screen and then move back and forth along the line of sight until the tongs surround the block. This is a difficult procedure. Furthermore, it is not evident whether the success of this operation depends on stereo cues or perspective cues.

The fact that stereo does not provide a method of decomposition along the familiar axes, and the additional fact that a non-trivial portion of the population cannot <u>see</u> stereo pictures, indicate that stereo should not be used as the only cue for three-dimensional selection.

. . . .

It is not particularly surprising that the users adopt a strategy of decomposition. This is quite consistent with our experience gained while operating our pair of ARMs in normal master-slave mode. In our laboratory, the user views the slave ARM, a table, and several blocks through a closed circut television system. A standard procedure for picking up blocks is to slide along the table until the tongs collide with the blocks.

We also see evidence of this strategy of decomposition in real manipulator operations. It is common to have the manipulators equipped with several cameras that can be used to give different views of the work area . . . .

. . . .

<u>Individual Differences</u>: A second conclusion is that different people prefer to use the various [3-D perception] cues in different ways. These differences in use can lead to a different order of preference.

. . . .

<u>Guidelines</u>: I propose the following guidelines that are based on these two conclusions.

1. When faced with a task that requires moving a point probe to a target in a three-dimensional space, the user of a graphics system will be most aided by cues that provide a natural mechanism for decomposing the task into a 2-D fit and a 1-D fit. If a three-degree-of-freedom orientation must be specified, the user will also profit from cues that

allow the orientation task to be decomposed into a 2-D fit and a 1-D fit.

2. How useful a particular cue is to a particular user is a complex function of how the user attacks the problem at hand. Any real system should either be custom tailored to the user(s), or have enough redundant cues that a variety of users can adopt satisfying strategies.

Some of Kilpatrick's results have been summarized by Brooks [1977].

2.4.18 Feldmann

A molecule viewing and manipulation system is currently under development at the National Institutes of Health in Bethesda, Maryland, by Feldmann and Porter. Its main purpose is to support raster graphics for the AMSOM microfiche atlas [Feldmann, 1976], but it also can manipulate molecules by on-line idealization with atoms targeted to points in space. That is, the user marks points in space (targets) where selected atoms ought to be. He then invokes a command to advance the selected atoms towards their target positions iteratively. The atoms drag some of the rest of the molecular structure with them while the program attempts to maintain any rigid connections in approximately the correct shape.

An Evans and Sutherland PICTURE SYSTEM 2 display unit driven by a DEC PDP-11/70 computer is used to orient dynamically a stick-figure representation of a large molecule. A color surface representation of the structure can be displayed in about a minute on a color monitor or recorded directly on film.

Three knobs orient the picture through Euler angles. Three more knobs translate the structure in user-space. The stereo disparity is fixed. Atom selection is done by keyboard.

2.5 ELECTRONIC RICHARDS BOX SYSTEMS

The display techniques of these systems are summarized in Table 2.2 by approximate order of publication or use. Publications (if any) are indicated by year.

TABLE 2.2

Electronic Richards Box Systems

Key

- D: Dynamic
- I: Interactive
- P: Present
- *: Absent

Viewpoint-Rotation

- M: Manual
- T: 90-degree Turn
- R: Auto. Rocking
- S: Auto. Spinning

Two-Image Display

- O: Orthogonal views
- R: Rotating shutter
- S: Side-by-side ster.
- P: Polarized mirror

Structural Cues

- F: Relative Fading
- D: Intens. Depth-cue
- Z: Z-clipping
- W: Windowing
- P: Perspective
- C: Color

Picking (selecting)

- L: Light pen--tablet
- S: Sequential bright.
- K: Numeric Keyboard

		Rotation				Structural						2-Image				Pick		
Systems		M	T	R	S	F	D	Z	W	P	C	O	R	S	P	L	S	K
1 Oxford	1972	I	*	*	*	D	*	*	*	*	*	P	*	P	P	*	*	P
2 Leeds		I	I	*	*	P	P	*	D	*	*	P	*	P	P	*	*	P
3 MMS-4	1974	D	*	D	D	P	P	*	D	*	*	*	*	P	P	*	P	*
4 NIH-Adage	1974	D	*	*	D	*	I	D	I	I	*	*	*	*	P	*	*	P
5 Princeton	1975	D	*	P	D	*	P	*	*	P	I	*	P	*	*	*	*	P
6 CRYSNET	1975	D	*	D	D	P	D	D	I	*	*	P	*	P	P	*	P	*
7 GRIP-75	1975	D	D	D	D	D	D	D	I	*	*	*	P	P	P	P	*	*
8 Yale PIGS		D	*	*	*	D	D	D	D	P	*	P	*	P	*	P	*	*
9 MMS-X	1976	D	I	*	D	I	*	*	D	*	*	*	*	P	P	*	*	P
10 San Diego		D	*	*	D	I	*	*	D	P	*	*	P	*	*	*	*	P
11 Bilder	1978	D	I	*	*	D	D	D	D	I	*	*	*	*	*	P	*	*
12 RING	1978	D	*	*	*	D	P	D	D	D	*	*	*	*	*	P	*	*

Published information is incomplete. I apologize in advance for omissions and misreadings. References to many systems were edited to those papers that best describe 3-D display techniques. Some of the information in this section is from Barry [1976] and personal communications from J. Andose, P. Evans, S. Dempsey, A. C. T. North, and A. Olson.

2.5.1 Oxford (1972)

At Oxford University in England, Barry and North [1972] built an electronic Richards box system using a Ferranti Argus 500 computer and a Ferranti model 30 display. Input is by means of a keyboard, a knob, and a tracker ball.

The tracker ball allows dynamic viewpoint-rotation in user-space of a small molecule, and there is an option for automatic spinning. The tracker ball can also be used manually to dynamically rotate a bond to alter the shape of the molecule. The contoured electron density map is rotated interactively.

A knob sets the relative intensity of the molecule and the map. The display tube has two beams whose intensities can be separately controlled. The display will not clip lines that fall outside the screen boundaries.

They report good experience with simultaneous orthogonal views. Side-by-side and polarized-mirror stereo are also available.

2.5.2 Leeds

At Leeds University in England North, Ford, and Watson have built a system using a DEC PDP-11/40 computer driving a Hewlett-Packard display. Input is through 16 knobs and a keyboard.

Six knobs rotate and translate a small molecule in the user's frame of reference. The orientation of the electron density map may only be adjusted interactively. A 2-D spring-return joystick translates the entire image and a knob adjusts its size.

Two knobs set the rate of automatic bond rotation which dynamically builds a chart of allowable angles for two bonds. This chart can then be used as a guide for subsequent manipulations. As structure is modified, atomic contacts are indicated dynamically by dotted lines and by sound.

Successive orthogonal views or any pair of simultaneous orthogonal views may be selected with a rotary switch. Side-by-side stereo and polarized-mirror stereo are also available. A knob controls the separation between the two images. The stereo disparity is changed interactively.

2.5.3 MMS-4 (1974)

At Washington University at St. Louis an electronic Richards box system has been built [Barry et al., 1974; Fritch et al., 1975]. It is called MMS-4 (Man-Machine System) and it runs on a Link 300 computer and an LDS-1 display. Hand-made digital modules implement many of the dynamic operations. Interactive operations are performed on a keyboard attached to the Link 300 video display.

The dynamic operations include viewpoint-rotation controlled by a mechanical linkage like the GRIP-75 manual viewpoint control, automatic rocking and spinning, and multiple bond-twisting with rotary potentiometers. Intensification of the molecule relative to the electron density map is provided by redrawing the molecule several times each refresh. Dynamic intensity depth-cueing and windowing are available. Side-by-side stereo and polarized-mirror stereo have dynamic control of the picture size and the horizontal distance between (percent overlap of) the two images. Bond selection is controlled by a knob and indicated by brightening.

2.5.4 NIH-Adage(1974)

Cohen and Feldmann [1974] at NIH in Bethesda, Maryland, built a system for fitting a molecule to its electron density map with an Adage display unit and a PDP-10 computer.

The user is presented a keyboard and a box of push buttons for atom selection and other interactive commands. Atom selection is by typing an atom number.

Knobs are provided for dynamic operations. Viewing and substructure rotation is by six nested Euler angles. These may be automatically incremented to produce continuous tumbling. Substructure translation is in viewing-space.

Alternatively, a single bond may be twisted. The system interactively generates construction lines that show close atomic contacts.

Intensity depth-cueing is stepped up and down through eight-levels by push buttons. "Z-clipping" in any one of six directions is performed smoothly by turning a knob.

Perspective is available but not used much. Polarized-mirror stereo allows dynamic control of the stereo disparity and of the horizontal separation of the two images.

2.5.5 Princeton (1975)

At Princeton [Stellman, 1975] a PDP-10 computer and an LDS-1 display unit provide a fitting system for small molecules.

With knobs the user can dynamically control the view-point-rotation (three Euler angles), bond-twisting, and scale. The picture can be made to rock or to spin through Euler angles (tumble). Intensity depth-cueing is fixed. Windowing is provided. The relative intensity of the molecule and the map is not an issue since the map is not shown. It is represented by an "R factor", a figure of merit for the fit of the molecule to the map. This is periodically displayed as the molecule is moved about. The chemist attempts to decrease its value.

Stereo is available by use of a spinning disk (Lorgnette), which with colored sectors can also display color. Stellman comments:

> The stereo and color options, though available, did not seem to enhance the image and were not often used.

2.5.6 CRYSNET (1975)

At Texas A&M University [Collins et al., 1975], at the Brookhaven National Laboratories, and the Institute for Cancer Research in Philadelphia, there are systems built with a DEC PDP-11/40 computer and a Vector General Series 3, Model 3DR, display. A box of knobs and a keyboard on a video terminal allow interaction.

They have dynamic viewpoint-rotation, rocking, spinning, residue fitting, bond-twisting, intensity depth-cueing, and z-clipping controlled by rotary potentiometers.

Side-by-side orthogonal views, side-by-side stereo, and polarized-mirror stereo are provided with dynamic control of image size. The stereo disparity is fixed. Hardcopy plots are available on-line from a Versatec printer-plotter.

2.5.7 GRIP-75 (1975)

At the University of North Carolina at Chapel Hill, an electronic Richards box system called GRIP-75 was built. It is described in Foley and Wright [1975], Tsernoglou, Petsko, McQueen, and Hermans [1977], Brooks [1977], Britton [1977], Britton, Lipscomb, and Pique [1978], and Pique [draft].

A Vector General Series 3, Model 3DR, display is driven by a DEC PDP-11/45 computer, which is connected by a 250 kilobaud communications link to a 400 kilobyte region in an IBM System/360 Model 75.

The dynamic features are viewpoint-rotation with a mechanical linkage, smooth 90-degree turn, automatic rocking and spinning, 3-D translation of either a jack-shaped position indicator or a molecule substructure using a 3-D box provided by M. Noll at Bell Labs, unconstrained rotation of the substructure with 3-D spring return joystick used as a velocity control, eight-level nested bond-twisting, relative fading of the molecule vs. the density map or of one molecule superimposed on another, intensity depth-cueing, z-clipping, stereo disparity, and picture size during stereo display.

The interactive viewing features are push-button selection of rocking or spinning, windowing, selection of stereo mode (rotating-shutter stereo, side-by-side stereo, and polarized-mirror stereo), and the light pen. A hardcopy pen plot of the screen is also available interactively.

2.5.8 Yale PIGS

At Yale University a system called PIGS (Protein Interactive Graphics System) uses a stand alone DEC PDP-11/70 with an Evans and Sutherland PICTURE SYSTEM 2 display unit. A data tablet is used for picking and for most dynamic actions. Some knobs are provided as an alternative.

Dynamic viewpoint-rotation is provided by a velocity control device. A relative fading slider is provided. Intensity depth-cueing is linked to the z-clipping. Dynamic windowing and perspective are allowed.

Three simultaneous orthogonal views can be shown (top, front, and side) simultaneously with a perspective view. Side-by-side stereo can be viewed through a viewing device that swings in from the left of the workstation.

2.5.9 MMS-X (1976)

The successor to the MMS-4 system at Washington University is under development [Rosenberger et al., 1976]. The "X" in MMS-X stands for "exportable" and "expandable". Indeed, a few complete systems have been shipped.

A TI 980B computer drives a Hewlett-Packard 1321A display through locally developed coordinate-transform hardware. PERTEC D3341 disk drives store the molecule and map data, and a Beehive video terminal handles all interactive commands from the user.

A pair of 3-D, spring-return, rate joysticks translate or rotate, by switch selection, the overall view and a substructure. The current viewpoint-rotation rate may be held fixed by flipping a switch, which maintains the continuous spinning. Six knobs control bond-twisting.

A knob controls the scale. Four push buttons perform instantaneous 90-degree rotation left, right, up, and down.

The molecule can be drawn brighter than the map with a switch-selected option that slows the beam while drawing the molecule.

Side-by-side stereo and polarized-mirror stereo are available. A stereo map reader as well as prism glasses with large polarized sheets are used for side-by-side stereo viewing. The polarized-mirror technique uses two display tubes and a combining mirror. Rotary knobs control the stereo disparity and percent of image overlap.

2.5.10 San Diego

At the University of California at San Diego, Cornelius and Kraut have a DEC PDP-11/40 emulator (CalData 135) driving an Evans and Sutherland PICTURE SYSTEM display. Input is by a keyboard and a box of knobs.

Six knobs rotate and translate the picture in user-space. A switch allows for continuous rotation and translation, whose rate is set by the knobs.

The user can invoke commands which increment and decrement the molecule's brightness. He also can change the amount of perspective distortion. Lines that fall outside the screen are clipped.

The stereo viewer, a Lorgnette, is apparently quite popular with the users.

2.5.11 Bilder (1978)

R. Diamond [1978] at the MRC Laboratory for Molecular Biology in Cambridge, England, has built a system called Bilder with a DEC PDP-11/50 with an RP04 disk driving an Evans and Sutherland PICTURE SYSTEM display. G. Cohen has imported a version of this system to run on the PDP-11/70 and PICTURE SYSTEM 2 at the NIH in Bethesda, Maryland, and another copy has been installed at the European Molecular Biology Laboratory in Heidelberg, Germany. All input, interactive and dynamic, is by means of a data tablet.

Molecules are initially built on-line one residue at a time with dynamic bond-twisting preserving ideal geometry. Later passes through the structure employ bond-twisting of side chains and on-line idealization.

Dynamic viewpoint-rotation is by means of a 2-D velocity control on the tablet. The user can also command instantaneous 90-degree turns. A velocity slider sets the picture scale.

The molecule is drawn brighter than the map. A single slider adjusts a combination of relative intensity and depth-cueing. There is an on-off perspective display switch.

2.5.12 RING (1978)

T. A. Jones [1978a], [1978b], at the Max-Planck-Institut fur Biochemie in Munich, Germany, has a Vector General 3404 display driven by a DEC PDP-11/40 computer, which is, in turn, connected to a Siemens 4004 computer. On this hardware he built a system he calls the set of RING programs.

Three knobs rotate the picture about the x, y, and z axes.

An arbitrary molecular substructure may be translated and rotated by six knobs with respect to the rest of the structure, or, alternatively, up to three bonds may be rotated by three knobs. On-line idealization and energy minimization are used often to help correct the molecular geometry.

A knob adjusts the intensity of the molecule. The intensity of the combined molecule and map can be raised or lowered by the intensity knob provided on the display unit. Jones [1978a] reports that the manual rotation of the picture produces a 3-D effect so potent that

> By also using intensity depth cueing, we remove the need to display stereoscopic pairs. . . .

Knobs control the front and back z-clipping planes and the overall picture size (windowing). A knob sets the perspective, and Jones [1978a] further reports that

> Perspective has been little used, the combination of windowing, depth cueing and immediate rotations makes the display easy to work with, even with a cluttered screen.

An atom can be picked by a single stroke of a data tablet pen.

Chapter 3

PROJECT HISTORY

3.1 INTRODUCTION

My work on this project spans several years. Only that portion which concerns three-dimensional display and manipulation is described in detail in later chapters. Many people have worked many years on this project, and much of their work overlaps mine. An overview of the project will help clarify these relationships.

F. P. Brooks, Jr. at the University of North Carolina at Chapel Hill, has long been interested in man-computer interactions. He defined and led a number of projects that investigate this in computer graphics, where the man-machine interface is quite intimate.

One of these undertakings is a series of molecular graphic systems called the GRIP project. The computer plays the role of an "intelligence amplifier", easing the user's work but having little built-in intelligence. Of course, the gain of this amplifier must be greater than "one", or else it becomes an "intelligence attenuator".

The system has evolved directly with three successive goals producing three stages in its development. These three checkpoints in GRIP's evolution are called GRIP-71, GRIP-74, and GRIP-75.

3.2 GRIP-71

In 1971 W. V. Wright built a graphic system for energy minimization of small protein structures as part of a Ph. D. dissertation [Wright, 1972a; Wright, 1972b] (page 29 above). GRIP stood for <u>Graphical Representation of Interacting Proteins</u>.

This system ran on an IBM 2250 display unit attached to an IBM System/360 Model 40 computer. It involved about 3,000 lines of source program, nearly all in PL/I. (Blank lines and comments are included in this and later line counts.)

The system's major features and limitations were:

1. Its purpose was interactive energy relaxation of proteins. The potential-energy formula needed to be fast and thus considered such things as the 6-12 potential and polar charge, stopping short of molecular orbital calculations.

2. Other interactive manipulations, bond-twisting and general substructure movement, were made available to the user. He could thus set up arbitrary perturbations of an existing conformation.

3. The molecules available were presented in sequential order on command.

4. A postfix notation, such as current Hewlett-Packard calculators use, was employed throughout for all commands.

5. The operands for all commands were composites embodying a direction vector, a scalar, and an atom number. These three types of operands served to support all the available functions.

6. The verbs (system functions that one can invoke by interactive commands) were, in many cases, implemented in terms of more primitive verbs. These primitive verbs could be executed in sequential order or with conditional branches and loops. This implementation was invisible to the user who saw the same syntax for all commands. All the primitive verbs that were useful when directly invoked by the user were placed on separate menus. This established a hierarchy of verbs, the powerful verbs that make many assumptions about what the user normally wants, down to the primitive verbs from which they were made. The user could select any combination of these he wanted as it suited his purpose.

7. The commands sometimes took more than one parameter, and the user had a stack to push them on as neeeded. Thus, while setting up for one command, the user could push some operands for a second command, invoke it, and continue setting up the first. Actually, users seldom did this. If setting up and invoking a command can be thought of

as speaking a sentence [Foley and Wallace, 1974] (page 16 above), then the user is able to make parenthetical remarks nested many levels deep.

8. It was optimized for polymers. It was not designed to cope well with general molecular structures.

9. Amino acid residue boundaries were not visible to the software. The data structure described a homogeneous tree of atomic coordinates.

10. No dynamic operations were available, a limitation of the display unit.

11. Only small structures could be handled. The largest naturally-occurring structure tried was an eight-residue loop. The largest artificial structure had ten residues or about 100 atoms.

12. "The most obvious weakness of the . . . system for the selected application is the absence of facilities to help the user visualize the three-dimensional structure of his model." [Wright, 1972a].

3.3 GRIP-74

In 1973 new graphics hardware was acquired for a different project [Kilpatrick, 1976] (page 30 above), but it was also suitable for continued GRIP development: a Vector General Series 3, Model 3DR, display and a DEC PDP-11/45 computer with a 250 kilobaud communications link to a host IBM System/360 Model 75 computer. We have continued using this equipment for GRIP-75.

M. E. Pique and I were directed by Dr. W. V. Wright to reimplement the GRIP-71 system on the new hardware and to institute some improvements. By early 1974 I had selected enhanced 3-D techniques as an M. S. thesis topic although difficulties with the other enhancements put off a start until mid-1974.

This system, but not my thesis, was substantially complete in late 1974, hence the name GRIP-74. GRaphics Interaction with Proteins. At that time it involved about 4,000 lines of source code in PL/I on the host IBM 360 and about 3,700 lines of source, mostly in a locally developed PL/I, on the PDP-11. There was substantial overlap in time between this project and GRIP-75, but the GRIP-74 goals were these:

1. This was to be a direct extension of GRIP-71, serving the same need.

2. The user could select molecules for viewing directly with the light pen.

3. Larger molecules had to be allowed. The target molecule was rubredoxin with 52 residues, or about 900 atoms.

4. Residue-level operations were installed. For example, residues could be displayed as points connected by lines, by only the connecting bonds, or by the full atomic detail of each residue.

5. Dynamic viewpoint-rotation and other 3-D aids were installed.

3.4 GRIP-75

In mid-1974 E. G. Britton [1977] selected the electronic Richards box problem for his Ph. D. dissertation. Working under the direction of Dr. Brooks, he analyzed the requirements for such a device and designed a set of commands which, together with what was already available in GRIP-74, would be sufficient for this task. M. E. Pique [draft] selected the M. S. thesis topic of nested dynamic rotations, and I continued to develop 3-D viewing aids.

First production use by crystallographers came in mid-1975, hence the name GRIP-75. GRIP now stands for <u>GRaphics</u> <u>Interaction</u> <u>with</u> <u>Polymers</u>. The projected features were substantially complete by late 1975. This involved about 8,200 lines of source program on the host IBM 360, mostly in PL/I, and about 7,800 lines on the satellite PDP-11, mostly in our local PL/I. The features can be summarized as follows:

1. A new user, the crystallographer fitting a molecular model to an electron density map, is served.

2. GRIP-75 is a direct descendant of the GRIP-74 system.

3. Multiple libraries of molecules are provided. This helps the various user groups keep out of each other's way.

4. Dynamic manipulations of user-selected substructures is provided. Overall translation, rotation, and single-level bond-twisting are supported.

5. Interactive contouring allows the user to select his electron density map level and style of display at the terminal.

6. Interactive geometric idealization serves as a different energy relaxation procedure. It considers only the bonded interactions.

7. Nucleic acids as well as proteins can be displayed and manipulated.

8. Large molecules can be handled without substantial degradation of performance. The target molecule was aspartate carbamoyltransferase with about 450 residues in the repeating unit. This is about 4,000 atoms.

9. Large electron density maps can be handled. The size is limited by how much disk space we are willing to use. In practice we have not gone much beyond about 100 x 100 x 100 points. For a three angstrom resolution map, with points spaced every angstrom, this is a cube 100 angstroms on a side.

10. The user can command display lists to be stored on the host System/360 computer's disks for off-line conversion to pen-drawn plots.

11. Chemists with two conformations of the same, or similar, molecules can superimpose them in space for visual comparison.

12. Continued development of 3-D viewing aids.

In the years 1976 through 1978, existing features were upgraded, and support for routine production was enhanced. Indeed, without this work, routine production would not be possible. During this period the source lines that run in production grew to about 10,200 on the host IBM System/360 and to about 8,200 on the satellite PDP-11. This involved the following:

1. A set of electron density maps now comes with each molecule library. This further insulates the various user groups from each other.

2. Eight-level nested dynamic bond-twisting was installed. This is usually active during free-substructure dynamic fitting.

3. Usage of the system is automatically logged for accurate tally by the project administration.

4. More batch programs were written to simplify data conversion for visiting clients.

5. Visiting users are assigned shepherds (department students) to train them and to help them on-line.

6. Several new features were prototyped. I mention only the two that were used productively by chemists during production tests:

   a) An on-line facility displays Ramachandran plots showing the shape of the protein chain in an abstract way that allows easy identification of many conformational features, such as helical and sheet sections of the structure.

   b) Shaded opaque plots of molecular surfaces can be produced off-line on a Versatec Printer/Plotter.

GRIP-75 at its various stages of development is described in these papers (see bibliography):

1. Britton [1977];
2. Britton, Lipscomb, and Pique [1978];
3. Brooks [1977];
4. Foley and Wright [1975];
5. Pique [draft];
6. Tsernoglou, Petsko, McQueen, and Hermans [1977].

During all this time many people worked on many overlapping tasks. Where this occurred in my thesis work, it is covered in detail in the appropriate sections. Table 3.1 may provide a useful overview of the components that make up the system as well as the extent to which I was involved in building them.

TABLE 3.1

GRIP-75 Credits

UNC GRIP-75 MOLECULAR GRAPHICS SYSTEM

Produced by F. P. Brooks Jr.

Directed by W. V. Wright

HUMAN FACTORS DESIGN
E. Britton, W. Wright, J. Lipscomb, M. Pique

APPLICATION SOFTWARE

System nucleus:	W. Wright, M. Pique, J. Lipscomb, E. Britton
Dynamic fitting:	M. Pique, J. Lipscomb, E. Britton
Dynamic viewing:	J. Lipscomb, M. Pique, E. Britton, W. Siddall
Stereo display:	J. Lipscomb
Molecule display:	M. Pique, W. Wright
Density maps:	E. Britton
Plotting facility:	D. Tolle, G. Hamlin
Idealization:	D. Tolle, J. Hermans, J. McQueen, S. Wei
Data structure:	W. Wright, M. Pique, J. Hermans
Geometry primitives:	W. Wright
Data library mgmt:	R. Motley, J. Crawford, L. Brown, G. Kennedy
Software lib. mgmt:	R. Motley, T. Dineen, W. Siddall

SUPPORT SOFTWARE

Coordinators:	J. Foley, J. Leonarz
Inter-computer commun.:	P. Kilpatrick, W. Kerr
CHAT interface:	S. Bellovin
PDP-11 memory management:	S. Bellovin, D. Tolle, R. Motley
PL/C adaptation for PDP-11:	D. Kehs, T. Dunigan
Graphics subroutines:	G. Hamlin, R. Hogan, P. Mullen
Analog to digital conver.:	R. Hogan
Software library mgmt:	L. Nackman, T. Williams

HARDWARE SUPPORT
P. Nichols, J. Ross, P. Reintjes

COLLABORATING USERS
S. Kim, J. Sussman, J. Richardson, D. Richardson, J. Hermans

FINANCIAL SUPPORT
NIH #RR00898, NSF #GJ-34697, AEC #AT-(40-1)-3817, IBM

3.5 RESULTS

From the start of production in mid-1975 to the end of 1978, the GRIP-75 system had 2,725 hours of production time logged by 38 chemists from 17 institutions working on 20 molecules. Production time is defined as time spent at the graphic terminal by a chemist while engaged in chemical work. No programming time by system builders or demonstration time is included in this figure.

Publications thus far by users are given in the bibliography. They are:

1. Beem, Richardson, and Rajagopalan [1977];
2. Carter [1977];
3. Ferro, McQueen, McCown, and Hermans [submitted];
4. Girling, Houston, Schmidt, and Amma [submitted];
5. Holbrook, Sussman, Warrant, Church, and Kim [1977];
6. Kim and Sussman [1977];
7. Kimball, Sato, Richardson, Rosen, and Low [1979];
8. Love, Fitzgerald, Hanson, and Royer [in press];
9. Monaco, Crawford, and Lipscomb [1978];
10. Richardson [1977];
11. Schevitz, Podjarny, Krishnamachari, Hughes, Siglar, and Sussman [1979];
12. Stenkamp, Sieker, Jensen, and McQueen [1978];
13. Sussman, Holbrook, Church, and Kim [1977];
14. Sussman, Holbrook, Warrant, Church, and Kim [1978];
15. Sussman and Kim [1976a];
16. Sussman and Kim [1976b];
17. Tsernoglou and Petsko [1977];
18. Tsernoglou, Petsko, and Hudson [1978];

19. Tsernoglou, Petsko, McQueen, and Hermans [1977];

20. Tsernoglou, Petsko, and Tu [1977];

21. Warrant and Kim [1978].

Clients estimate a two- to sixteen-times speed-up (in calendar time) over use of the traditional Richards box for fitting their molecules to their maps.

GRIP-75 appears to be the first electronic Richards box system with which crystallographers have solved the atomic structure of a protein without the construction of a physical model [Tsernoglou, Petsko, and Tu, 1977].

Chapter 4

HUMAN FACTORS

4.1 INTRODUCTION

This chapter discusses decisions that span more than one class of three-dimensional cues. Human engineering considerations unique to any one three-dimensional cue or to any one class of such cues are covered in later chapters.

4.2 NEED FOR VARIETY

In some cases users of a system are best served by offering them a rich menu of redundant techniques.

The system builder's work is made particularly difficult if individual clients appear at random from an enormous pool of potential clients, if they arrive never having used any computer graphic system before, and if they come on short notice for only a brief visit. There is no time for the man to adapt to the machine or vice versa. The machine must have been adapted to the man in advance of his appearance for the first session. Inexperienced users need to be able to learn techniques quickly. Experienced users arrive with strong prejudice towards their way of working.

In these cases the builder can adapt the machine to the preferences of the users by offering "one of everything". The system builder must also assume that his client will concentrate fully on the application, not on the computer system.

4.3 ACTION LANGUAGE DESIGN

Orthogonality among functions implies that the invocation of one function should not affect the result of a previous or subsequent function. This principle can be generalized to suggest that orthogonal functions may be interleaved in time (simultaneously active) without mutual side effects.

After all, one need not turn off a car's radio to turn on the headlights, or turn off the windshield wipers to roll up a window.

Our users perform the following acts most frequently during substructure fitting (in decreasing order):

1. a series of similar 3-D motions (e.g. repeated translations or rotations) each decomposed into successive 1-D and 2-D motions;

2. alternating one motion with another (e.g. viewing-rotation vs. translation, or translation vs. rotation);

3. re-positioning the hands on different dynamic-motion controls;

4. alternating a series of dynamic motions with a series of interactive commands.

These are really extended forms of motion decomposition [Kilpatrick, 1976] that were not anticipated.

When the user switches from one dynamic motion to another, orthogonality of function allows many dynamic functions to be active at once, so the user, by alternating motions with each hand, may switch between any two functions without re-positioning his hands and without looking away from the molecule. He can also switch pairs of devices simply by repositioning his hands on different devices without the need for a distracting interactive command, although he usually looks away from the molecule.

The user dislikes breaking his rapport with the 3-D portion of the picture (molecule and map), which brings about an interesting adaptation in the task of repositioning his hands (step 3 above). The inexperienced user:

1. looks down at the dynamic device he wants to grab;

2. watches as he reaches for it;

3. looks back up at the molecule as he closes his fingers on the device.

As he gains experience, he reduces the time he is looking away from the molecule by using an interesting intermediate strategy. He:

1. reaches for the device;

2. glances at his hand as it bumps into the device;

3. looks back up at the molecule as he closes his fingers on the device.

Finally, after some weeks, he can often grope for devices without looking.

The use of few devices (or stations for the hands), each controlling many motions, avoids delay for switching hands. Orthogonality of function is not violated, and the user's need for motion decomposition is satisfied, if each of the multifunction devices controls a set of related functions (e.g. rocking rate vs. amplitude, x,y,z translation, etc.), if each degree of freedom of the device corresponds to a degree of freedom on the screen (e.g. x,y,z orthogonal axes), and if each degree of freedom of the device can be actuated separately, without spilling over into another.

When users intersperse dynamic manipulations with interactions, simply leaving the dynamic devices active during and after these interactions avoids any delay in returning to the dynamic manipulations.

A guiding principle in the button-language design, and indeed for all other interactive functions, was set by E. Britton [1977]. All buttons were to be <u>push for on</u>. There were to be no <u>push-on--push-off</u> buttons, and none were to behave in different ways in the context of another button having been pushed. In other words the button language was to be <u>context-free</u>. Hence, for example, the actions of rocking, spinning, etc. each have an <u>on</u> button; the set has a <u>reset</u> button.

Orthogonality among the functions further avoids the side-effects of context. An interaction always has the same effect regardless of the status of other features. For example, the stereo display mode may be changed regardless of the mode of contour display or molecule display, regardless of whether dynamic fitting is in progress or not, and independently of the rocking or spinning that may be in effect.

For each button and joystick there is exactly one function, the one indicated by its label. No device serves more than one function. There are no "shift" keys as on a typewriter. This avoids the complexity of knowing the context in which an action is made, and has worked well for this application. Commands with many parameters are often really multi-purpose functions, whose action depends on the context formed by their operands. Seeking to avoid this context encourages short commands (few parameters) and a proliferation of buttons. Having many simple buttons, instead of a few complex ones, helps the user select, rather than evoke, commands, because he does not have to remember as many operands.

The stereo-mode selection and the rocking or spinning mode change are the only interactive operations of the three-dimensional package. They take no interactive operands. Thus, they are completely orthogonal to the selection of a prefix or a postfix interaction language, and they are easy to operate.

More generally, the most frequently used interactive commands in GRIP-75 were designed to require only one parameter each. (We omit further discussion of the dynamic controls, whose existence makes possible this reduction in parameters, from here to the end of this section.)

First, the commands are so short that the choice between a prefix or a postfix interaction language matters little.

Second, the commands are sentence-structured by virtue of being short. Foley and Wallace [1974] (page 16 above) do not allow parenthetical remarks in their model of a sentence-structured command language, perhaps because it raises the specter of "context". GRIP-71, GRIP-74, and GRIP-75 allow parenthetical remarks. During command setup, the user may think of something else he wants to do. He may either defer it until he completes the first command at the cost of having to remember the second, or he may perform the second command before continuing with the first at the cost of juggling the context of his actions. Neither is desirable, and both derive from having long command sequences. With short commands these disadvantages are far less significant and seldom arise. The user is thinking about his application. If the commands are short, he is distracted less by mentally deferring a chemical goal than by physically altering his input to the system to insert a parenthetical remark. Indeed, many users avoid long commands, and some try to use them and fail, because they refuse to be distracted from their application by the mechanics of the computer system.

Selecting, say polarized-mirror stereo, causes the system to leave the "home state" (the state between sentences) of monocular display for some time, perhaps for the entire session. The user, however, believes he has returned to the home state if all his commands work as before (if they are orthogonal to the stereo-mode command). The light pen or dynamic fitting may have to do some strange things indeed, internally, if they are to behave the same way as before, but if one does not tell the user, he will never know. The important "home state" is that of the user, not of the system.

4.4 DEVICE LAYOUT

Figure 4.1 shows how the rocking and spinning and the stereo buttons of Figure 4.2 are laid out. The defaults are RESET (no rocking or spinning) and MONO (no stereo). These were chosen as the defaults for three-dimensional viewing when the scene is first presented because they are what the chemists normally use in production. They are also distinguished as the "off" states of these features. The defaults are indicated by "X" in the diagram. They are together on the right-hand column of the button box which can be easily found by touch. Immediately to the left of each of these "push-off" buttons are the mutually exclusive "push-on" features of automatic-rotation and stereo.

Actually, RESET may be a poor label for the "off" state of the rocking and spinning package. Several users have mentioned a certain reluctance to "try it" with several hours' work seemingly hanging in the balance. The odd thing is that some of them appeared unconvinced by my explanation that it was "harmless". F. Brooks suggests that HOME VIEW might seem less intimidating.

The lowest row of the button device is most easily reached by touch alone. It is therefore reserved for rocking and spinning, which preserve "visual continuity" [Foley and Wallace, 1974] (page 17 above). That is, the picture flows smoothly from one viewing-direction to another.

Switching stereo modes involves visual discontinuity in GRIP-75, because the user must change the mechanical viewing aids. Having to look down to find the button to activate the software therefore introduces no hardship.

Figure 4.3 shows the layout of the slider controls' seen in Figure 4.4. The user sits somewhat to the left of the controls and reaches them with his right hand. The most popular functions are placed where they may be found by touch. The leftmost slider is the most conveniently reached

----

' Model 221 SL 10K, SLIDELINE(TM) potentiometers.

Duncan Electronics Inc.
Costa Mesa, California

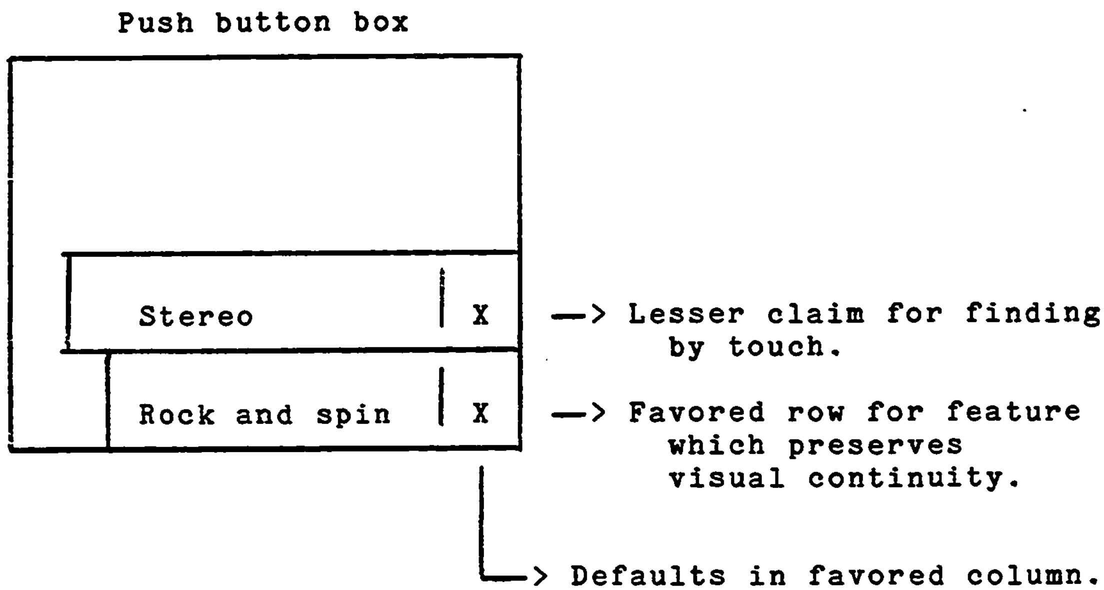

Figure 4.1: Push Button Layout

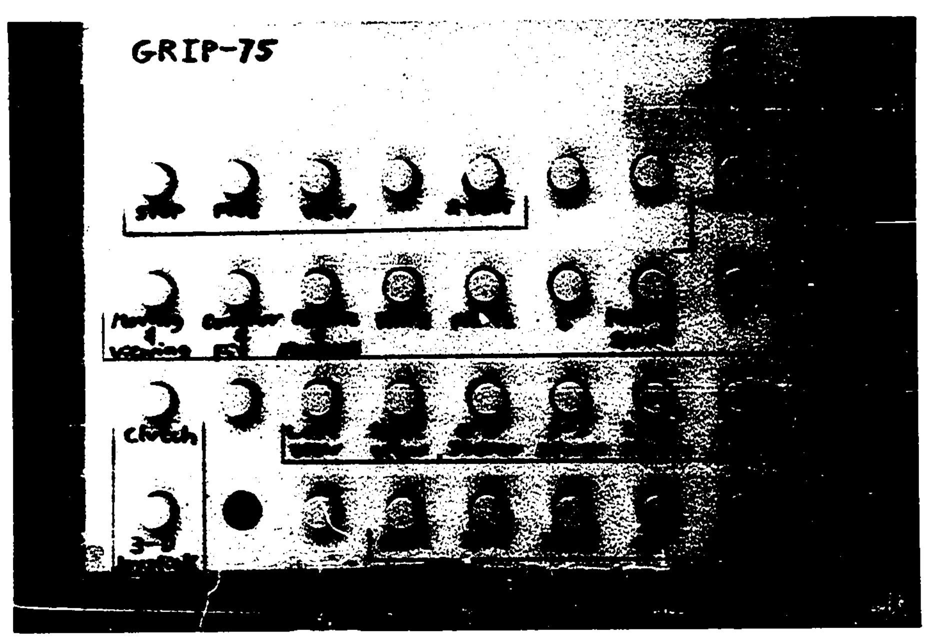

Figure 4.2: Push Buttons

and can be found by touch. The rightmost slider also can be found by touch. Chemists sometimes make good use of this. Even when they do not, they can find these sliders rapidly, and they always look back up to the screen when their hand touches the slider, keeping their eyes on the screen as they make the manual adjustment. Figure 4.3 also shows that similar controls are grouped together [Woodson and Conover, 1964] (page 14 above).

4.5 THE LIGHT PEN

GRIP-75 uses a light pen (Figure 4.5) for selecting verbs from a menu on the screen, for picking atoms or bonds as operands for these verbs, and for picking contour lines to be erased. Touching a switch mounted on the light pen invokes the picking action.

Because the light pen is involved in so many functions, substantial effort was applied towards its usability. Joyce's [1967] (page 21 above) ideas on a reactive display were implemented, both in giving an instantaneous indication of what the light pen sees and in suppressing that indicator when the light pen points to an object that is not pickable.

We have a circle that at any time marks a pickable object as the light pen points to it. When the light pen sees nothing pickable, the circle winks out. If the circle marks something the user was seeking to pick, before it winked out, then he is stirred to immediate action. He lunges for the light pen switch even though it is often too late still to make the pick. All our users behave so.

To accommodate them I installed a grace period of approximately one-half second during which the light pen remembers what it last saw, whether or not the indicator is still visible. It will pick that object if the light pen switch is depressed. If during this time it sees another pickable object, it immediately positions the circle indicator on it to announce that it will pick this new object instead if the switch is depressed.

Keeping the circle character visible during the grace period is worse than not having a grace period. If the user has just stopped the pen over an object he wishes to pick, the pen may have traveled too far. The user thinks he has plenty of time for the pick, but he may be on the grace period and only have a fraction of a second.

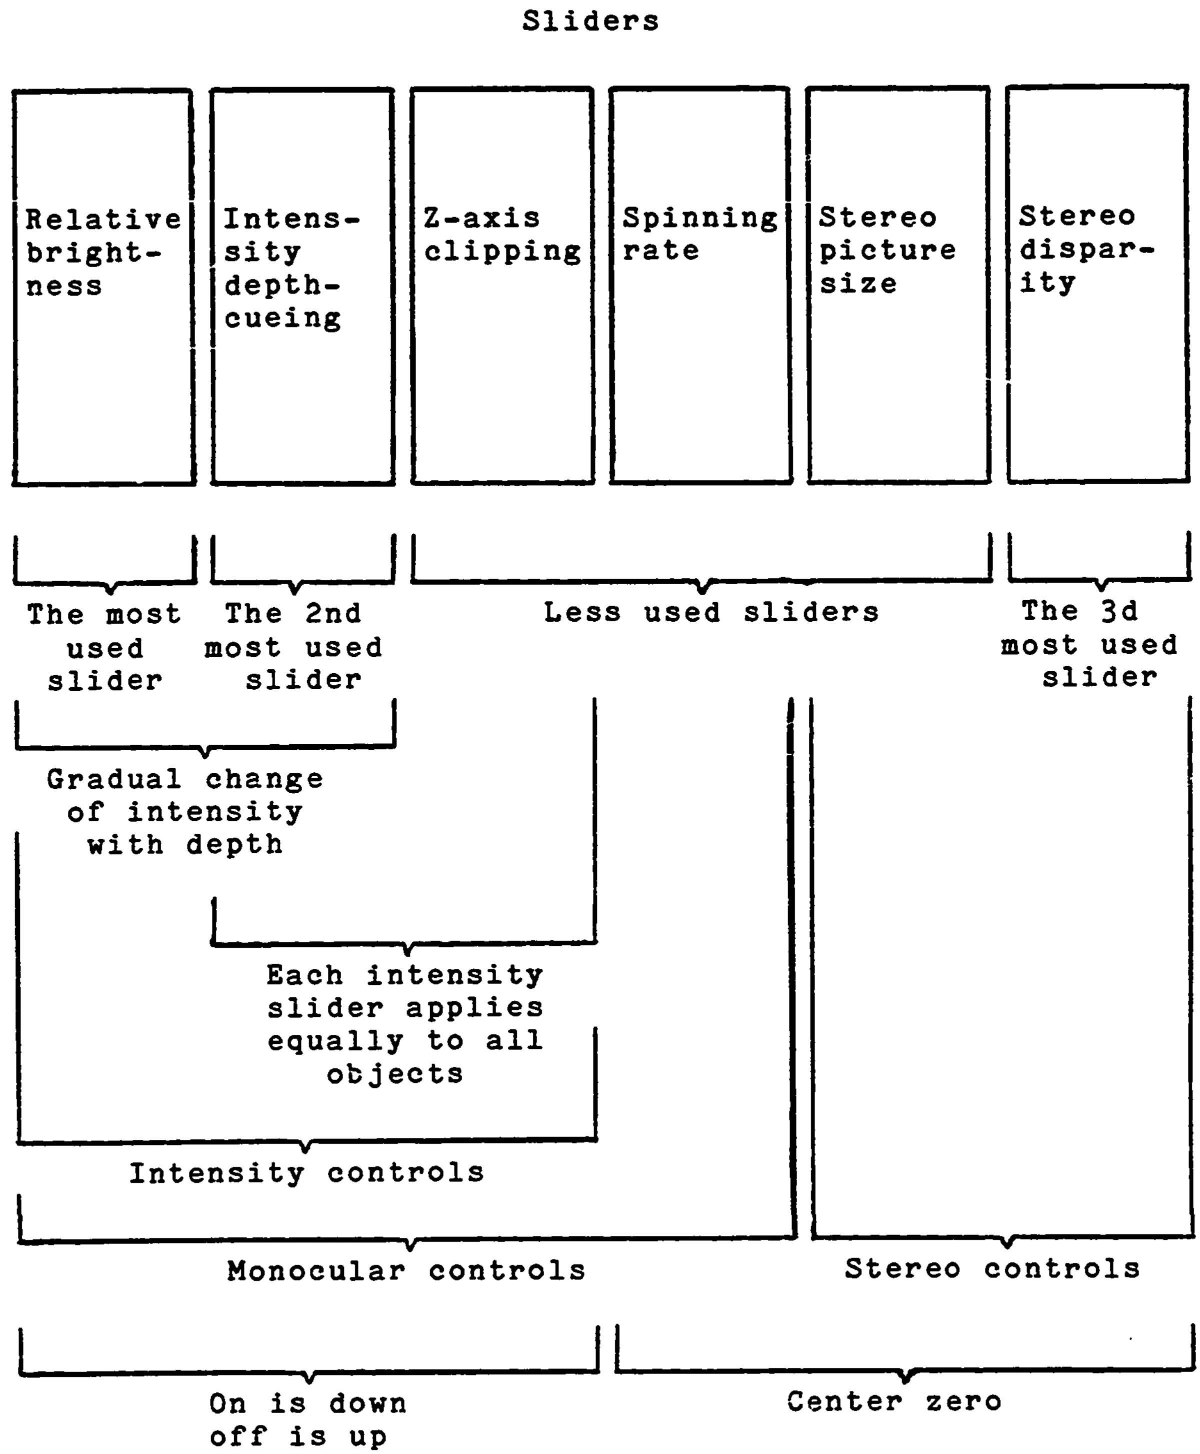

Figure 4.3: Slider Grouping

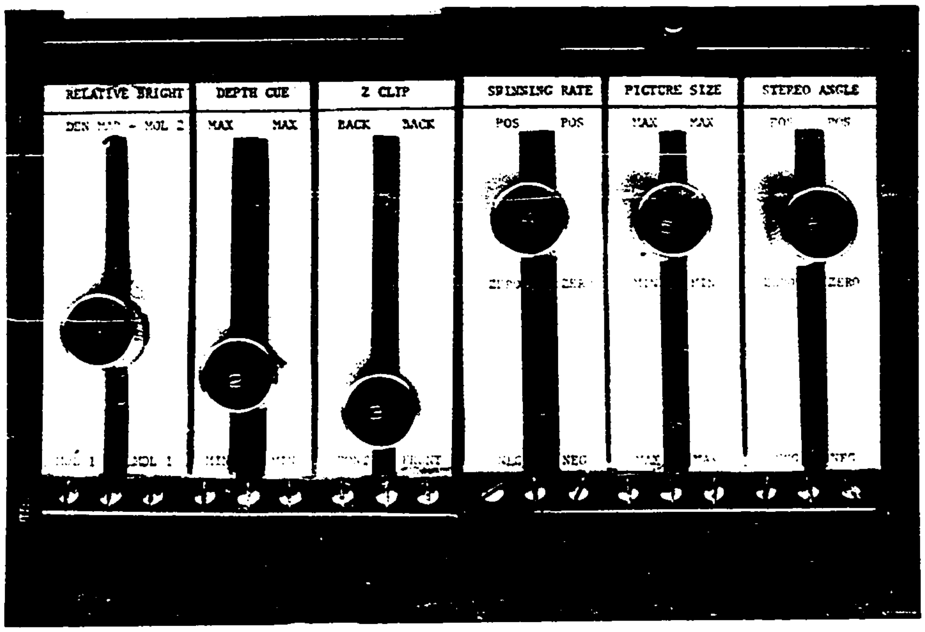

Figure 4.4: Sliders

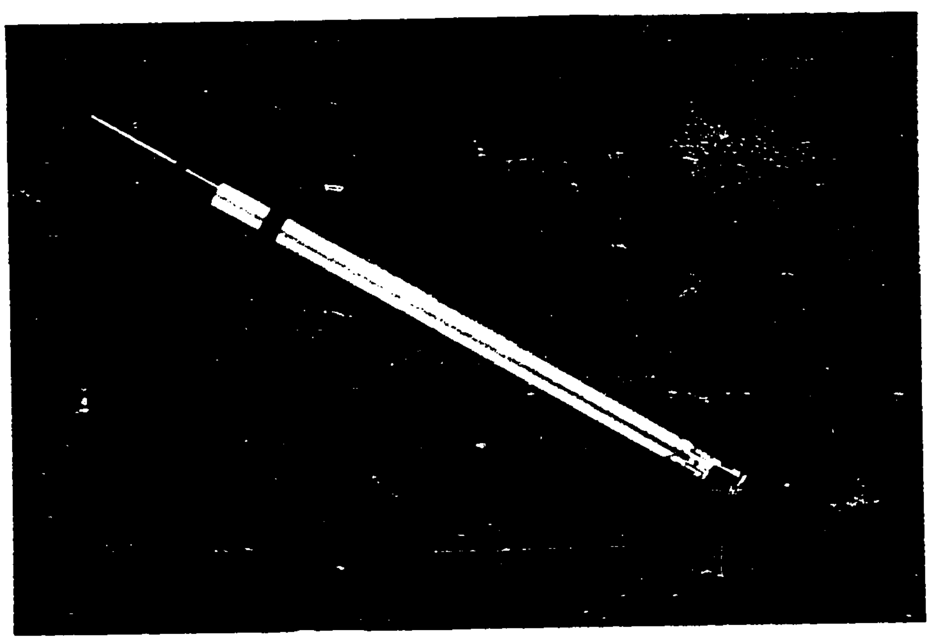

Figure 4.5: Light Pen

E. Britton [1977] made an interesting use of the grace period. With it the user may pick a blinking line with confidence. Because of this, Britton was free to use blinking as a preview for light pen editing (deleting) of contour lines.

When a pick is made, the circle changes to a square character that disappears when the host computer acknowledges the interaction. The square and the circle (but not the grace period) were implemented by M. Pique.

I believe the above issues of feedback and grace period apply to picking devices in general, not just to the light pen.

4.6 MISCELLANEOUS HUMAN FACTORS

The idea of a "reactive display" mentioned above can apply to devices other than the light pen. The rocking and spinning and the stereo mode buttons light up directly under the user's finger as he pushes them.

Reasonable reaction to user input sometimes requires specialized technique. One approach widely applied to velocity controls in other systems is to present a response of the sign-preserving square ($d * |d|$) of the control's displacement (d) rather than simply the displacement. This allows fine control of small values and coarse control of large values. I use it in the substructure rotation control, in the spinning rate control, and in the rocking rate and amplitude control.

The button-lighting convention for rocking and spinning selection is the same as for stereo-mode selection. That is, if the rocking and spinning buttons are enabled then one of them is illuminated, the one currently active. If the rocking and spinning package is disabled, none of its control buttons are lighted.

As one might expect, the rocking and spinning controls and the stereo controls are disabled when in a two-dimensional mode, for example the library mode used for storing and retrieving molecules. The built-in lights within the control buttons are all off in this case, indicating that the buttons are disabled. Pressing them in this state has no immediate effect on lights or system. Such depression also has no effect on the "memory" built into these features that returns the picture to the previous automatic rotation and stereo condition when a three-dimensional mode is re-entered.

Robustness ("bullet proofing", [Martin, 1973]) prevents user errors from shooting down the system. This problem is side-stepped by the interactive 3-D viewing commands, because all user activities are valid or do nothing. When disabled they do nothing, and when enabled they have no error conditions to trap out because they take no interactive operands and because any button push is valid any time. The dynamic controls may always be set in any position desired; there are no rules to follow. The software ranges of the devices allow a generous margin at each end against the possibility of voltage drift significantly altering the behavior of the physical devices.

Right-hand rotation of objects (not the coordinate systems in which they are embedded) is defined as positive for rocking and spinning, consistent with traditional usage in mechanics and with various interactive features in the system. The rate axis of the rocking control is the up-down axis so as to be consistent with the slider spinning rate control.

Execution speed is critical for smooth dynamic operations. When no input motion is in progress for an active function, say substructure rotation or intensity control, the program shuts itself off or goes to sleep. The system continues to use the last values produced for the function. When the user again moves the control, the function wakes up and calculates new values for the desired effect.

W. Siddall recoded some of the stereo programming in PL-11, a high-level assembler language, for greater efficiency, yielding smoother motion. Other than that, all of the programming for the three-dimensional viewing techniques I built was done in PL/I. The IBM System/360 programming was done in "F"-level PL/I, and DEC PDP-11 programming, the bulk of the work, was done in PLCD, a locally developed variant of Cornell University's PL/C.

Response time for the interactive operations of the three-dimensional techniques was tuned to what the user requires.

There is no great hurry to switch stereo modes, and so that is handled by the host IBM 360. This follows the general design decision to perform all interactive operations in GRIP-75 on the host and all dynamic operations on the satellite computer. Changing stereo modes takes about 1 to 2 seconds. But if polarized-mirror stereo is involved, the time climbs to 5 to 10 seconds, because the characters marking the atoms must be removed lest they reflect upside-down in the image combining mirror.

Rocking and spinning maintain visual continuity and so must react quickly. For example, when the user punches the HALT button, he expects the motion to halt "here" and "now". The interactive response must be immediate before the rotation proceeds any farther. This forced implementation on the satellite PDP-11.

Chapter 5

DYNAMIC MOTION

5.1 INTRODUCTION

Of the techniques I have tried, smooth rotation of varying or of constant speed is the single most effective aid to three-dimensional perception in molecular graphics. The three-dimensional perception arises from the kinetic depth effect (KDE) (Wallach and O'Connell [1953], page 12 above). Smooth manual motions are also the best tool for three-dimensional eye-guided manipulations.

Dynamic manual manipulation, which is only as accurate as the user's hand, is inherently less precise than numerically coded interactive manipulation. Yet it carries with it motion, the best cue for perceiving depth. Indeed, the human visual system has special processors for motion that help the viewer understand what he sees [Sekuler and Levinson, 1977].

Our clients tend to alternate 3-D dynamic viewing with dynamic manipulation. Yet that manipulation itself often sustains the 3-D perception of the object being moved. The memory effect of the KDE (Wallach, O'Connell, and Neisser [1953], page 13 above, and Ortony [1971b], page 26 above) allows the user to perceive the still portion of the picture in 3-D for a short time (for pictures as complex as ours) after the viewpoint-rotation stops. Even so, dynamic 3-D rotation of one part of the picture against the remainder does in a certain sense reduce the problem of seeing what one is doing, when only the manipulated object is in motion. Thus, the issues of viewing are inseparable from those of manipulation. This section therefore describes some of both.

Dynamic motions should be effortless so that the user may think about his application, rather than about the system. Realism is one way of achieving this. The user should be lulled into believing that he holds real, physical objects in his hands, not, as is actually the case, hand controllers for a representation of the objects. There are limits of course. The user wants realism without the disadvantages of the real world, and with the advantages of computer implemented models.

5.2 MANUALLY-OPERATED VIEWPOINT-ROTATION

Dynamic viewpoint-rotation is the single most effective 3-D cue in this system. Its usefulness is greatly augmented by building the physical rotation controls into a mechanical linkage that the apparent motion on the screen mimics. One of our users has commented that she did not think that the mechanical linkage was important until she used a graphic system that did not have one. It is so effective that users typically will rock the picture constantly by hand for a half-hour at a time rather than invoke the automatic rocking feature. With brief interruptions for interactive operations they work this way for hours.

We call our linkage' the toothpick (Figure 5.1) in honor of the thin shaft extending from it. It has two degrees of freedom. Complaints by users strongly suggest that a third degree of freedom should be supplied so that the molecule may be turned "right side up". Crystallographers become so familiar with their molecules that they develop a sense of "up" and "down".

The motion of the physical device is apparently one-to-one with the picture. The illusion is that of having seized the picture bodily and that the picture is a real object held in one's hand. There is a simple directness to this position-based scheme which aids the illusion. The effect was not as strong when we tried a spring-return velocity device as a substitute. A position device allows one's kinesthetic sense to help swing from one view to another and back again swiftly and accurately. This is related to "muscle memory" [Hansen, 1971] (page 15 above).

Dynamic motion is achieved by rapidly generating successive transformations (translation vector or rotation matrix for example). Each new transformation supersedes the previous one and maps the picture data onto a slightly different place on the screen. The image thus moves in discrete steps, but if they occur rapidly, smooth motion appears to

----

' Modified Model 525 joystick.

Measurement Systems, Inc.
121 Water Street
Norwalk, Connecticut 06854

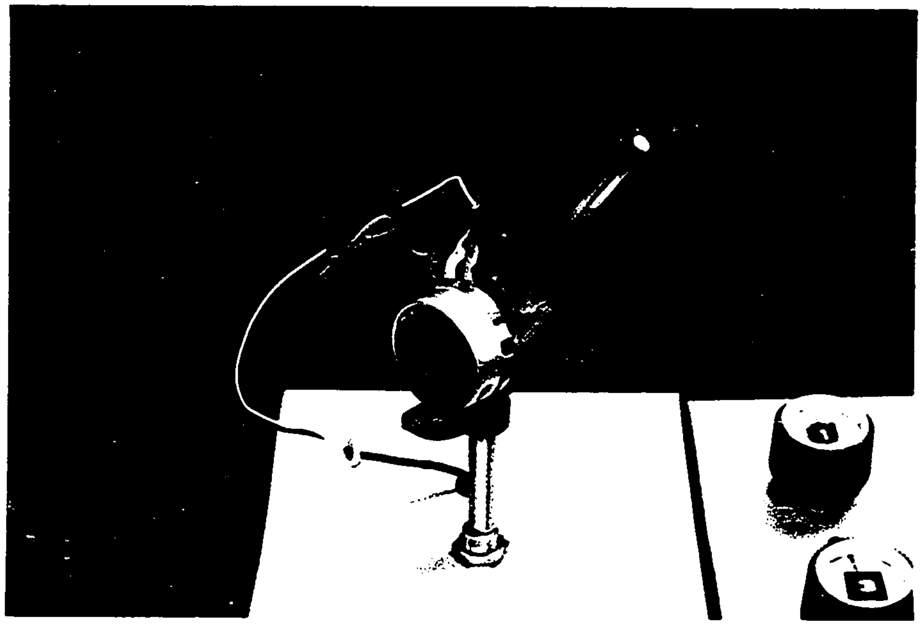

Figure 5.1: Toothpick for Viewpoint-rotation

take place. Just how rapidly is "smooth" is not a simple matter. This question is covered at length in a later section.

We now enter the realm of mathematics where single large letters represent matrices (R), isolated small letters represent scalars (d), small letters adjacent to a large letter represent subscripts of a matrix (uRv), and underlined small letters represent vectors ($\underline{v}$). A blank space between letters indicates non-scalar multiplication. " := " indicates assignment.

The rotation transformation of our display is

$$\underline{v}' := R\ \underline{v}$$

where $\underline{v}$ is a three element vector describing a point in the model frame of reference (the picture data), built by an interactive operation, and $\underline{v}'$ is its dynamic position on the screen (x, y, and z). The program sets $\underline{v}$ and R, and the user sees $\underline{v}'$. x is to the right, y is up, and z is towards the user (Figure 5.2). This forms a right-handed coordinate system.

The program generates successive rotation matrices. These are 3x3 arrays by which all points ($\underline{v}$) must be multi-

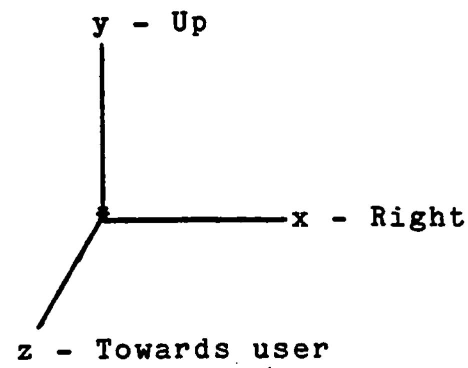

Figure 5.2: Coordinate System

plied to appear on the screen ($\underline{v}'$). The simplest rotation matrix is

$$I = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

the identity matrix, which results in

$$\underline{v}' = \underline{v}$$

with the picture on the screen displayed in the orientation of the picture data.

Rotation about the $\underline{x}$, $\underline{y}$, and $\underline{z}$ axes is achieved with the appropriate rotation matrices. Right-hand rotation of the image about the $\underline{x}$, $\underline{y}$, and $\underline{z}$ axes separately by angles x, y, and z is shown below.

$$X = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos x & -\sin x \\ 0 & \sin x & \cos x \end{bmatrix}$$

$$Y = \begin{bmatrix} \cos y & 0 & \sin y \\ 0 & 1 & 0 \\ -\sin y & 0 & \cos y \end{bmatrix}$$

$$Z = \begin{bmatrix} \cos z & -\sin z & 0 \\ \sin z & \cos z & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Nested rotations posed a problem because our display unit can pass segments of the display list through only one rotation matrix on the way to the screen ($\underline{v}' = R\ \underline{v}$). That rotation matrix must contain within it the rotation nesting. For the toothpick we want

$\underline{v}' := Y\ X\ \underline{v}$.

Fortunately the product of two rotation matrices is itself a rotation matrix. Thus, if the satellite computer can perform

$R := Y\ X$,

then the display unit can perform

$\underline{v}' := R\ \underline{v}$.

X Y may be analytically combined once to avoid matrix multiplication at run time.

I assisted M. Pique in doing this. I shall not call forth the details since Pique [draft] describes a superior scheme he later built alone.

5.3 AUTOMATIC VIEWPOINT-ROTATION

5.3.1 Introduction

Automatic rotations are commonly provided in molecular graphic systems. Three forms of automatic viewpoint-rotation in GRIP-75 complement the toothpick, a device for manual viewpoint-rotation. Rocking and spinning are continuous once started and must be stopped by user command. 90-degree turn rotates the picture smoothly through 90-degrees. A push button is provided for interactive selection of each function, and dynamic controls determine their speed and direction.

5.3.2 90-degree Turn

90-degree turn is the most useful automatic-rotation feature. About a half of our clients use it. It is a successful alternative to side-by-side orthogonal views [Ortony, 1971b] (page 27 above). The 90-degree turn is used during substructure fitting or, more generally, in positioning a point in space. Occasionally chemists use it strictly for viewing where the halt every 90-degrees seems to help them keep track of where they are. Then too, they sometimes line up an interesting feature into or across the screen and gain some benefit from a second view exactly 90-degrees away.

As with Kilpatrick [1976] (page 32 above) our users tend to decompose their 3-D positioning into a series of 1-D and 2-D motions, in our case in the user's frame of reference exclusively. Before the 90-degree turn feature was provided, they often positioned in the plane of the screen, turned the scene 90-degrees as best they could with the toothpick, and then continued fitting in the plane of the screen as before. These steps alternated rapidly, but the motions were inaccurate because the 90-degree turn was done by eye, so the successive steps inadvertently interacted. Now that this is automated, fewer iterations are required with a concomitant increase in productivity.

Wright [1972a] (pages 29 and 43 above) showed that if the turn is made suddenly, the chemist loses track of the features he sees and will have to spend some time at the new orientation re-understanding the picture. The user's rapport with his picture is broken when it suddenly disappears, reappears, or jumps from one place to another. This is called loss of visual continuity [Foley and Wallace, 1974] (page 17 above). Visual continuity is maintained in the 90-degree turn feature by rotating smoothly through the turn rather than taking it in one jump.

Some users persist in doing the 90-degree turns manually. They get too wrapped up in what they are doing to look down for the 90-degree button. For them we have a three-axis gnomon in the lower right corner of the screen and a three-axis jack-shaped position indicator (Figure 5.3) at the Noll box position (when not fitting), which turn with the picture. The axes of these can be successively lined up into the screen to allow more accurate 90-degree turns than kinesthesia alone permits.

Figure 5.3: Jack-shaped Position Indicator

5.3.3 Rocking and Spinning

Automatic rocking and spinning are not very popular. Mrs. Jane Richardson (Anatomy Department, Duke University), her students, and other users they trained use these features routinely and like them. Much as with rotating-shutter stereo used by Dr. Kim (Biochemistry Department, Duke University) and his research group (among others), the "liking" of a given system feature is often infectious among close collaborators. The source of this infection merits further study. The Richardson group, however small and loyal, is nearly the only group interested in rocking and spinning.

The idea is to provide the user with a "third hand" for his viewpoint-rotation while his other two hands are busy elsewhere with fitting or relaxing at his side while he contemplates the picture. Since users tend to alternate viewing and manipulation, and since manual viewpoint control provides powerful kinesthetic feedback, automatic rotation is of only limited popularity. In practice spinning is unattractive at all times because it changes the viewpoint too drastically for the user to keep track of the spatial relationships he is contemplating. Rocking is better behaved for viewing, and therefore more often used than spinning. Also, rocking is a help to some people during bond-twisting but causes trouble with unconstrained translation. The usual reason people give for having trouble is that the axes of manipulation fixed in the user space change constantly with respect to the viewing space. (This is not a problem with bond twisting which is nested within the viewing space.) But actually the errors due to this are very small, and I think that the more fundamental problem is that users have trouble perceiving two motions at once.

Since most users want to alternate viewing with manipulation, then perhaps rocking and spinning should automatically disengage whenever a manual input is made on a dynamic device.

5.3.4 Human Factors

The term phantom device [Pique, draft] best describes automatic rotation. In between the toothpick and the picture, the ghostly hand of the system rocks or spins the picture as if the device were being moved. Indeed, it simply adds or subtracts from the setting of the toothpick. The user's hand is no longer connected directly to the picture. It turns under his grip as if his fingers were slippery and numb. This may be a clue to the unpopularity of rocking and spinning.

Users of the 90-degree turn feature do not suffer from this feeling of having greasy fingers since they tend to leave the toothpick alone during the turn.

When the 90-degree turn is complete or when rocking or spinning is halted, a new problem appears. The toothpick is again rigidly connected to the image of the model, but the horizontal axis of the device no longer coincides with the motion on the screen. The horizontal axis of model rotation has been rotated with respect to the mechanical linkage of the toothpick. In about equal proportions users

1. do not notice that this has happened,

2. notice it but do not care one way or another, or

3. do not like the sensation of the model turning one way when one axis of the hand control moves the other.

For the benefit of this last group a reset button is provided. When pushed the picture rotates back in line with the toothpick. Our clients who need it find the reset feature easy to get along with.

There is a back-front ambiguity inherent in the picture. The user's perception is deficient even with dynamic viewing and intensity control helping. He may accidently and frequently perceive the near objects as far and the far objects as near.

The back-front ambiguity problem does not detract much from the usability of dynamic translations, since the user performs translations mainly in the plane of the screen. Even if he incorrectly perceives the near objects as far and vice versa, the benefit he receives from translations is unaffected.

Bond-twisting while rocking is similarly unaffected by a back-front exchange of the image in the user's mind.

The vertical axis in the user's frame of reference was selected for automatic rotation and for the highest level of manual viewpoint-rotation. Dynamic viewpoint-rotation is perceived partly as turning the model in one's hand and partly as moving one's head from side to side around a stationary object. Left-right motions are particularly effective 3-D cues because of an effect called "head motion parallax". Experiments with random dot stereograms [Ross, 1976] demonstrate that the brain has a special horizontal motion 3-D processor that aids 3-D perception when the head is moved from side-to-side. We are hoping to cash in on it or at least to pick up any spare change it tosses our way.

The automatic rotation package maintains visual continuity throughout all of its features and their mutual interactions. No combination of button pushes can make the picture jump suddenly through a perceptable angle. After all, real objects do not behave so.

Real objects also have mass and so does one's head, whichever one thinks of as moving. Here visual continuity can be extended to embrace velocity as well as position.

Thus, automatic rocking is sinusoidal. Horizontal sinusoidal rocking is said to be perceived as linear with a quick reversal at its endpoints (pages 32-33 of Johansson [1950b]).

Also, all forms of automatic viewpoint-rotation take almost a second to accelerate to full speed. I installed this feature in response to complaints from co-pilot users. Many of our users work in teams. There is the pilot, who operates the controls, and the co-pilot, who makes notes and advises the pilot. The pilot pushes the button, usually the 90-degree button, and is grateful for the immediate response. The co-pilot's senses, however, are jarred by the sudden motion. He "was not expecting it" in quite the same way the pilot was. Sudden changes in velocity often cannot be distinguished from gradual changes [Levelt, 1962], but this appears not to be the case when starting from a full stop. A slow startup keeps everyone satisfied and is easily programmed. Halting is a different matter. When commanded by the user it really ought to be immediate so as to stop where the user wants it to stop. As for the automatic halt at the end of a 90-degree turn, I have no such excuse. A gradual halt seemed more difficult to program than it was worth. Incidentally, slow startup and stop come free with a spring-return rate-control device.

One problem remains with the one-second slow startup. The acceleration simply does not appear to be uniform. It is. The velocity is ramped linearly, but it looks more like a one-third second acceleration, followed by one-third second of constant velocity, followed by one-third second of acceleration, peaking to the terminal velocity. Independent programming and observation by M. Pique reached the same unhappy result, largely excluding the possibility of programming error. Actually though, such faulty perception of accelerated motion is normal (page 174 of Schmerler [1976]). Further studies involving eye-movement recordings show that the eye tracks the target accurately (page 182 of Schmerler [1976]). The problem must be further back in the brain.

5.3.5 The Button Language

The buttons HALT, SPIN, ROCK, 90 DEG, and RESET are laid out in a row with the RESET button at the edge of the button box where it is easiest to find by touch (Figures 4.1 and 4.2).

The SPIN button causes the picture to rotate continuously about the vertical axis, ROCK causes it to oscillate sinusoidally about its orientation when the ROCK button was pushed, and 90 DEG turns the picture smoothly through 90 degrees and then halts. HALT stops the automatic rotation instantaneously, even though that may leave the horizontal axis of the toothpick out of kilter with its effect on the picture. RESET rotates the picture back in line with the toothpick orientation and then halts.

The RESET button is distinguished as the _off_ button for automatic rotation. It is more "off" than the HALT button, which simply stops the motion, because RESET does not leave behind evidence of previous actions, as HALT does when one discovers that the effect of the toothpick has been rotated. RESET is also the true "off" button internally, because after RESET has realigned the picture with its viewpoint controller, the phantom device for automatic rotation may be disconnected until needed again. It has been "put to sleep" meaning that it no longer requires any update loop time for itself.

The the button layout was determined by pairing up related functions. Here are some examples. If the picture is either halted or rocking, it may be set spinning with the SPIN button. Then, when a satisfactory orientation is reached, it may be reset to HALT or ROCK. Alternatively, if the picture is rocking, it may be commanded to rotate 90 degrees. Then when it halts, it may be set back to rocking again. Finally, 90-degree turns may be made one after another, as they usually are, or alternately with RESET, which is more like what I had expected. The idea was that users would look down, rest two fingers on the desired pair of buttons, look up, and operate the buttons by touch, with the other hand still free to continue whatever it was doing. In practice users almost never work this way. They look down, poke at one of the buttons, and immediately look up.

Why does a rocking picture not revert back to rocking after a 90-degree turn? Why does it always halt? Why does it not remember? After all, the 3-D viewing mode of the system remembers if the picture was spinning and which stereo mode was in effect when it returns from an excursion to library mode (a different picture, 2-D with lots of text) where the user may have loaded or saved a molecular structure. Does not the principle of orthogonality suggest that a 90-degree turn should return the picture to rocking if it was rocking to start with? Ah, but what about a 90-degree turn of a spinning picture reverting to spinning when it is done? And what about simultaneous rocking and spinning? Unfortunately this matter of "smart programs" is not something one can do half-way without running the risk of confusing the naive user. I did not have a clear vision of what constitutes "good smarts" here, and so I opted for none at all.

The rocking and spinning buttons have n-key rollover. That is, the depression of a button invokes its function independent of whether or not one's fingers have been lifted from the other buttons. Lack of rollover decreases orthogonality among functions, it encourages users to "look and poke" rather than to rest their fingers on the buttons and work by touch, and it greatly amplifies the failure of a single sticky key while simultaneously enlarging the area of search for the sticking key, since none respond. N-key rollover has long been standard on keyboards of distinction. Its implementation in PL/I requires less than a half-dozen lines of source program and enough storage to hold one bit per button (its previous state, up or down).

5.3.6 Dynamic Controls

There are two dynamic controls for automatic rotation. A slider sets the rate of spinning and of 90-degree turn. A 2-D non-spring-return joystick[z] (Figure 5.4) sets the rocking rate and amplitude.

The reset rate is the greater of these two. However, if both are small, a minimum speed (20 degrees per second) is imposed so that the reset does not take agonizingly long to complete if the user has not troubled with the rate control. This comes to a worst case 9 seconds for 180 degrees. The reset direction is whichever way is closest to the toothpick.

The 90-degree turn feature is given the same minimum speed, which works out to 5 seconds for a right angle turn. The continuous spinning rate can be set arbitrarily low.

Work by Shephard and Metzler [1971] suggests a maximum turning rate of 60 degrees per second. In this system the upper limit is set at 70 degrees per second, which with a slow startup comes to 90 degrees in 2 seconds. When set faster the user is startled. Indeed, one did cry out in alarm.

The slider for the rate of 90-degree turn and continuous spinning has its zero in the center of its range. In the upper half of its range the spin vector is up

----

z Quadraphonic balance control by Panasonic.

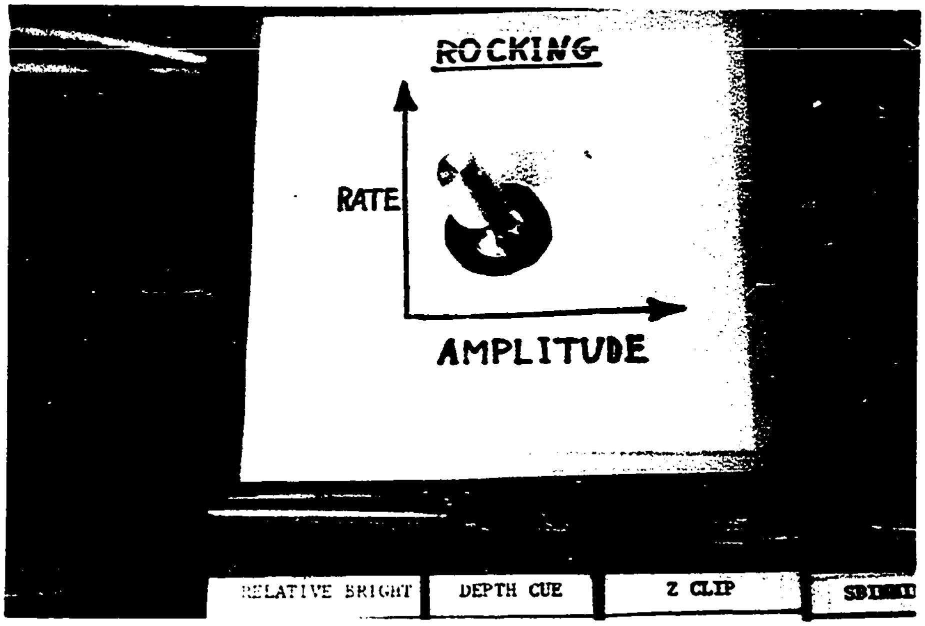

Figure 5.4: Rocking Rate and Amplitude Control

(counter-clockwise when viewed from above); in the lower half it is down. This is the spin vector of the object on the screen, not of the coordinate system in which it rotates.

Like the fist control, the spinning-rate control is set for d*|d| effect, the sign-preserving square of its displacement.

For rocking, the two degrees of freedom, rate and amplitude, are placed on a single 2-D control. Rate is the vertical axis to match the spinning control. This leaves amplitude going left to right.

Center zeros were found to be too confusing for the rocking control. The problem is that the rocking direction is meaningful only for the first cycle. After that the advantages evaporate leaving only the disadvantages. Thus I settled for offering only the first quadrant with zero at the lower left. Why there? Zero amplitude and rate correspond to "off". "Off" is down and to the left and "on" is up and to the right [Woodson and Conover, 1964] (page 14 above).

Naturally the speed control is d*|d|. The amplitude control is also d*|d|, because one wants sensitive control over the more common small amplitudes and coarser control over the larger ones.

Finally, consider the question of amplitude compensation. At first I had the RATE control simply set the period of oscillation. When this control is adjusted satisfactorily, and then the amplitude is decreased, the perceived motion quite suddenly becomes disappointingly slow.

I guessed that one's perception of "speed" is therefore the rate of motion across the screen. I programmed that to remain constant by increasing the frequency of the oscillation as the user decreases the amplitude. Unfortunately, that compensates too much, and the user perceives the motion as increasing with decreasing amplitude. I speculate that the proper change in the period to compensate for amplitude changes must be based on some other perceptual feature. I have no guesses as to what. Some light has been shed on this matter by Johansson [1950b]. This is a promising area for future work.

5.4 REVERSED AND ERRATIC APPARENT MOVEMENT

5.4.1 Introduction

Motion appears to be smooth when the update rate (frequency) is a fixed submultiple of the refresh rate, where the refresh rate is constant and the update rate is fast enough. The breakdown of user satisfaction caused by a progressively reduced update rate is a complex event that is the subject of this section.

If an update cannot be completed with each refresh, then one will have stumbled directly into a zoo of perceptual oddities, which are mildly annoying to see and irritatingly difficult to subdue.

We shall meet them now one at a time. We start with one refresh per update and no problems, then the effects of lower update rates, and finally, the consequences of not synchronizing the refresh and update cycles.

5.4.2 Smooth Motion

Figure 5.5 is an annotated diagram of the eye following a moving image on the graphics screen. It details the case of one refresh per update seen in Figure 5.6.

The left half of these figures shows the view from a reference frame stationary with the screen; the right, from a

```
Scene as viewed by              Scene as viewed by
high-speed stationary           the eye tracking
camera.                         the image shown
                                at left.


 | <- Vertical line.                  | <- Vert. line.
 X <- X marks eye's                   X <- Eye's
 |  center of attention.              |  center of
 1 <- Refresh number.                 1  attention.


  Vertical line       |               |
moving to right -> X                  X <- Eye still
                      |               |  centered on
                      2               2  vertical
                                         line.


   Axes:                          Scene as
                                  perceived by
Left -- Space -- Right            the user:
     <------------>
           |
           |                      | <- Smoothly moving
           V                      X  line perceived as
                                  |  if it were still.
      Increasing                  1 <- Refreshes.
         time                     2
        (down)
```

Figure 5.5: Terminology for Eye Tracking Motion

frame fixed to the eye's center of attention (indicated by the "X"). The eye moves in an attempt to keep the line on the screen stationary in this frame. Lines are represented in these figures as disappearing immediately after they are drawn. This is effectively the case with a fast phosphor (e. g., P4). I have had no experience with slow phosphor. The known features in each figure are the average number of refreshes per update and the scene as perceived by the user. All other features are deduced from these. The motion of the eye as indicated is really the motion of the eye-brain combination that forms a person's center of attention. Actually, most researchers believe that multiple-image phenomena are retinal, but I know of no experiments on the matter.

Figure 5.6 shows one refresh per update producing motion so smooth that a slowly moving object is perceived as clearly as if it were standing still.

5.4.3 Multiple Images

The effect of having one update every two refreshes is shown in Figure 5.7. The crystal clarity of the one refresh per update motion is lacking here. The motion appears to be smooth, but the image is blurred. Close inspection of this "blurring" reveals that the image now has twice as many lines as before. What was once a single line sliding across the screen now appears to be a flickering doublet, a pair of closely spaced lines, having a separation directly proportional to the speed across the screen. Typically they are a few millimeters apart. These double images, arising from temporal aliasing, have been well studied on raster displays [Brown, 1977; Pearson, 1975; Szabo, 1978] and in animation [Williams, 1979].

The eye tracks smoothly rather than jerking and jumping across the screen. Understanding the problem does not really help. If one cannot reduce the refresh rate without flicker and one does not have enough computing power to speed the update rate, then one may as well enjoy the view.

5.4.4 Reversed Apparent Movement

With three refreshes per update things get worse (Figure 5.8). The diagram shows a bizarre rippling motion which is evident to the eye with an update rate of 10 Hz and a refresh rate of 30 Hz. The refresh order 1,2,3 is <u>right to</u>

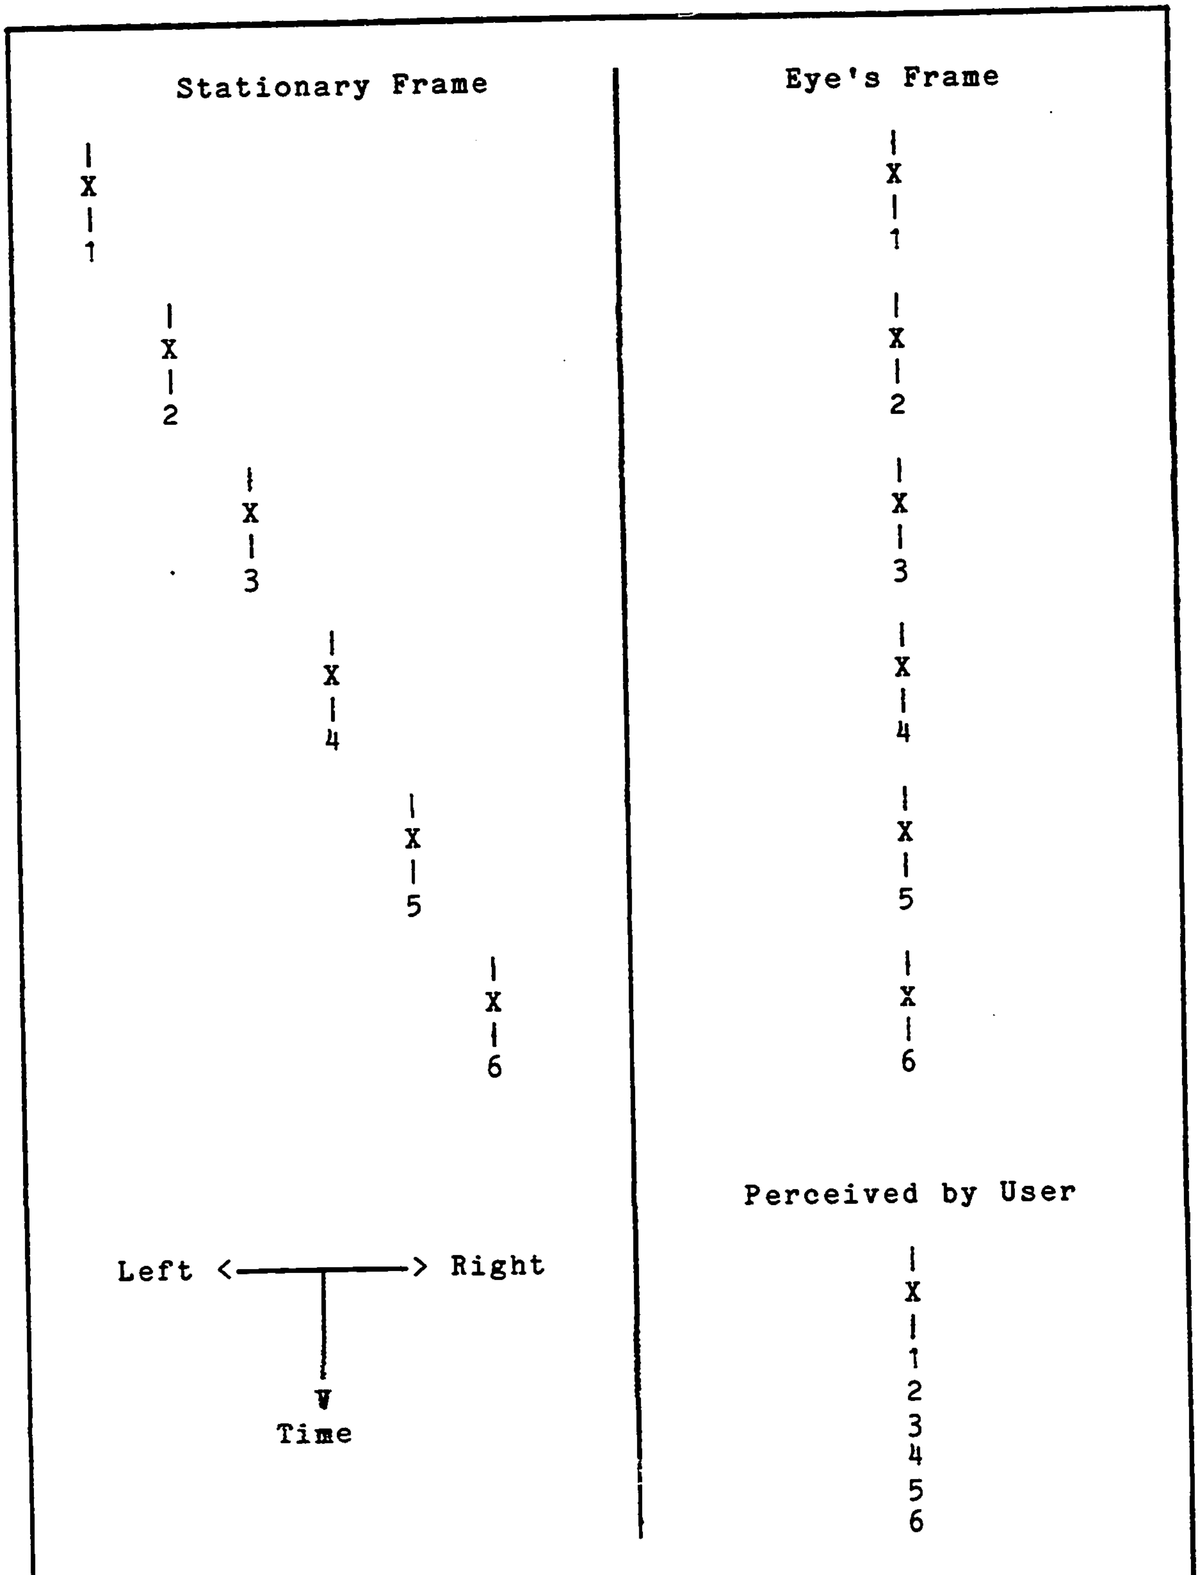

Figure 5.6: One Refresh per Update

Dynamic Motion -- -- Reversed and Erratic Apparent Movement left for the eye tracking left to right. Refreshes 4, 5, and 6 then repeat the effect. Even though the eye is tracking smoothly, as evidenced by the even separation of the lines, the user begins to be aware that the motion is no longer smooth. This rippling effect opposite in direction to the motion across the screen (i. e., reversed) and the proliferation of lines become increasingly annoying (and visible at lower speeds) as the update rate is lowered further.

5.4.5 Jerky Motion

By about ten refreshes per update (with a refresh rate of 30 Hz) the eye's motion can no longer be smooth. It turns and stops with each update. Figure 5.9 shows this jerky motion simplified to four refreshes for purposes of illustration. The spacing of the multiplet becomes uneven. That, plus the difficulty of looking at the rightmost lines of the multiplet as it moves to the right, indicate that the eye's motion is jerky.

5.4.6 Erratic Motion

Letting the refresh and update processes run unsynchronized, results in a varying number of refreshes per update. Figure 5.10 shows an example of erratic motion where the structure appears to move steadily except for occasional pauses. For this figure only, the right-hand columns are constant-speed frames of reference rather than the eye-brain frame of reference, to more faithfully represent how the motion is perceived in this case.

Some, or all, of the effects described above, when taken together, are called optical jitter by animators [Halas and Privett, 1958].

5.4.7 Unsuccessful Smoothing

The pauses that occur when an update misses its scheduled refresh can be eliminated by jumping a little farther the next update to make up for having been left behind once. But this really improves matters not at all (Figure 5.11). The motion is erratic in a different way, but erratic nonetheless.

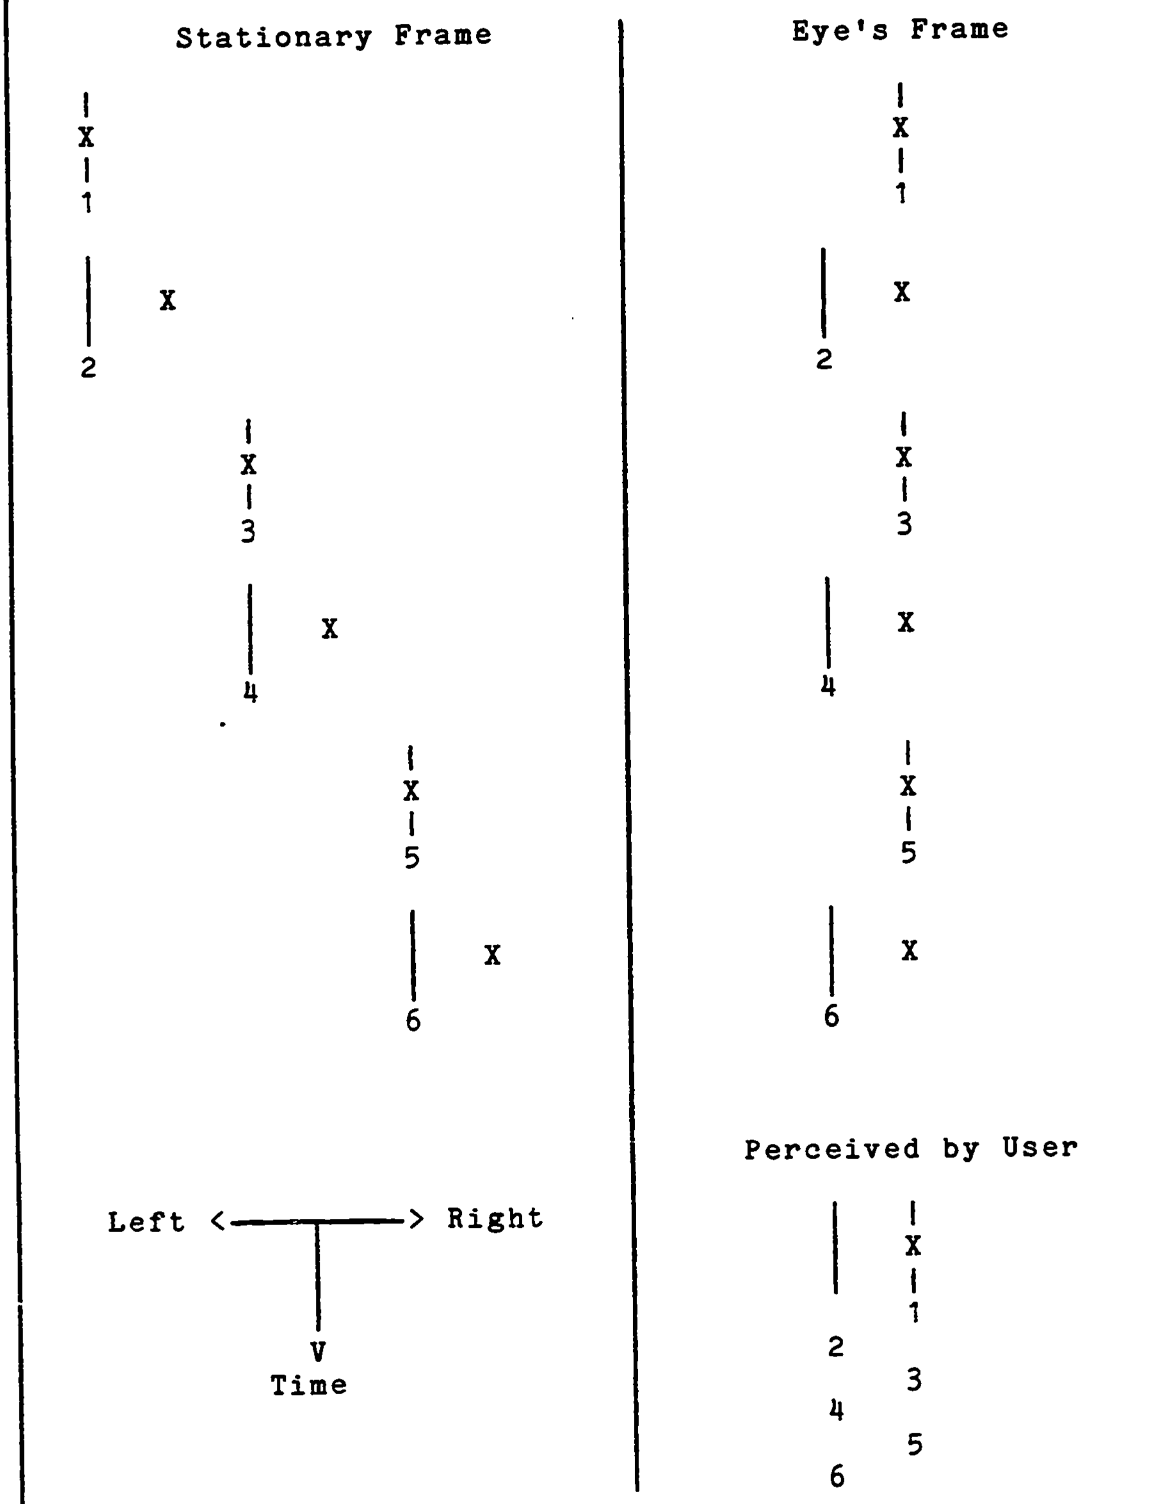

Figure 5.7: Two Refreshes per Update

The update process could, upon completion, wait for the refresh process to pick up the result of the update. This would be the "starting gun" for the next update cycle. Unfortunately, there is no guarantee that the end of the update cycle will be reached after a predictable number of refreshes (n). In a marginal case the number of refreshes per update may vary somewhat as the race between the update process and refresh n+1 is variously won and lost. The "starting gun" only assures that erratic motion will not occur very often. However, it can happen, and it can look bad. Small changes in the update time are quantized to multiples of the refresh time.

Deliberate variations in the amount of computing needed for each update can also cause erratic motion. One can attempt efficiency by shutting down dynamic processes not required for a given update. When an active device is moved, it can "wake up", causing its code to be executed then only. If the user does not ordinarily move all the joysticks simultaneously, then this is a great timesaver, but it does introduce variations into the update rate. Even when these variations are only a fraction of a refresh long, their sum is quantized to a multiple of the refresh time.

5.4.8 Smoothing Erratic Motion

A software low-pass filter can be built to smooth the refresh per update ratio and to assure reasonably smooth motion. The filter keeps the number of refreshes per update from varying by more than one from a "true" or "base" number. If the number of refreshes per update tries to drop by one, it must keep trying for several seconds before the filter will allow it. If the number of refreshes per update increases by one or changes by more than one, or if an interaction occurs, then the new number of refreshes per update is taken to be the new "base" value. This scheme limits spurious changes in the number of refreshes per update to at most a few times per minute.

5.4.9 Experience

The use of a software filter to reduce the spurious small jumps was worthwhile only for automatic viewpoint-rotation. It appeared unnecesary for manual viewpoint-rotation with a position device or for any other user-controlled dynamic motion that was not rapid (e. g., substructure manipulation). The unsteadiness of the hand seems to mask the spu-

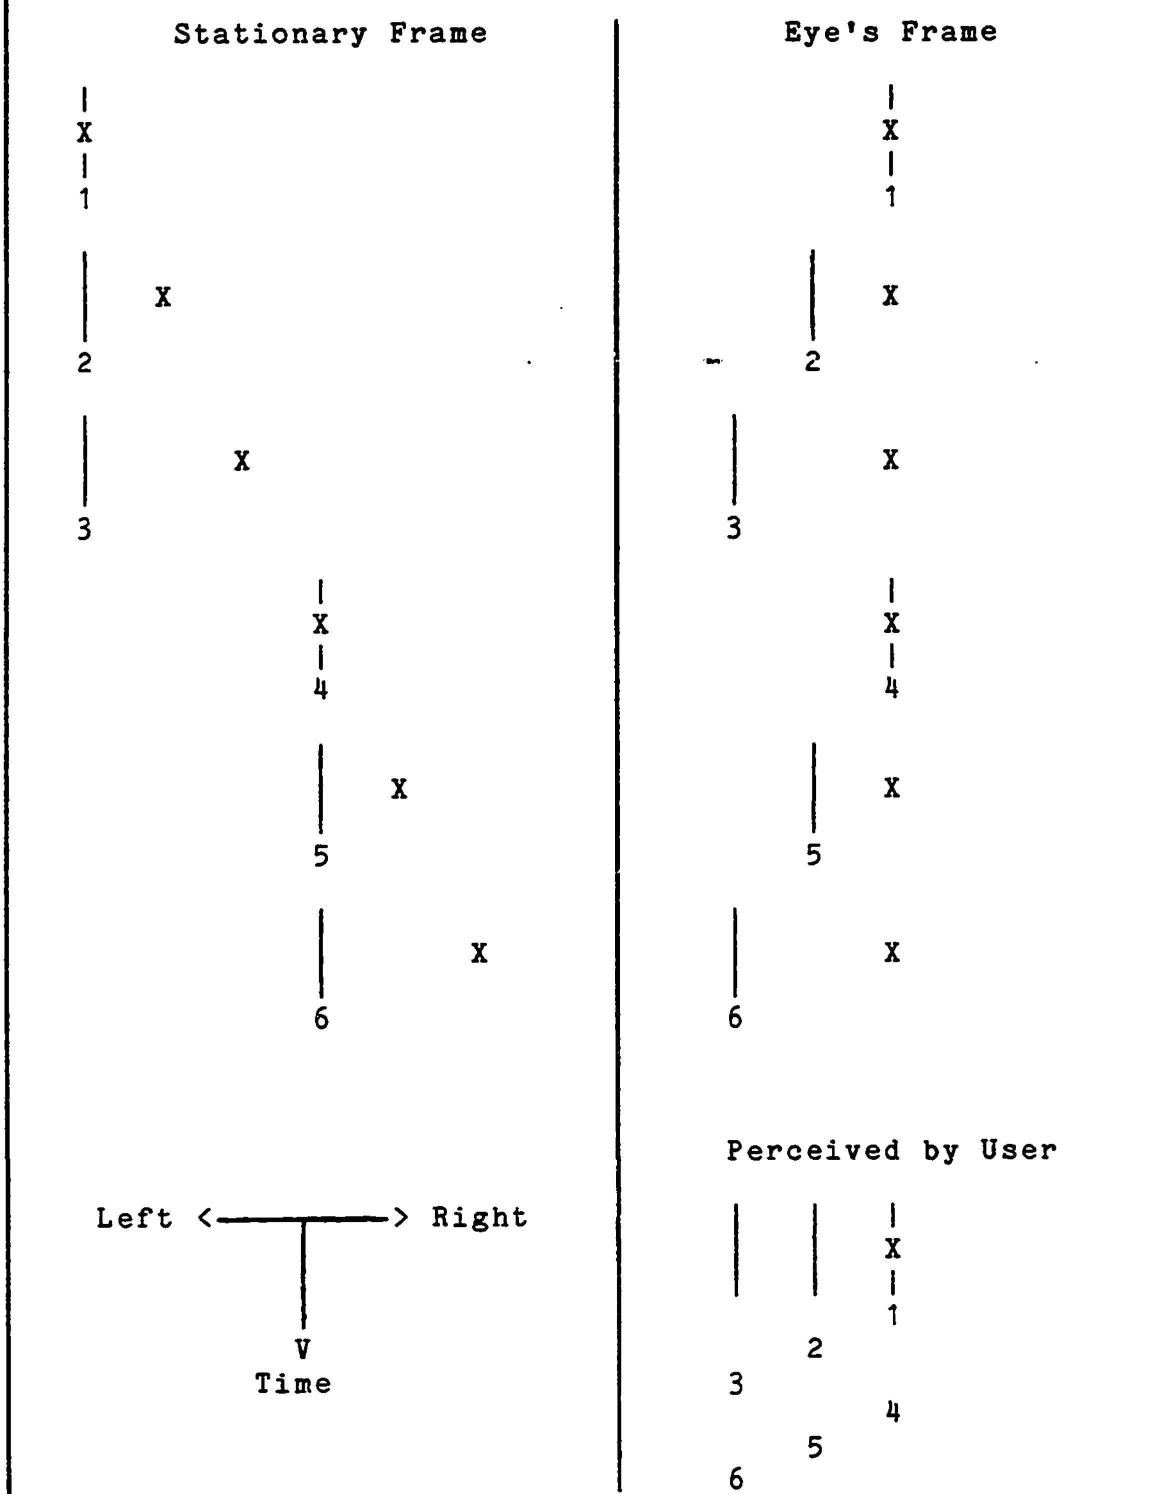

Figure 5.8: Three Refreshes per Update

rious jumps during manual viewing, and manipulated subimages are too small and moving too slowly for these jumps to be apparent.

With automatic, as opposed to manual, motion there is one more choice to be made. On the scale of seconds (rather than milliseconds as in Figures 5.10 and 5.11), motion can be a constant increment in angle with each pass through the update process, or it can be a constant increment in angle per unit of real time with the angle of increment becoming greater when the update rate is low. GRIP-75 uses per-pass increments for substructure rotation but constant per-second increments for automatic viewpoint-rotation. The benefit of the constant per-pass increment is that motions become slower and more deliberate as more dynamic devices are made active, if, indeed, this is a benefit. The benefit of the constant per-second increment is that it preserves orthogonality between the rotation rate and the number of active devices. Based on my experience with these two alternatives and on user complaints, I recommend using the constant per-second increment and preserving orthogonality between rate of motion and complexity of manipulation.

5.4.10 Conclusions

Reversed apparent movement can arise with three or more refreshes per update. Erratic motion can be reduced by means of a software low-pass filter, but, depending on the type of motion and the application requirements, one may not need to do so.

5.5 UNCONSTRAINED TRANSLATION

When translational motion has three degrees of freedom, it is said to be <u>unconstrained</u>.

A. M. Noll of Bell Laboratories has very kindly provided us with one of his 3-D positioning devices, a <u>Noll box</u> (Figure 5.12). A handgrip (knob on vertical shaft in figure) may be moved to any position within a cubic region of space.

We need 3-D translation for two things. The user may slide around a 3-D jack-shaped position indicator (Figure 5.3) whose position he can record interactively by pushing a button. The position may then be employed for many functions. For example, one might interactively center the viewing cube there, target an atom to head towards it during

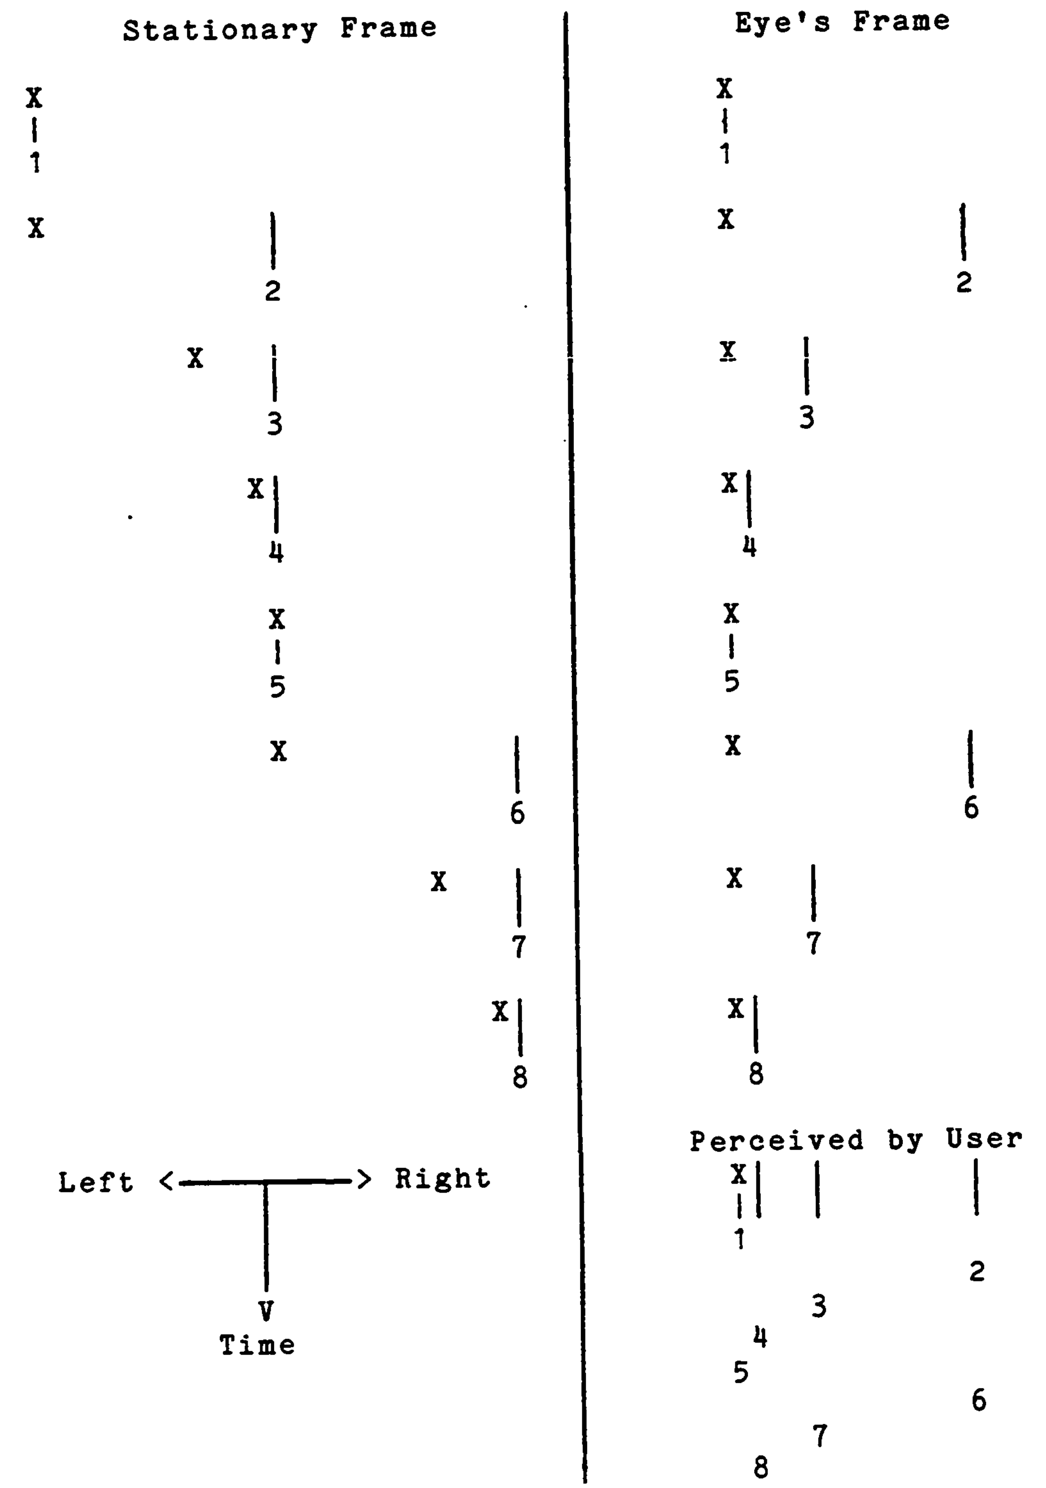

Figure 5.9: Many Refreshes per Update

automatic bond length and angle idealization, or write down the position on paper after decomposing it into its three x,y,z components. The second and more important use of the Noll box is to fly a substructure of a molecule through space and into its associated region of electron density.

The space in which the Noll box moves may have been rotated by the toothpick. Even so, left-right motions of the hand grip result in left-right motions on the screen. The same holds for up-down and in-out motions, etc. We say that this device operates in the _user's frame of reference_. Consistent with Kilpatrick [1976] (page 31 above), we observe that the scale of the motions on the screen can differ substantially from the scale of the device's motion without the user's noticing and without interfering with his productivity. It is important only that the directions of the motions be the same. We tried the alternative of the x,y,z Noll box motions resulting in x,y,z _viewing frame of reference_ motions, where the viewing frame is the one which rotates with the picture, and the result was disastrously unusable. This and other issues are discussed in Britton, Lipscomb, and Pique [1978].

Because the Noll box is a position device that acts in the user space, yet is nested within a rotation, it can get locked up (Figure 5.14). If an object is pegged on the right of the screen by the full travel of the Noll box (Figure 5.14 a,b), then a 180 degree rotation of the picture (Figure 5.14 c) will place the object where it can be moved only off the screen and out of view (Figure 5.14 d,e). Therefore M. Pique and I provide a _clutch_ button which temporarily disengages the physical device so that it can be reset at will.

We tried several schemes for holding the substructure at the edge of the screen as the Noll box moved as if to push it off, but these were inferior to a push-button clutch. These _automatic clutch_ schemes required the user to move his substructure away from where he wanted it over to the edge of the screen from time to time. This was unacceptable.

In practice the manual clutch is a distracting necessity. It improves over the automatic clutch in that the push-button clutch requires only that the user occasionally move his hand in the direction away from his goal, rather than also moving the substructure on the screen.

We would probably be better off with a spring-return velocity device, which would not need a clutch. The directness of a position device is more benefical to rotation, where it helps resolve the back-front ambiguity, than it is

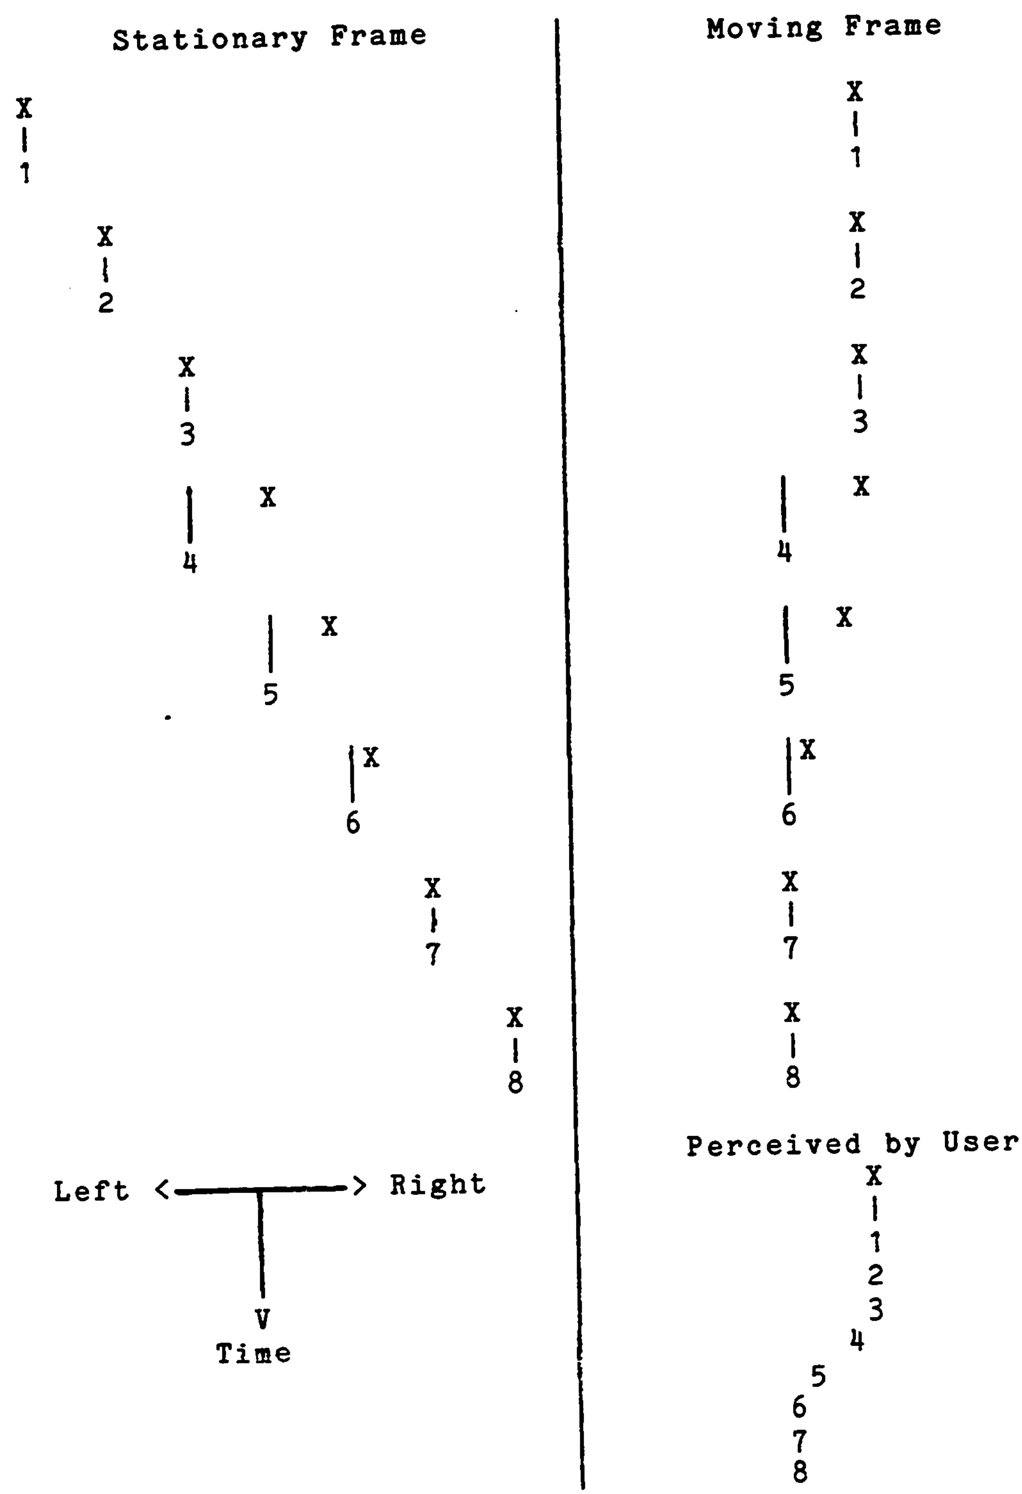

Figure 5.10: Uneven Number of Refreshes per Update

to translation. Also, a position device for translation is more difficult to grope for than one for rotation, because a rotation device is constrained to move in a circle and therefore tends not to wander far.

Users tend to decompose their 3-D motions into sucessive 1-D and 2-D motions consistent with Kilpatrick's [1976] results (page 32 above). They also tend to alternate between viewpoint-rotation with the toothpick and positioning motions. Even when not using stereo aids, they make motions into and out of the screen (after first moving in the plane of the screen) even though they cannot see their effect immediately. Actually they do see it in their mind's eye if not on the display screen. They depend on the effect of their motions being consistent along and into the plane of the screen. This is also consistent with Kilpatrick's [1976] observations.

5.6 UNCONSTRAINED ROTATION

5.6.1 Introduction

Rotation is said to be <u>unconstrained</u> when there are three degrees of rotational freedom.

We call our rotation controller[3] (Figure 5.13) the <u>fist</u>, because it is grasped by the user's nearly closed hand. It has three degrees of rotational freedom, each with a spring-return. When each of the push, twist, or lean axes of the device is actuated alone, as is usually the case, rotation about the corresponding x, y, or z axis on the screen (Figure 5.2) occurs. Twisting it simultaneously about more than one orthogonal axis sets an appropriate intermediate axis of rotation, as if the device were isometric. It is not, because it moves slightly when pushed, but most users do not notice the discrepancy.

The direction of action on the screen is made to be approximately one-to-one with the device in user-space, as with the toothpick. But, whereas the toothpick is a position device, the fist is a velocity device.

----

[3] Model 521/3 joystick.

Measurement Systems, Inc.
121 Water Street
Norwalk, Connecticut 06854

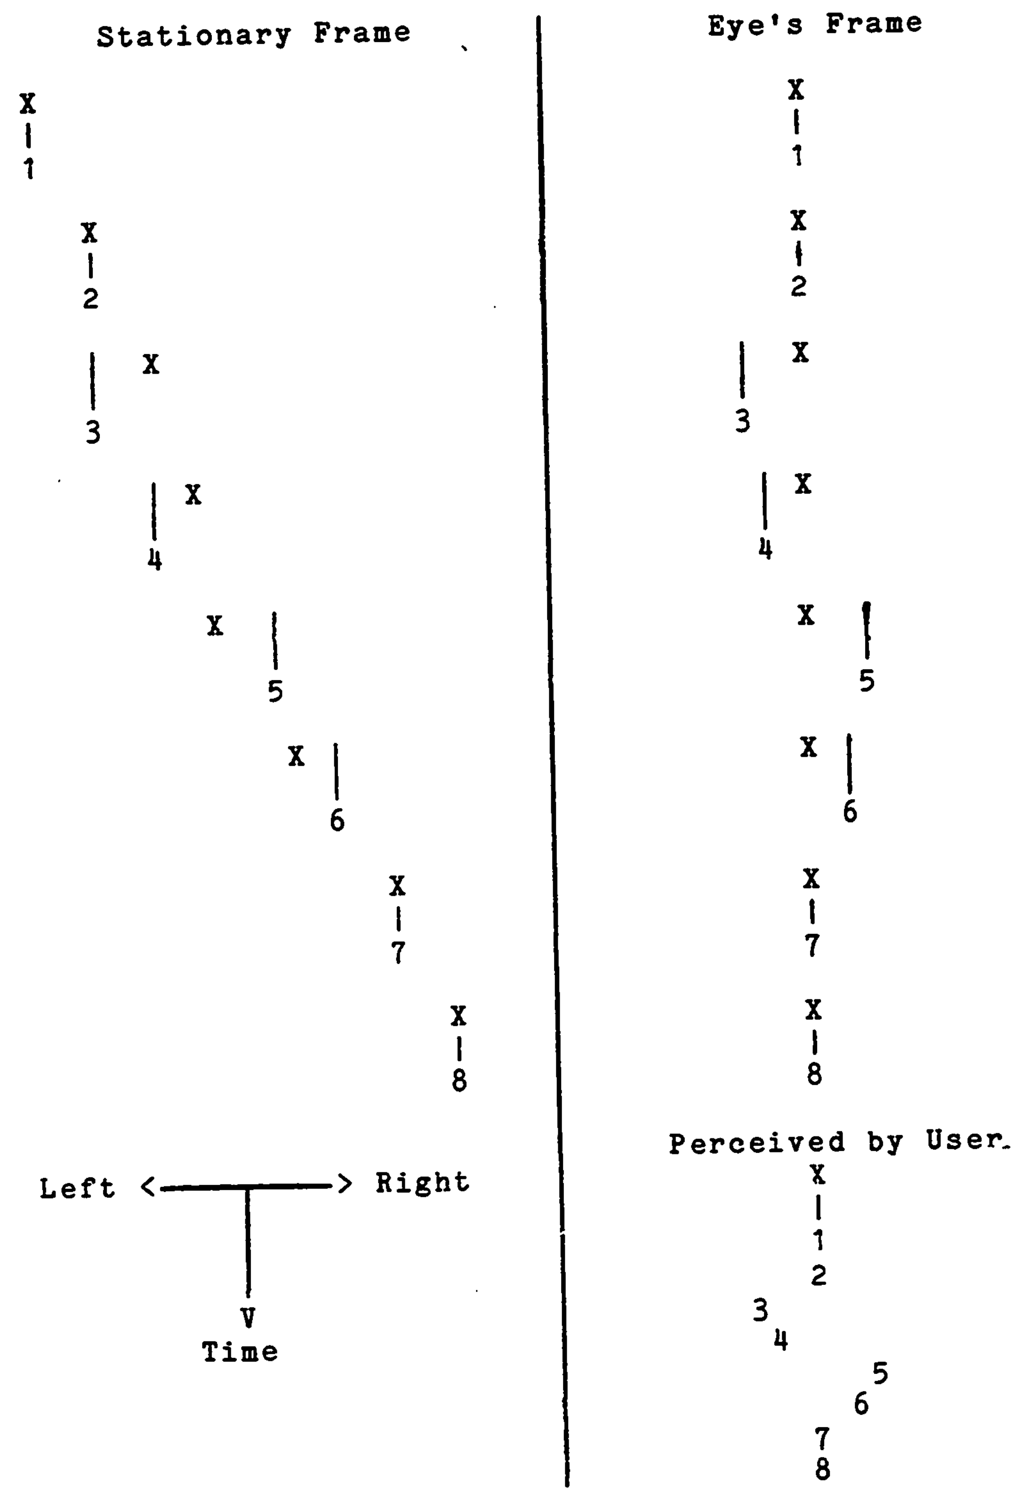

Figure 5.11: Making up for Lost Time

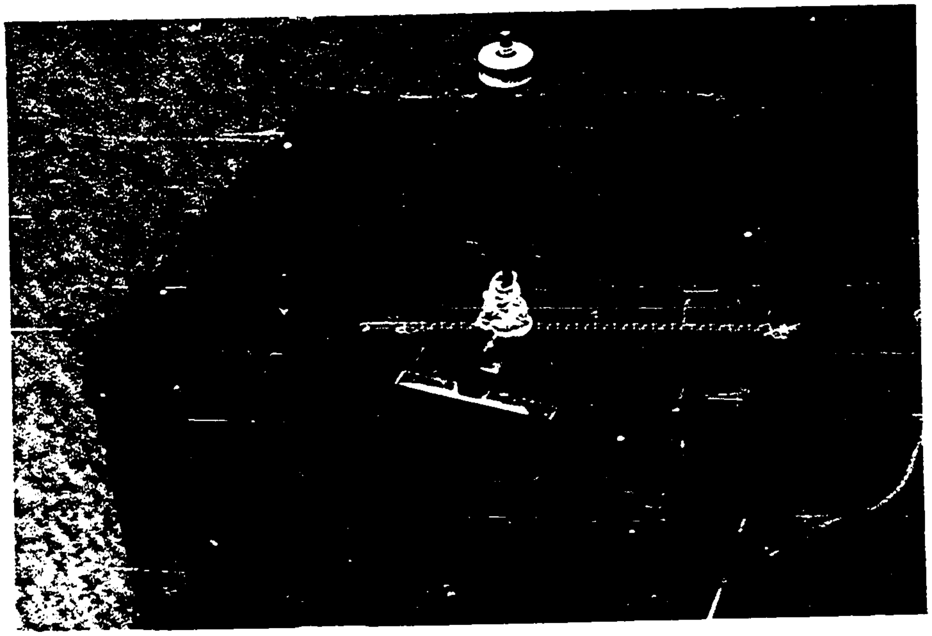

Figure 5.12: Noll Box for Substructure Translation

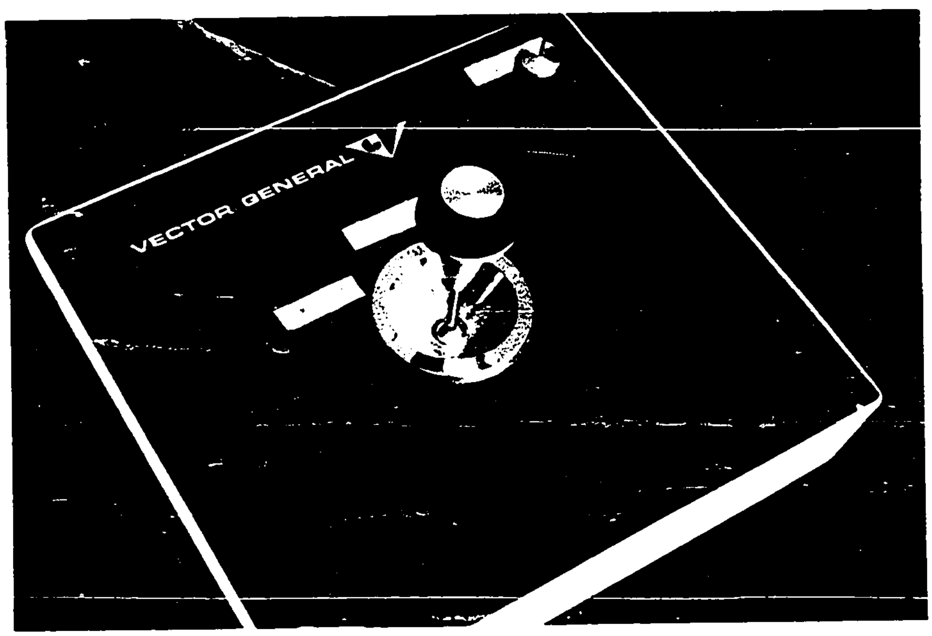

Figure 5.13: Fist for Substructure Rotation

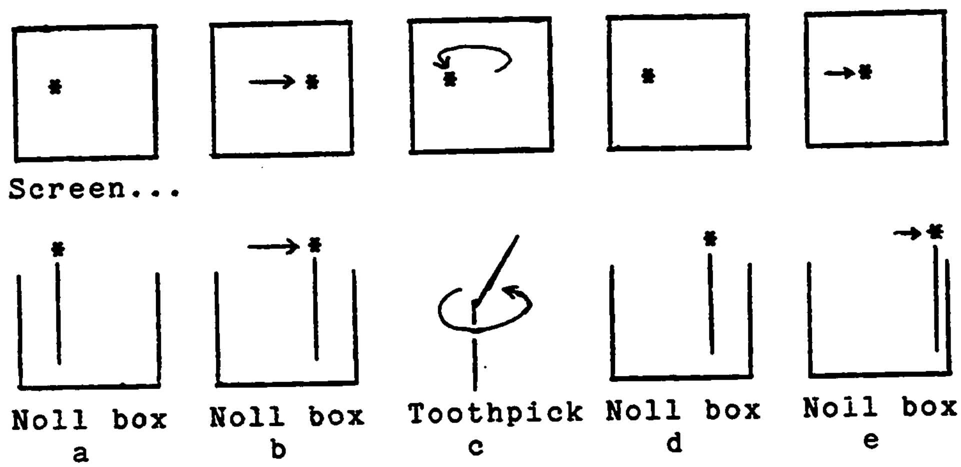

Figure 5.14: Noll Box Trapped

A velocity device is needed for substructure rotation, because a position device, other than a three-axis tracking-ball, would need to follow the substructure during viewpoint-rotation by being motorized or mounted on the viewpoint-control. The same question comes up for unconstrained translation. The choice is between a velocity device and a position device with a clutch. A velocity device eliminates the need for a clutch in both cases.

While a velocity device is suitable for substructure manipulation, it is not good for viewpoint-rotation. It lacks the directness of a position device.

The square of the device's displacement (d) sets the velocity with the sign preserved ($d * |d|$). This is standard practice for velocity devices in graphic systems. With velocity devices the user wishes to control equal steps of percentage change of the rate.

The hand controller must allow rapid rotation for gross orienting while preserving enough resolution at low rotation rates for delicate touch-up work. A maximum rotation rate of 40 degrees per second at maximum deflection appears to yield a satisfactory compromise.

When active (during substructure fitting) the unconstrained rotation program can "go to sleep", that is, it is only called when the user's hand is pushing on the device. This frees the processor to update other manual motions more rapidly.

5.6.2 Rotation Nesting

The molecule fragment undergoing unconstrained translation and rotation must rotate in concert with the rest of the structure under dynamic viewpoint-rotation.

The mathematics of this is covered in Britton, Lipscomb, and Pique [1978], but for comparison I shall cover it here again with two important differences. I shall use pre-multiplication by the rotation matrix as our display unit demands ($\underline{v}' = R\ \underline{v}$). This will demonstrate that Pique's [draft] notation can handle pre-multiplication as well as post-multiplication. Also, I shall employ the method that Pique and I installed at the time rather than the newer formulation described in Britton, Lipscomb, and Pique [1978]. The reader may be entertained by comparing the two treatments of the same subject (See Appendix A).

Using Pique's [draft] notation, rotation matrix uRv corresponds to viewpoint-rotation by the toothpick control. The subscripts "u" and "v" show that matrix "R" maps the <u>viewing space</u> (v) onto the <u>user space</u> (u). That is, the display list is mapped onto the graphic screen in the displayed orientation. The matrix uRm maps the <u>manipulated substructure</u> (m) onto the screen. Dynamic viewing changes uRv. Matrix uRm must be made to follow along because it shows a substructure nested within the viewing space. This can be achieved by holding the relationship between these two rotation matrices constant. That relationship is itself a rotation, vRm. The relationship is

uRm = uRv vRm.

Note the notational rule that the letters in the middle (v) cancel out and one "R" goes away, so:

```
uRv   vRm
|       |
 -->uRm<-
```

Thus, changes in uRv can be easily reflected in changes in uRm, making both the viewed image and its substructure appear to be rigidly connected while dynamic viewpoint-rotation is in progress. When the substructure is rotated with respect to the rest of the picture, the relationship between them changes. This matrix (uRm) must be updated accordingly.

The nesting algorithm follows below. The symbol " := " indicates assignment. This is a loop executed about 10 times per second under normal conditions. Each pass through the loop makes a single change or update to all the elements in the picture.

1. vRm := I.

   "I" is the identity matrix. At the start of substructure fitting there is no difference between its orientation and that of the surrounding picture. That is, uRm will be set to uRv.

2. uRv := dynamic viewpoint controller.

   The toothpick is interrogated, its setting recorded, and a rotation matrix corresponding to that orientation is generated.

3. uRm := uRv vRm.

   The manipulated substructure is made to follow along with the viewpoint-rotation.

4. uRm := change(uRm).

   The substructure's orientation is changed slightly with respect to the rest of the structure. The algorithm for this unconstrained rotation will be covered in the next section.

5. vRm := vRu uRm.

   The new relationship between the substructure and the rest of the picture is recorded. vRu is the inverse, or simply the transpose, of uRv. This is a special property of all rotation matrices that the inverse of the matrix equals its transpose.

6. Display uRv and uRm.

   Now that the mathematical transformations for this update cycle are complete, update the picture on the screen. Continue refreshing the picture as often as necessary until this point is reached again. This is typically three refreshes of 1/30th second each.

7. Go to 2.

There are some obvious efficiencies to be had here. If the toothpick control has not moved since the last pass through the loop, then uRv is not recalculated. Also, if the substructure rotation control is zeroed, then "change(uRm)" is the identity matrix and need not be calculated.

5.6.3 Rotation Algorithm

The unconstrained rotation algorithm, "change(uRm)" in step 4 above, follows techniques laid out in Tountas and Katz [1971] (page 28 above) with some improvements. In GRIP-75 it is subroutine SCREEN. M. Pique and E. Britton assisted in interfacing this routine to the GRIP-75 system.

Assume that the substructure is in an orientation described by a rotation matrix. During one update cycle (about 1/10th second) it is changed by a matrix multiplication to represent a new orientation which will be the initial orientation of the next update cycle.

Since the output of each iteration is the input to the next, the iterations accumulate error. A position control would gradually lose register with the image it was moving. However, since the substructure is under velocity control rather than position control, the user senses the motion incrementally and therefore cannot detect error buildup.

The matrix describing the substructure's orientation must, however, retain those special qualities that make it a rotation matrix. The buildup of iterative error mentioned above must be distributed carefully within the matrix. If the rows are kept mutually orthogonal and the columns are kept mutually orthogonal, but no constraint is placed on the length of these vectors, then the picture will change size, generally by shrinking slowly out of sight. If the rows and columns are held to unity in length, but no constraint is placed on orthogonality, then the picture will stay the same size but it will change shape until overflow occurs, and the image "flaps like a wounded bird" [Wipke and Dyott, 1973]. If all the orthonormalization conditions hold, but the determinant has the value of minus one, then the picture will be displayed inside-out. The rotation matrix must retain its characteristic qualities at least well enough for the user not to see any change in the size or shape of the object he is manipulating. Given this, the execution of the rotation routine should be fast and uniform with each invocation to assure smooth motion.

The procedure, "change(uRm)", executed once per update, has three parts. " := " signifies assignment.

1. Fu := X Y Z.

Fu is the incremental (one subscript) user space (u) rotation matrix to be applied this update, from the setting of the fist.

X, Y, and Z are rotations about the push, twist, and lean axes (Figure 5.2) of the fist by angles x, y, and z. Their ordering is unimportant since the incremental rotation angle is small and rotations are approximately commutative over small angles. Ordinarily, if three rotation matrices are multiplied together, only one of them can operate in a given space. However, if for some reason the order of these rotations is of no particular significance, then any one of them can be thought of as applying to a single space, here, the user space.

The approximations "x = sin x" and "1 = cos x" are satisfactory for small angles.

M. Pique suggests that when analytically combining X, Y, and Z, the square and cube terms can be dropped as an approximation comparable to the one above. This suggestion came too late for me to try, but it seems sound. I present it here for the reader's approval:

$$Fu := \begin{bmatrix} 1 & -z & y \\ z & 1 & -x \\ -y & x & 1 \end{bmatrix}$$

2. uRm := Fu uRm.

   Rotate the substructure from its old to its new orientation. This will be the "old" orientation in the succeeding update.

   The incremental rotation (Fu) is applied in the user's frame of reference by multiplication on the user-space side of matrix uRm.

3. Clean up uRm. Make it look a little more like a rotation matrix.

   Repeated Gram-Schmidt orthogonalization and normalization of the columns or rows is adequate to keep the matrix in good condition. The orthogonalization may be speeded by calculating only one pair of columns (or rows) per invocation of "change(uRm)".

Tountas and Katz [1971] (page 28 above) suggest that

n := 1/2 + (v*v)/2

is a sufficiently good approximation to the square root of v*v for the normalization of vector v in the expression

v := v/n.

The square root approximation relies on the magnitude of the rows or columns (v) being near unity (differing from unity by "e"). The first two terms of a binomial series are used to generate the expression for "n" having an absolute error on the order of the square of "e".

With all of these approximations in force (except for eliminating the square and cube terms), the matrix will avoid overflow for rates up to 15 degrees per increment in 16-bit scaled arithmetic. Beyond that, the substructure turns inside-out as overflow occurs and the determinant flips and stabilizes to plus or minus one. At ten increments per second this is 150 degrees per second, far faster than we need in production work. The maximum rate, at maximum deflection of the physical device, is set four times lower to facilitate delicate work as discussed on page 92, and rigid-body motion prevails.

Chapter 6

STRUCTURAL CUES

6.1 INTRODUCTION

We use three structural cues: relative fading, intensity depth-cueing, and z-clipping. A fourth technique, color, was available using a spinning disk but was not implemented.

A related group of items may be uniformly dimmed to enhance the visibility of the remaining items on the screen. This is called relative fading (Figure 6.1). Similarly, the items apparently close to the viewer may be emphasized by dimming out those farther away. This is called intensity depth-cueing (Figure 6.2). Finally, a "plane of invisibility" may be swept through the viewbox, a space, fixed in the user's frame of reference, that encloses the entire picture. This plane is parallel to the screen, and items between it and the user are blanked out. This is called z-clipping (Figure 6.3).

All three techniques dim back a portion of the picture, making it more understandable by simplifying it. This probably accounts for their popularity. Fading is used by all our clients, intensity depth-cueing by about two thirds, and z-clipping scarcely by any (because it effectively disables the other two).

The challenge to this system builder was to make the user think that these functions are mutually orthogonal, although obviously they are not.

The user is given sliders to control these functions. They are treated as position devices. That is, the effect of each function holds steady at each setting of its respective device. Velocity control of these functions is inappropriate, as M. Pique and I discovered in an early implementation. It was difficult indeed to tell what one was doing. As soon as possible, we made these functions position-control sliders, which are much easier to use. I then added some function to this package and a certain amount of human engineering.

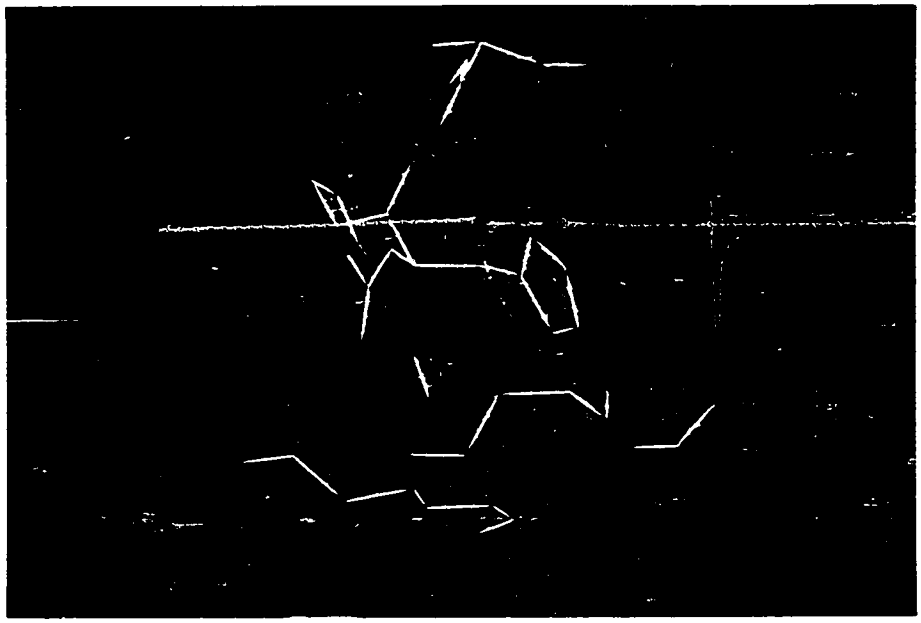

Figure 6.1: Molecule Brighter than Map

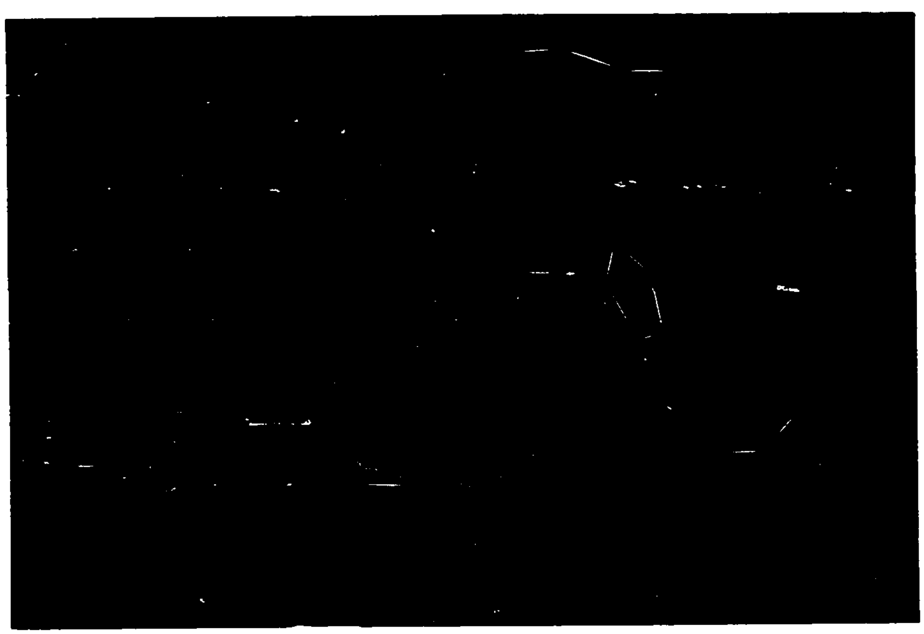

Figure 6.2: Intensity Depth-cueing

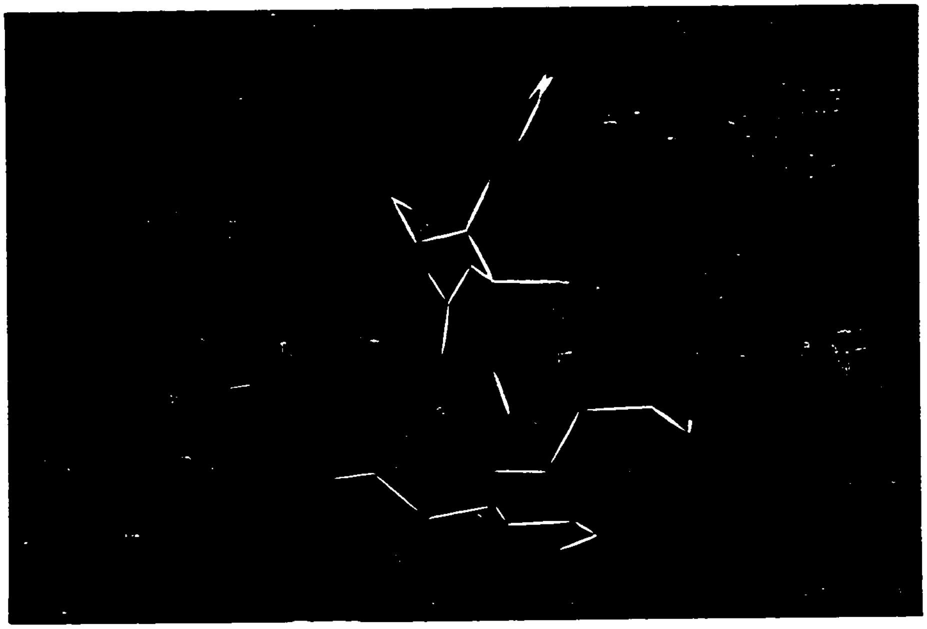

Figure 6.3: Z-clipping

To the builder, intensity control of two objects implies four degrees of freedom, the slope and intercept of the intensity drop-off with depth for each object.

One of our clients did not fully grasp the implications of this when he pulled me aside one evening to say that his difficult work required my help by giving him, the user, full access to all possible intensity states allowed by the hardware and that human engineering considerations must be strictly subordinate to this.

He was arguing for generality, rather completely at the expense of propriety and orthogonality. His problem is common among users, who often fear their creative range will be cramped by the narrow vision of the builder of overly specific tools.

Four sliders with four degrees of freedom would have been too unwieldy, so I supplied only two sliders, one for relative fading and one for intensity depth-cueing. These map in a complex way onto the four degrees of freedom that actually exist. Z-clipping is not involved in this controversy. The very fact that these ideas can be so simply stated implicates them as being "understandable". These concepts are display-independent and suitable for user controls. With the two degrees of freedom left over one can choose among several methods for implementing these features. I shall now offer two: one that I implemented, and one I call "best" that probably would have been superior had I time to attempt its implementation.

This folding of four degrees of freedom into two has been very successful. It took our clients two years to discover they were shortchanged.

6.2 RELATIVE FADING

Fading the density map is such an important cue for distinguishing it from the embedded molecule that molecular graphic systems whose displays cannot change their beam intensities implement it nonetheless by the simple expedient of drawing the molecule several times for each redraw of the map or by slowing the beam for a brighter line.

The picture is easier to understand because it is simpler. The molecule has fewer lines per unit volume than the map. By dimming the map its imposing detail becomes less obtrusive, drawing attention to the otherwise elusive lines of the molecule. This balances the effective visibility of the two. In systems where this must be done by redrawing or by slowing the beam, there is the happy coincidence (really a consequence of the need for balancing) that the extra time for the molecule is small in comparison to the map display, and so there is little increase in the flicker.

The labels on the user-controlled slider cast this feature as "relative brightness", the opposite sense of the implementation. This labels the extreme positions of the control by what the user does see rather than what he does not. As to why, imagine a living room light switch labeled "darkness: on-off" that has to be flipped to "off" when the sun goes down.

When both the molecule and the map are visible and a light pen hit on either is meaningful, the light pen should preferentially "see" the one that the user more often wants to hit. It must be programmed to identify either the first or the last line drawn in its field of view during a refresh. This assumes that the entire visible map is drawn before the molecule or vice versa. In our system the user is far more likely to be selecting an atom than editing a contour line, hence the system should preferentially see the molecule. This is easy to get right, but I never got around to it, so instead the system sees the map rather than the molecule, if given a choice. Consequently, the user must fade back the map before selecting an atom and then fade it up again. This irritating procedure must be performed several times each hour.

Some users solve this problem by dimming down the map sufficiently for the light pen not to see it. This makes it too dim for intensity depth-cueing to distinguish front from back. In any case the light pen is used in occasional flurries of heavy activity interspersed by five to twenty minute streches of dynamic work. It is not a constant irritation.

Molecule comparison is a different use of the relative fading control. When no density map is being displayed but several molecules are on the screen, the first of these in the molecule coordinate list is faded with respect to the others as if they collectively were the density map. Normally, the chemists display two versions of the same molecule superimposed in space and having nearly identical conformations. They usually slap the control back and forth between its two extremes. This makes the structure jump back and forth between its two conformations with a satisfying illusion of continuous motion. A toggle switch would do just as well as the slider, except that intermediate positions of the slider are occasionally employed for critical scrutiny of small portions of the visible picture. Once relative fading of the density map relative to the model was implemented, molecule comparison came practically free.

Figure 6.4 shows the intensity versus viewbox-depth space. Figure 6.5 shows a diagram of the portion of this space used by GRIP-75. The intensity levels described in this figure will be used in subsequent figures.

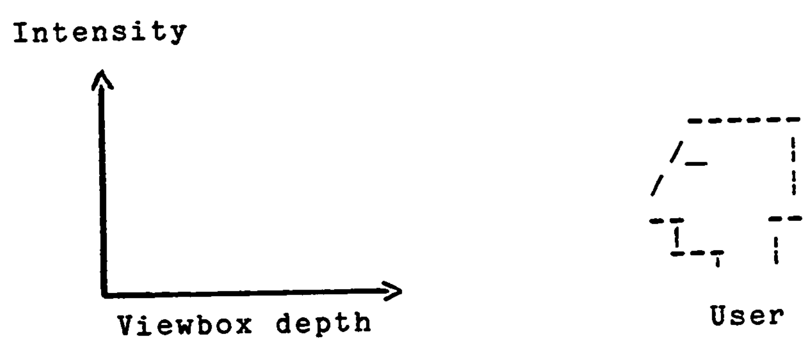

Figure 6.4: Intensity versus Viewbox-depth Space

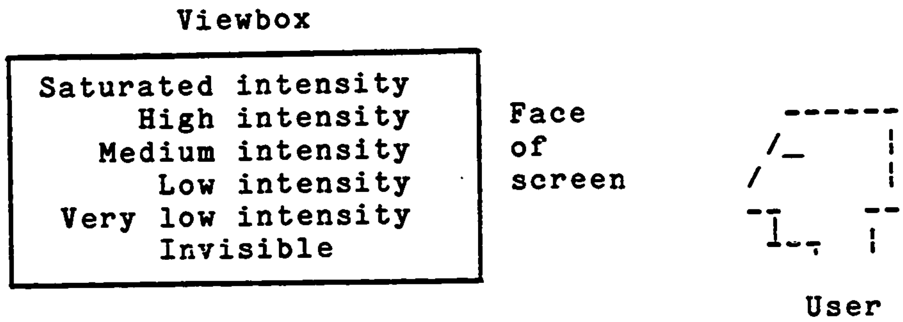

Figure 6.5: Intensity Viewbox

Fading is shown in Figure 6.6. "I M A G E" indicates the back to front extent of the visible picture. The slider is scaled in software to make the "dead zones" at the extremes of intensity as small as possible.

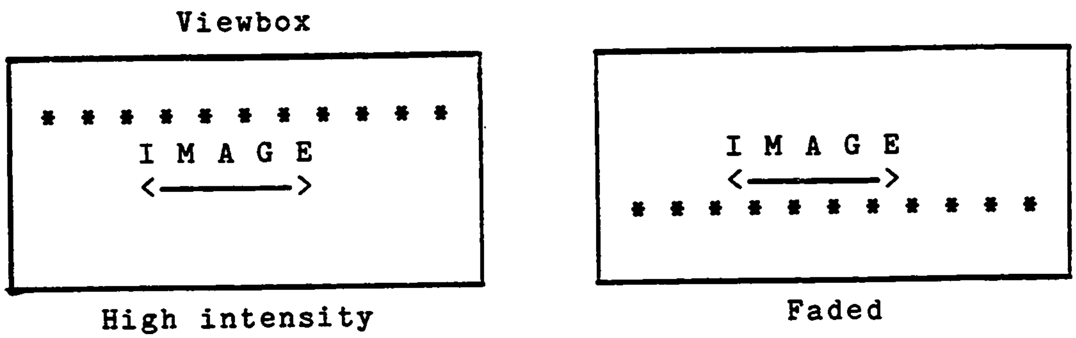

Figure 6.6: Intensity Fading

Relative fading is shown in Figures 6.1, 6.7, and 6.8. Note that the one feature is brought up to full intensity before the other is faded back. I tried simply trading off the intensities linearly over the full range of the slider's travel, but at intermediate positions of the control the picture was unpleasantly dim.

The z-clipping plane has a surprising use in implementing relative fading. Just before the relative fading slider pegs at one limit of its travel, I have the z-clipping plane blank out the dimmed feature (either the molecule or the

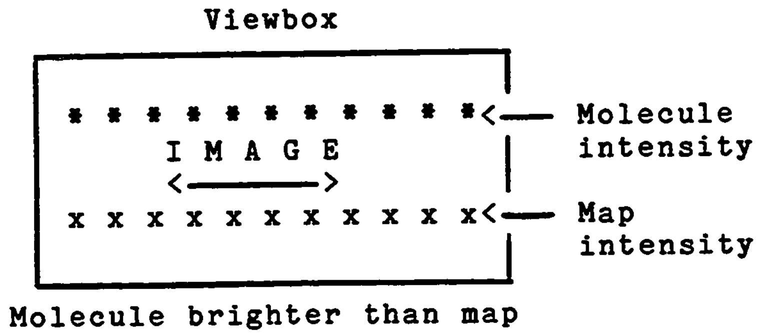

Figure 6.7: Scheme for Bright Molecule

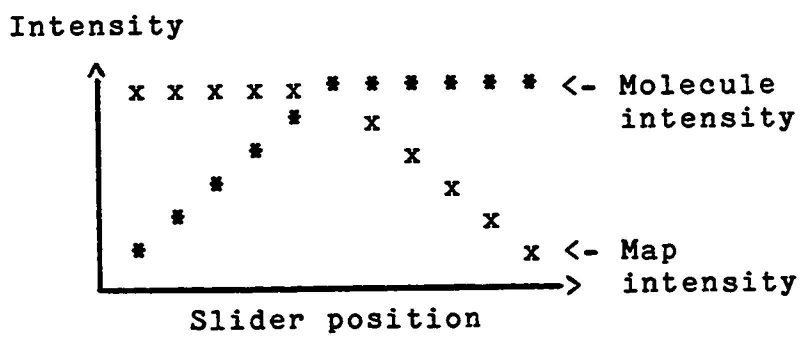

Figure 6.8: Relative Intensity Control

map) completely. When the display unit is properly adjusted, the feature being dimmed is already invisible. Otherwise, the dimmed feature blinks out. Thus, the device performs as expected (or as close as possible to the expected) when the hardware is uncooperative. This scheme is possible because the molecule and the map, although superimposed in x,y,z, are in separate intensity spaces, each with its own z-clipping plane.

As a beneficial side effect, the user gets what he sees when he commands a hardcopy plot of the picture on the screen, because the plotter software written by G. Hamlin and D. Tolle does take the z-clipping plane into account. The most common instance of this occurs when a user is doing molecule comparison. He commands a plot of one molecule, moves the relative fading slider to the other extreme, and commands a plot of the second conformation. He is not surprised to receive the plots just as he previewed them on the screen.

6.3 INTENSITY DEPTH-CUEING

Like fog or haze, intensity depth-cueing (Figure 6.2) reduces the contrast with depth between a visible object and its background. It is the poor man's hidden-line removal, and it simplifies perception of the picture by emphasizing the front at the expense of the back. Items in the center of the viewing region are obscured only by objects in front of them. Items in the back may be swung around up front, where they can be seen better; this at the expense of having to take the time to swing them forward.

Further, the kinesthesia of the dynamic viewpoint-rotation control (the toothpick) is enhanced by intensity depth-cueing. This is because for small angles near the home viewpoint orientation the motion of the bright parts of the image (forward) follow the nearly linear motions of the user's hand on the control stick (for small angles).

Together with relative fading, we have the same encoder, intensity, for two things, distinction and depth. Britton [1977] rightly calls this "a poor choice but the only one open to us". They do interfere with each other, and relative fading is too useful to be optional. Intensity depth-cueing can be shut off without doing much harm, because it does not help to resolve back-front ambiguity as well as one might imagine, and because it interferes more with the light pen by making objects in the distance hard to pick. Perhaps this explains why many of our clients do not use it. But intensity depth-cueing is popular enough to be worth building. The challenge is to make it appear independent of relative fading.

M. Pique implemented a simple intensity depth-cueing feature early in the system's development. The users were provided two sliders whose settings were written directly into the two hardware registers controlling the intensity slope and offset of the entire picture. Users did not like this scheme, but it held them at bay while I built an improved package based on observed use of Pique's. Given a free choice of all possible intensity slopes and offsets, they performed intensity depth-cueing as shown in Figure 6.9.

The intensity drop with depth is held constant while the points of saturation and invisibility are swept through the

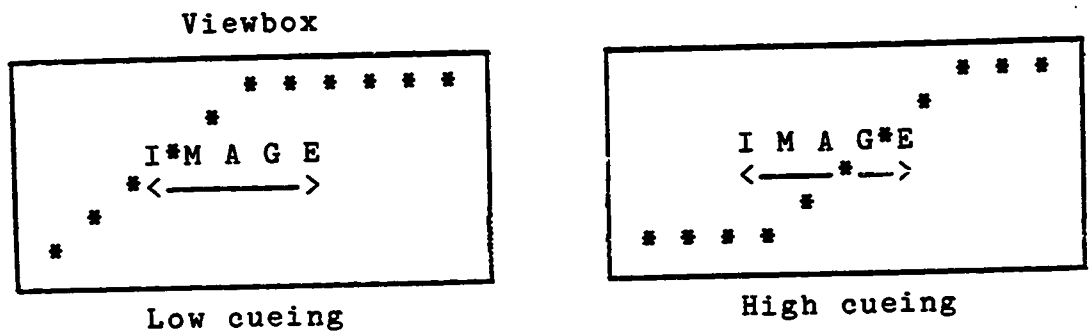

Figure 6.9: Scheme for Intensity Depth-cueing

picture. The slope is the maximum the hardware provides. By selecting it over all lesser slopes, the users clearly are asking for more. It is able to dim the picture from maximum intensity to only near invisibility for the size picture we have (as shown).

The display unit implements the intensity depth-cueing rate, or slope, as an exponential function. How the eye sees it is explained in Eastman Kodak [1950].

> However, the eye tends to see as equal tone steps not the difference in reflectance (e. g., 10, 20, 30, and 40 percent, where there is a constant difference of 10 percent), but rather equal ratios of reflectance (e. g., 10, 20, 40, and 80 percent, where the ratio of each reflectance to the preceeding one is 2). As a result, the gray which impresses the eye as falling midway between white and black actually has a reflectance of about 20 percent.

This property holds for luminance as well as reflectance, and it is called the <u>Stevens power law</u> [Stevens, 1957]. A discussion of this law and a comparison to similar laws (Weber's law and Fechner's law) may be found in Kaufman [1974].

Thus, I show a straight line for the exponential drop of intensity with depth. The hardware implements the action in Figure 6.9 as a motion of the picture in z through a stationary intensity field. To the user this appears to be quite the other way around, and that is how I show it.

As the intensity depth-cueing is first increased, only the back of the image is affected, while the front remains at a uniformly high intensity. When extreme, only the front of the picture shows.

As the popular choice, this method was selected for the advanced package. It appeared more that users were trying to simplify the picture by dimming the back of it, than they were interested in seeing all of the picture with different but uniform intensity drop off.

Fading, given intensity depth-cueing, I implemented as more intensity depth-cueing (Figures 6.10 and 6.11). This is not "correct", of course, but it has bothered only one user who could not get "a little bit" of intensity depth-cueing. The amount of intensity depth-cueing applied is the product of the settings of the fading and the intensity depth-cueing sliders. The product rather than the sum is taken to avoid dead zones on the controls. That is, one control will almost always have an effect independent of the setting of the other.

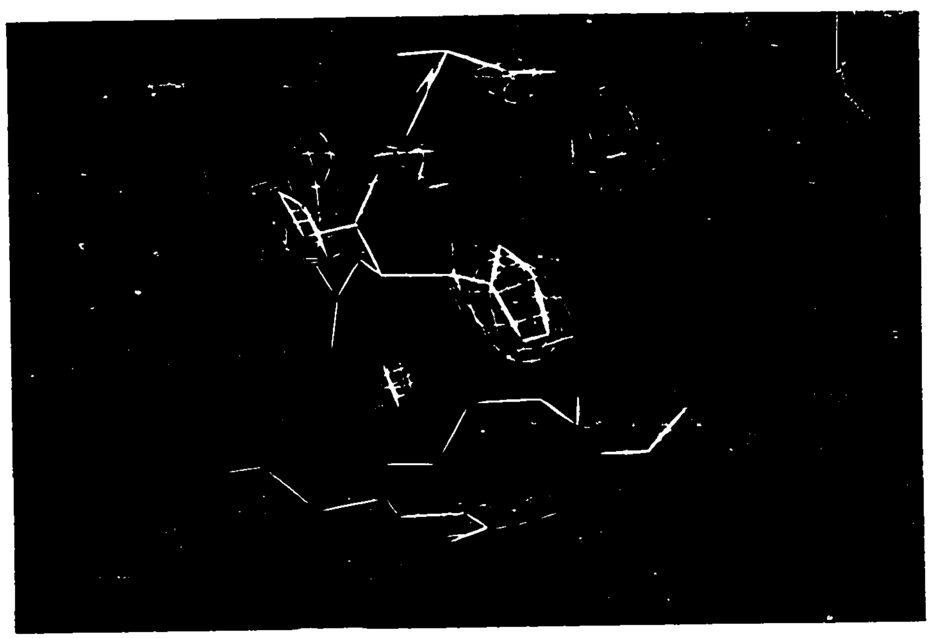

Figure 6.10: Relative Fading with Intensity Depth-cueing

Another issue is size compensation. The goal is to make picture size orthogonal to the apparent intensity depth-cueing. When the picture is shrunk for side-by-side stereo or

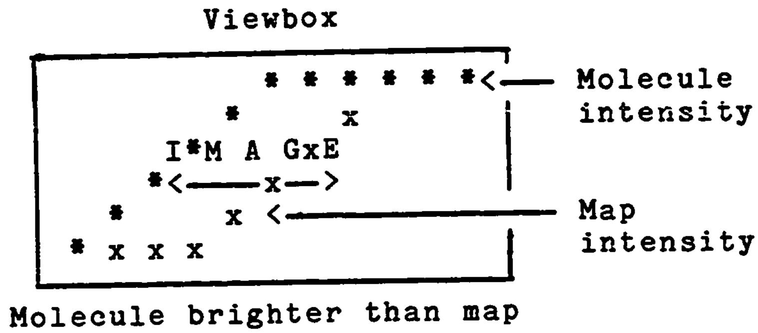

Figure 6.11: Scheme for Fading with Depth-cueing

for polarized-mirror stereo, the amount of intensity depth-cueing, already at its maximum, remains constant. With no correction (Figure 6.12) the amount of intensity depth-cueing across the image is effectively decreased. It could be corrected (Figure 6.13), except that the intensity depth-cueing slope cannot be increased on this display unit. Something had to be done though, because one effect of no size compensation (Figure 6.12) was for the smaller picture to appear unpleasantly dim, even though its total brightness was actually unchanged. I guessed that the apparent "brightness" of the image was driven in part by the true intensity of its near face. I installed a compensator that keeps it constant (Figure 6.14), and it seems to do the job, for our display at least.

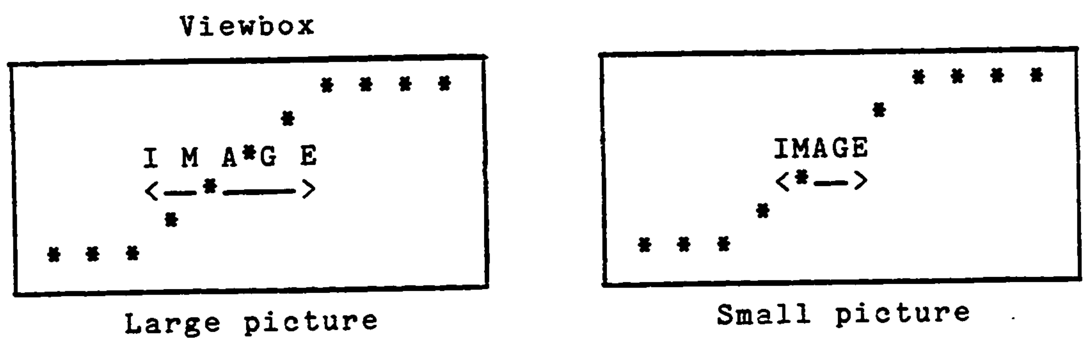

Figure 6.12: No Size Compensation

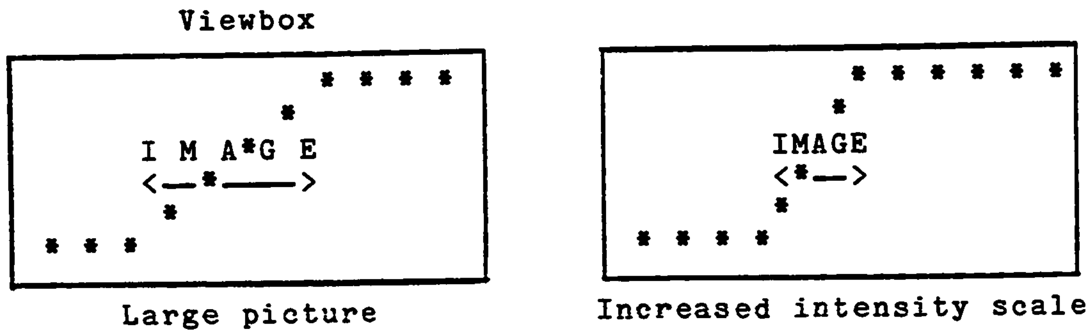

Figure 6.13: Compensated by Scaling

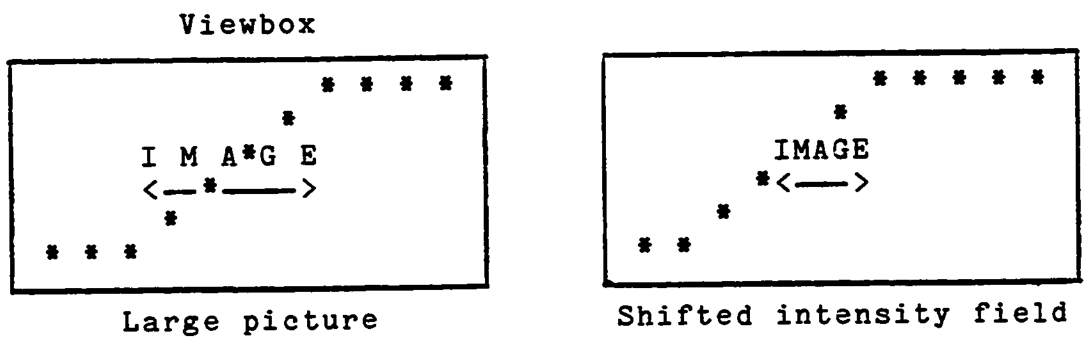

Figure 6.14: Compensated by Translating in Z

What I think is perhaps the most correct scheme for fading and for intensity depth-cueing, if steep intensity depth-cueing is not available, is shown in Figures 6.15 and 6.16 respectively. Fading pivots about the zero-intensity plane. Intensity depth-cueing pivots about the front plane of the image. When relative fading and intensity depth-cueing are combined in one picture, the sub-images share a common zero intensity plane. This scheme was too demanding for me to implement in the time I had.

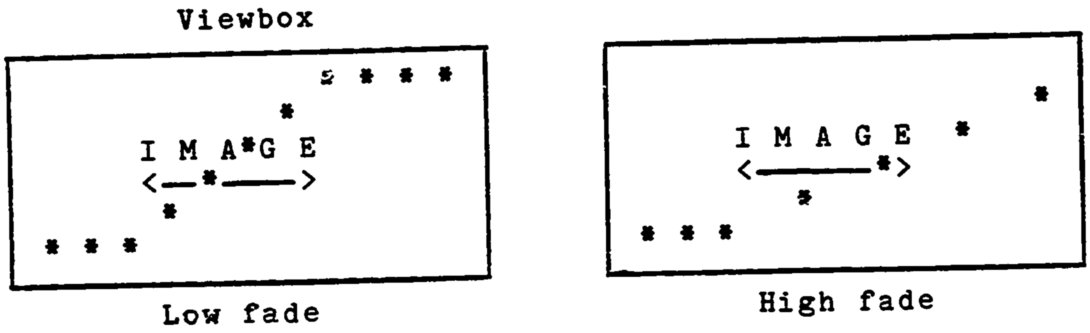

Figure 6.15: "Best" Fading

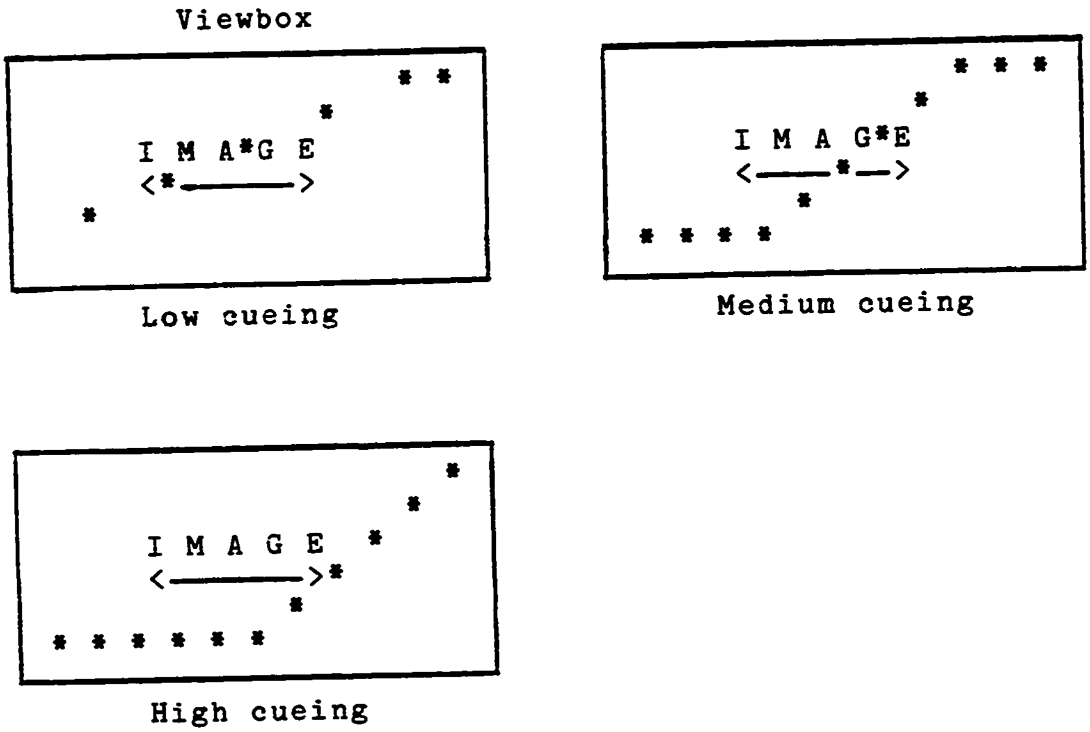

Figure 6.16: "Best" Intensity Depth-cueing

6.4 Z-CLIPPING

Z-clipping simplifies the picture by removing unwanted lines (Figures 6.3 and 6.17). A "curtain of invisibility" is swept through the viewbox in z. Only what lies behind it (away from the user) is visible.

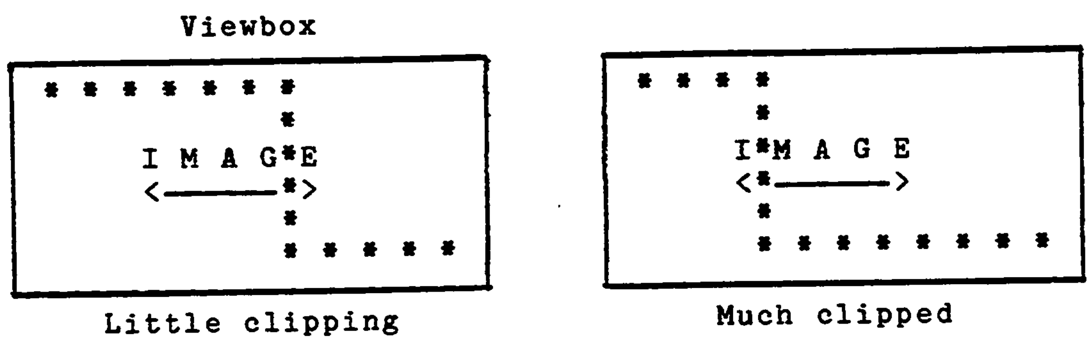

Figure 6.17: Scheme for Z-clipping

Unfortunately, the hardware requires the surface of the clipping plane to be fixed at the maximum (saturated) intensity. This nearly disables fading and intensity depth-cueing. Figure 6.18 shows depth-cueing doing its best. Although it is really hopeless, the fading and intensity depth-cueing controls are as orthogonal as possible to z-clipping. Relative fading works by relative intensity depth-cueing (Figure 6.19), exercising the only degree of freedom made available by the hardware in this situation.

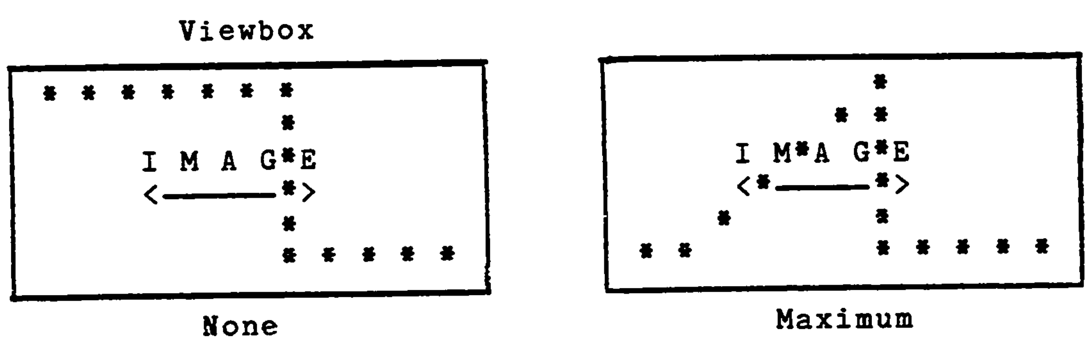

Figure 6.18: Intensity Depth-cueing with Z-clipping

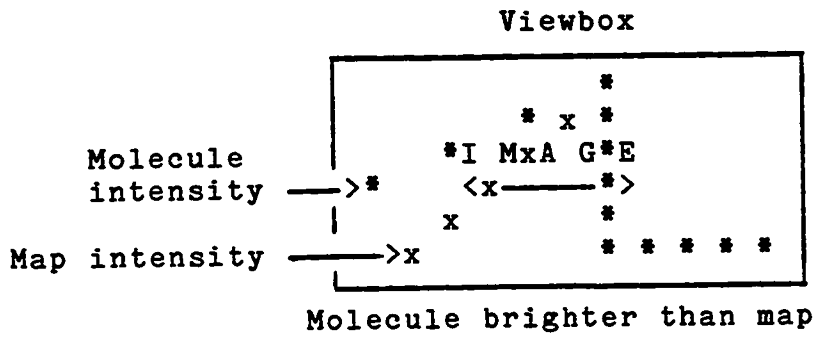

Figure 6.19: Relative Fading with Z-clipping

Blanking the front rather than the back is unnatural. This reverses the effective kinesthesia of the dynamic view-point-rotation control, the tip of whose joystick is usually connected to the forward (now invisible) part of the picture, because z-clipping only shows the back of the which moves in the opposite direction as the tip of the toothpick.

When certain settings of a dynamic control for some feature cause some special state to be entered, care must be exercised to avoid the danger of inadvertently leaving this special state and turning on the dynamic feature just enough to get its disadvantages (which the special state sought to avoid) but not enough to obtain its benefit. As I provide no positive indication that z-clipping is "on" other than the slider being off its peg, it causes perhaps more trouble than it is worth.

This change of state that turns on z-clipping, when its slider is moved from its "off" position, is far more accident-prone than the relative fading slider (page 103 above) which uses the z-clipping plane to enforce full fading of the molecule or map for hardcopy plots. The greater problems with the z-clipping slider arise not from a programming, electrical, or mechanical difference, but rather from a social difference in when and why the relative fading and the z-clipping sliders are moved.

Crystallographers use the z-clipping slider only occasionally and then mostly for preparing hardcopy plots. Our plotter package does not display intermediate intensities.

The z-clipping slider moves towards or away from the user mimicking the apparent motion on the screen. The effect is scaled for the minimum dead zones at each end of its travel.

The most picture is visible when the slider is nearest the user (down). This is opposite the usual convention of "on is up, off is down" [Woodson and Conover, 1964] (page 14 above), but I felt that having the control move in the same rather than the opposite direction of the controlled object took precedence. For uniformity I then defined "on" as "down" for the other two intensity controls. With all three controls fully "on", relative fading shows only the normally brighter object, the molecule, with no intensity depth-cueing.

Size compensation is necessary to keep the clipping plane fixed to the image as its sized is changed for side-by-side stereo and for polarized-mirror stereo (Figure 6.20).

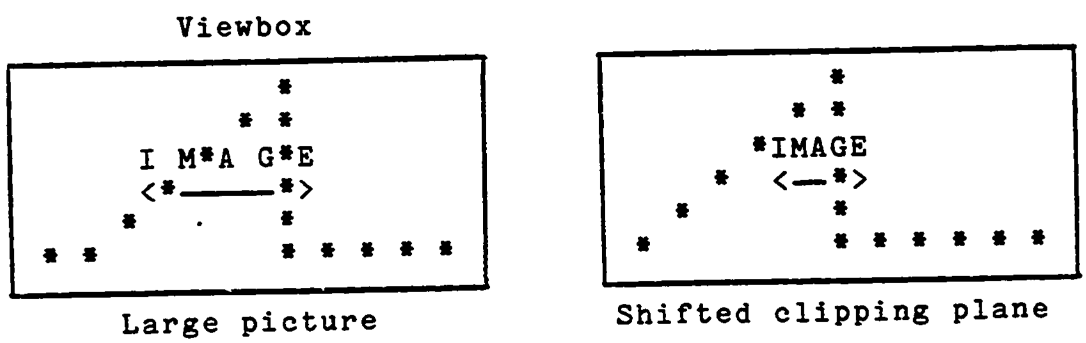

Figure 6.20: Size Compensation for Z-clipping

6.5 IMPLEMENTATION

The intensity package "falls asleep" when no intensity change is needed. That is, it is not invoked unless its sliders or the picture size slider are actually moving. At other times it would calculate the same display unit intensity control values over and over. This unnecessary work is not done, freeing the processor to produce more rapid rotation updates for smoother motion.

The subroutine that maps user actions onto intensity register values had to be interruptable, because in stereo the subroutine's output values (register pair values) may be seized at any moment. Therefore they are updated as close to each other in time as possible. If caught halfway through that process, the effect (which lasts one refresh) can only really be seen when z-clipping is used. I have seen it exactly once. Such short cuts are acceptable only if the users do not notice them.

The extreme positions of the intensity control sliders set the intensity package in special states. I supply some margin there to assure proper function will continue should the slider voltages go out of adjustment, mapping the full mechanical range of the sliders to possibly somewhat less that their expected electrical range.

The intensity package is compatible with and orthogonal to all other features of the system except the light pen.

6.6 OTHER USE

D. Tolle and M. Moore, also in the department here, incorporated this package of structural cues without incident into graphics projects for three-dimensional drafting and for 3-D display of electrical activity of the human heart.

6.7 COLOR

Color would be ideal for distinguishing the molecule from the map. Test photos by M. Pique and me show it to be better than relative fading. Unfortunately, our color hardware imposes stiff penalties for its use. Rather than fight that, I chose not to implement this feature. Even so, I can offer a few comments.

What about color blindness? Rushton [1975] makes two pertinent observations. First, only about 8 percent of all men and 0.4 percent of all women have defective color vision. Second:

> People with defective color vision can nearly always see a range of colors, particularly along the yellow-blue axis. Thus they tend to resent the implication that they are in some way blind. Their deficiency almost always shows up as abnormal red-green discrimination.

Indeed, only about 3 out of a million men and 2 out of a million women are totally color blind [Judd, 1943]. Proper color selection and system design should suffice. Even if a user unable to distinguish any colors can be found, he or she will see them as having different intensities, and will have the full benefit of relative fading.

More test photos suggest that color depth-cueing may be a poor substitute for intensity depth-cueing.

Color depth is a 3-D effect which arises because some colors seem inherently closer to the observer than others. Color depth should cause little trouble when color is employed to distinguish between objects superimposed in space. Land and Sutherland [1969] built a graphic system with a spinning disk for both stereo and color. It allowed dynamic movement of some simple colored objects into and out of the screen.

> One interesting observation is that all who experienced the stereo phenomena have been able to make the objects apparently coplanar (to within 0.3 mm separation), regardless of the color choice for each object and independent of both being in front or behind the actual display plane. The artistic sense of red advancing and blue receding is apparently more subtle than produced by this experiment.

We have a color Evans and Sutherland Lorgnette', but mechanical problems precluded routine production.

' Evans and Sutherland Computer Corp.
580 Arapeen Drive
Salt Lake City, Utah 84108

Chapter 7

STEREO

7.1 INTRODUCTION

Although stereoscopic display is the most obvious 3-D technique, users indicate that it is less effective than dynamic rotation. It appears to be superior to intensity depth-cueing. In general, no mater what combination of dynamic rotation, intensity depth-cueing, and stereo is in force, adding another of the techniques increases the 3-D effect. This suggests some orthogonality among these classes of 3-D cues.

Not everyone has stereoscopic vision. It is difficult to decide exactly what constitutes "deficient" stereopsis because of the spread of test scores on stereopsis tests, but various studies show that about 16 percent to 25 percent of the adult population has a definite shortcoming [Julesz, 1971; Kuhn, 1950; Snell, 1945]. This percentage changes with age. Assuming an average value of about 20 percent, the population with deficient stereopsis drops from 25 percent to 15 percent from ages 25 to 35, and then rises from 15 percent to 35 percent from ages 50 to 65 [Tiffin, 1942]. About 2 percent to 5 percent of the adult population is totally stereopsis-deficient, i. e. stereopsis-blind [Hofstetter, 1956; Julesz, 1971].

About half of our clients use some form of stereo at least occasionally. This is well below the percentage of those capable of stereovision. Most of this discrepency arises from the small number of vectors available in rotating-shutter stereo, from the lack of an adequate side-by-side stereo viewing device, and from the poor mounting and small image size of our polarized-mirror viewer. But there may be other reasons.

Individual users either strongly like or strongly dislike working with stereo. Feelings run high, whereas they do not with intensity depth-cueing. Perhaps stereo helps too much. Users of monocular viewing decompose their 3-D manipulations into a series of 1-D or 2-D motions largely in the plane of the screen. They dismiss stereo as "unnecessary" even though Kilpatrick [1976] states that such decomposition of 3-D motions arises from defective 3-D perception. Clearly stereo aids that perception. Dynamic viewpoint-rotation plays a much more active role in our system than in Kilpatrick's. Our monocular users alternate this motion with substructure manipulation. Could they be trying to limit their perception to 2-D in order to match it to their intended motions? And could stereo interfere with this? If true, then stereo is desirable only insofar as it aids 3-D perception during viewpoint-rotation, not during substructure manipulation. Also, the memory effect of the KDE (Wallach, O'Connell, and Neisser [1953], page 13 above, and Ortony [1971b], page 26 above), in which an image still appears 3-D to the user for a short time after viewpoint-rotation stops, is weak for complex pictures, and users may want to turn it on and off as it suits their purpose.

Kilpatrick's monocular users decomposed their motions in the viewing space (page 33 above), while the GRIP-75 monocular users decompose their motions in the 2-D projection of the user space onto the graphic screen. The difference no doubt arises from the different scenes in these two applications, blocks on a table versus molecules. Interestingly, the introduction of stereo drives both sets of users to a common ground for motion decomposition, the full 3-D user space.

Stereo effects are important to consider even when no stereo aids are used with dynamic rotation. W. Sloan (UNC graduate student) observes that without the stereo option the image really does look like a physically real object from across the room, whereas seated before the screen the impression is weaker. When viewed from afar, the difference between stereo and monocular display is the presence or absence of a negligibly small stereo disparity; no difference at all really.

I have observed users employing less relative fading between the molecule and the map in stereo than do users of monocular viewing. It could be that stereo helps distinguish the molecule from the map enough that users can then enjoy a better, brighter view of the map.

All of our stereo display techniques show the image alternately to each eye. Efron [1957] gives several results that apply to this case. The stereo fusion rate is lower than the flicker fusion rate. Like the flicker fusion rate, it drops with decreasing luminence and with a longer duty cycle (as a slow phosphor might do). Ogle [1963] gives more technical details.

The stereo techniques were made compatible with and orthogonal to dynamic fitting (even allowing stereo-mode change during fitting), the three intensity controls, all interactive commands, and the light pen.

7.2 STEREO DISPARITY

Neither the eye separation nor the screen viewing distance vary much from user to user. I intended to see what angle of stereo disparity (page 9 above) users liked and then to retire the slider for it. The "correct" angle is about 6 degrees, but users selected everything from 3 to 16 degrees with the average somewhat greater than 6, so I kept the slider. Kilpatrick [1976] had a similar experience (page 31 above). Exaggerated stereo disparity is "pleasant" and is for this reason often employed in stereo photography. M. Hollins speculates that it might also help relieve the absence of perspective in the display.

An obvious question is raised by the fact that users usually set the stereo disparity higher than the theoretically correct value. Does the (roughly) spherical volume of space presented on the screen take the shape of a football (prolate spheroid) pointed at the user within which the image undergoes a peculiar non-rigid rotation? The answer is an emphatic no. The kinetic depth effect immediately and completely dominates the stereo disparity by producing the illusion of <u>rigid</u> 3-D rotation [Wallach, Moore, and Davidson, 1963] and also [Wallach and Karsh, 1963a; Wallach and Karsh, 1963b].

The slider sets the stereo disparity negative as well as positive to allow for the rotating-shutter stereo viewers and their phase-angle problems. Also they are sometimes hung down instead of held up. This exchanges the sign of the stereo disparity. The slider is a "d" control rather than an "d*¦d¦" control, because large stereo disparity is common, and so should map onto a large region of the control.

The axis for the stereo effect passes through the origin for dynamic viewpoint-rotation, the center of the cubical viewing window. This minimizes the focus problem. The eye's focusing distance is driven in part by the amount of parallax. Parallax shows the image both in front of and and behind the display screen. Lines brought closer to the user by parallax tend to be focused closer by the eye than they really are, unless the eyes can decouple their focus from their parallax. Placing the axis of stereo rotation through the center of the viewing cube minimizes the maximum excursion of a line from the plane of the screen. The pivot point for viewpoint-rotation has to be on that same axis to minimize its use of the screen's surface.

Nearly all our users like to work with the room dimly lit. The screen is also mostly dark. In near darkness the eye tends to focus closer than the viewed object. This is called night myopia [Leibowitz and Owens, 1977]. M. Hollins speculates that night myopia does not occur here, but the question remains open.

The left eye and right eye views are sheared slightly to either side of the calculated monocular view (Figure 7.1). This avoids the appearance of rotation when snapping to or from stereo display. At one time W. Siddall attempted to save storage and increase the update rate, at the expense of the appearance of rotation just mentioned, by disabling stereo rotation for one eye. He substituted the more conventional scheme in Figure 7.2. Calculating only one set of sheared picture transforms saved so little storage and time that he actually restored the calculations as before (Figure 7.1), just to regain the benefit of no apparent rotation.

The rotation nesting data structure must be traversed to prepare the monocular view. Once this is done, preparing the left eye and right eye views takes little time. So little time is needed, the stereo calculations when bunched together doing their work once every update take less time than for one refresh to occur. But in stereo, the refresh loop picks up the current transformations at the start of each refresh, ready or not. If these are only half-updated from the previous update, then that is how they are shown for one refresh. It does not happen often, because the stereo calculations have been designed to be rapid. The effect is not desirable, but it is hard to spot when it does occur.

The rapidity of the stereo calculations relative to the entire update time approximates viewing the changes in the picture suddenly each update. This is important. The changes may instead be posted on the screen gradually as the update proceeds. We tried this and as the picture was rotated, it wiggled in a rubbery, caterpiller-like fashion which was remarkably unpleasant to see.

The stereo transformations in Figures 7.1 and 7.2 were shears (Figure 7.3) rather than rotations (Figure 7.4) as one might expect. A shear looks like a rotation to the user, is faster to calculate, and handles this display's z-clipping correctly (unlike a rotation).

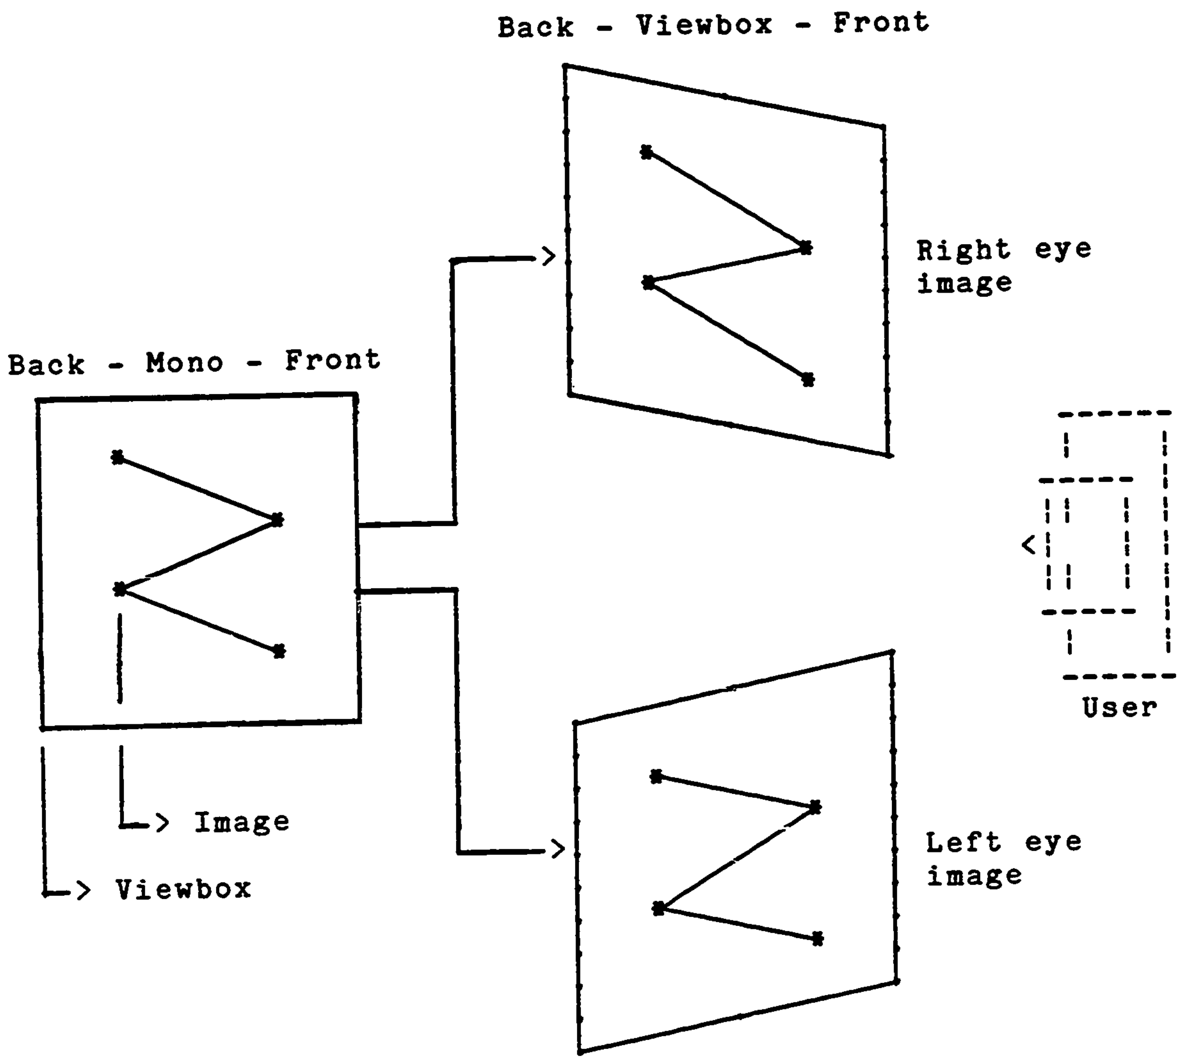

Figure 7.1: Production Stereo (Top View)

The shear matrix masquerades convincingly as a rotation matrix, because it is a rotation matrix to which two approximations have been made.

Let M be the monocular-view rotation matrix.

$$M = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix}$$

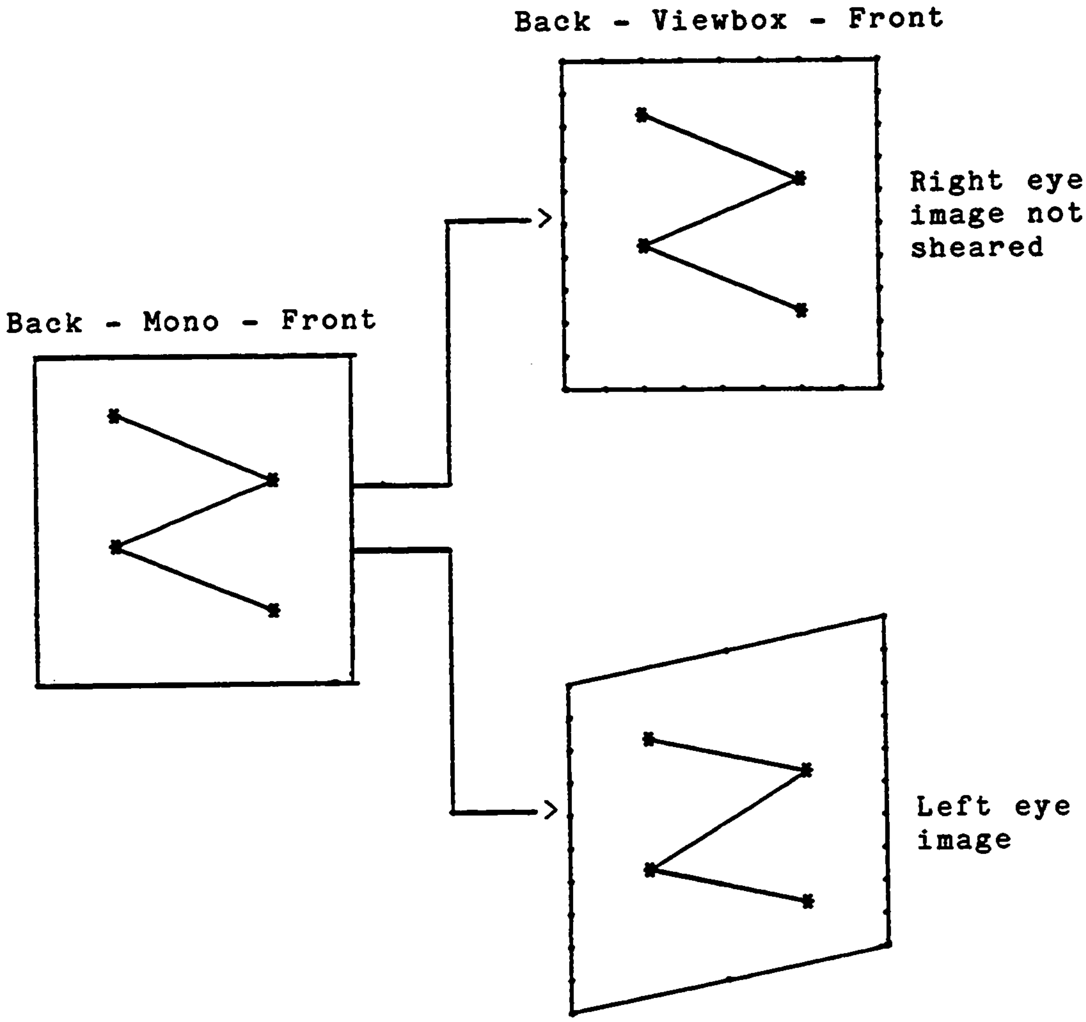

Figure 7.2: Attempted Shortcut

Let Y be the incremental stereo rotation matrix for the left eye about angle y, half the stereo disparity.

$$Y = \begin{bmatrix} \cos y & 0 & \sin y \\ 0 & 1 & 0 \\ -\sin y & 0 & \cos y \end{bmatrix}$$

Let L be the matrix which actually delivers the picture to the left eye.

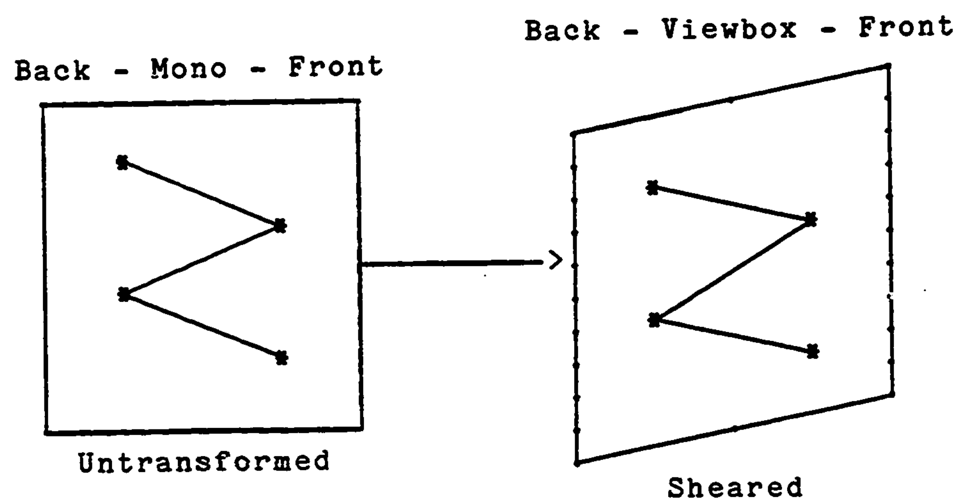

Figure 7.3: Sheared Image

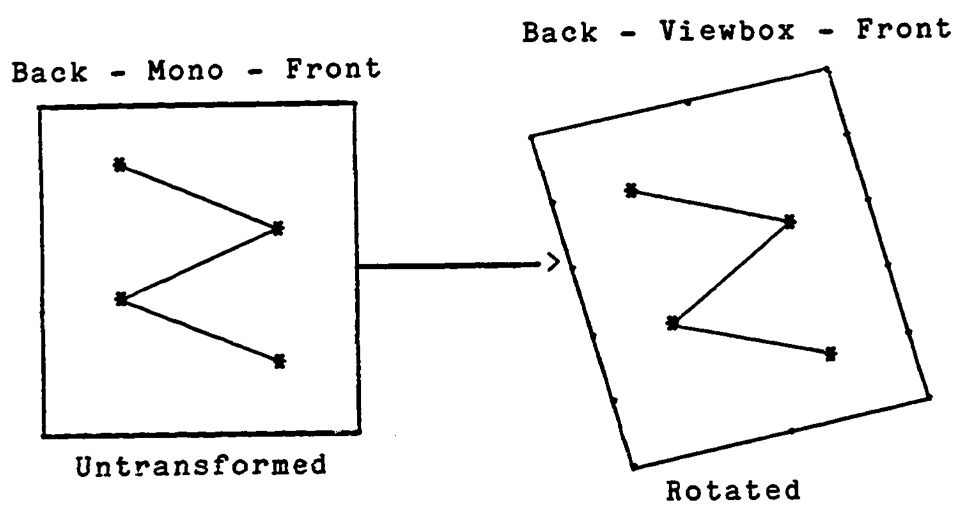

Figure 7.4: Rotated Image

$$L = M\,Y$$

$$= \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \begin{bmatrix} \cos y & 0 & \sin y \\ 0 & 1 & 0 \\ -\sin y & 0 & \cos y \end{bmatrix}$$

$$= \begin{bmatrix} a \cos y - c \sin y & b & c \cos y + a \sin y \\ d \cos y - f \sin y & e & f \cos y + d \sin y \\ g \cos y - i \sin y & h & i \cos y + g \sin y \end{bmatrix}$$

The first approximation to try is

$$y = \sin y$$

$$1 = \cos y$$

where y is in radians. This is valid for small y. Indeed, y is typically 0.05 radians (3 degrees) for a stereo disparity of 6 degrees. This produces

$$Y' = \begin{bmatrix} 1 & 0 & y \\ 0 & 1 & 0 \\ -y & 0 & 1 \end{bmatrix}$$

which gives

$$L' = M\,Y'$$

$$= \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \begin{bmatrix} 1 & 0 & y \\ 0 & 1 & 0 \\ -y & 0 & 1 \end{bmatrix}$$

$$= \begin{bmatrix} a - c\,y & b & c + a\,y \\ d - f\,y & e & f + d\,y \\ g - i\,y & h & i + g\,y \end{bmatrix}$$

a substantial saving.

The second approximation follows from observing that small motions into and out of the screen for each eye individually (the z-axis) should be visible only as small intensity changes if intensity depth-cueing is present. Otherwise, there is no change in the picture. We shall cover z-clipping in a moment. Neglecting the z component we have

$$Y'' = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -y & 0 & 1 \end{bmatrix}$$

and therefore

$$L'' = M\ Y''$$

$$= \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -y & 0 & 1 \end{bmatrix}$$

$$= \begin{bmatrix} a - c\,y & b & c \\ d - f\,y & e & f \\ g - i\,y & h & i \end{bmatrix}$$

L'' is simply a shear along the x (horizontal)-axis whose magnitude is determined by the z (out of the screen) component. The shear matrix requires very little time to compute.

The shear matrix has the added advantage over true rotation of correctly handling the z-clipping plane. With true rotation, small portions of the picture are clipped for one eye but rotate out of the clipping plane to be seen by the other eye. This is not so with a shear matrix. Each eye sees the same picture stretched in different ways. This gives the illusion of the clipping plane being in the viewing frame of reference rather than in the user's frame of reference. The irony is that with a shear matrix both the image and the clipping plane behave as if they rotate as the stereo disparity is changed, whereas neither one actually does.

7.3 TIME-MULTIPLEXED STEREO

The stereo effect can be achieved by means of a shutter that alternately blocks the view for each eye, while the display unit runs in synchrony. The picture appropriate for each eye is drawn on alternate refreshes. The pictures are broadcast through the same space, distinguished only in time. Therefore, we call this form of display <u>time-multiplexed stereo</u>. It is also known as <u>Field Sequential</u> stereo.

We have two such viewing devices for our system. We have a spinning disk called a Lorgnette made by Evans and Sutherland Corporation. We also have a rotating tube' made by Bausch and Lomb (Figures 7.5 and 7.6) which was interfaced to our PDP-11/45 computer by P. Reintjes. The rotating tube device gives a clearer view.

Among our users, Dr. Kim is especially partial to this form of stereo. Oddly enough, most of his graduate students like it too. The reason is probably the same as for Mrs. Jane Richardson, and therefore her students, using the otherwise unpopular automatic rocking and spinning (page 70 above). When one chemist teaches another to use the system, he also teaches his particular method of use to which the student, through practice, becomes personally attached.

' Bausch and Lomb Stereo Image Alternator System: Viewing Shutter and 3-Motor Controller.

TBR Associates Inc.
P.O. Box 8129
Rochester, New York 14617

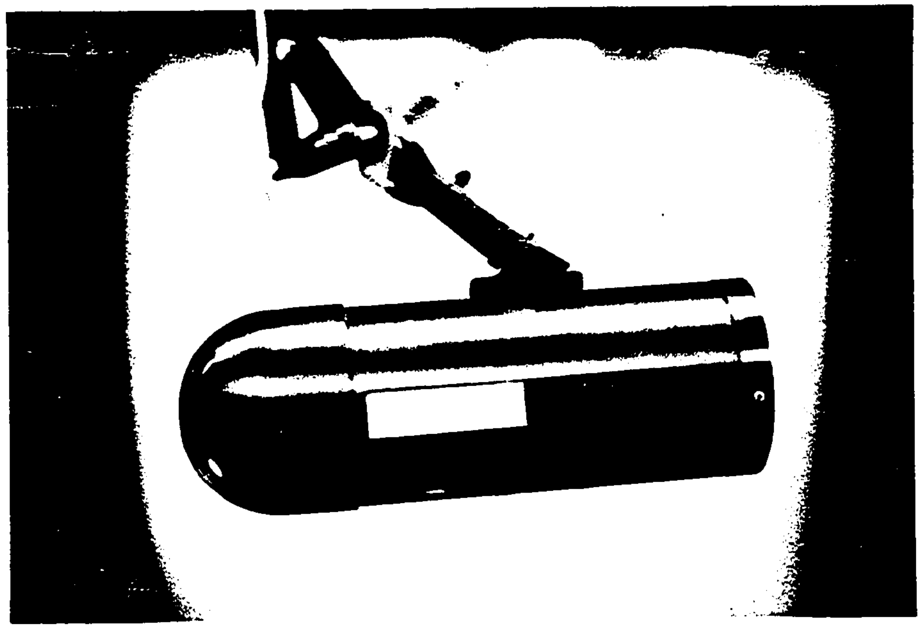

Figure 7.5: Rotating-shutter Viewer -- Left Eye View

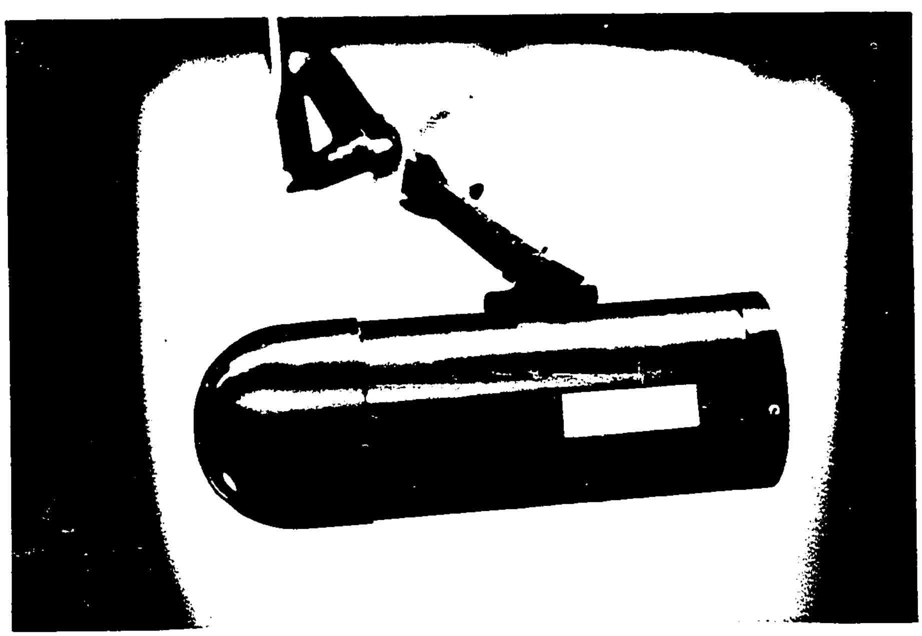

Figure 7.6: Rotating-shutter Viewer -- Right Eye View

Rotating-shutter stereo allows a full-screen stereo view as would two picture tubes and a combining mirror. It gives a brighter image and is perhaps cheaper at the expense of requiring the user to look through a closely-mounted viewing device which clips out lines that there is no time to draw. When the shutter comes around, it is time for the next refresh, ready or not. In practice this allows one, sometimes two, but rarely three sets of perpendicular contour planes to be shown for the 12 angstrom cube used in normal work.

Coordination of the light pen with rotating-shutter stereo was something of a problem because of the peculiarities of the Vector General Series 3 display unit. The user may wish to select something that is no longer visible to him because the display cannot get to it before the next eye is uncovered, and it is therefore time to restart the picture from the beginning. Also, the user can see some objects that the light pen cannot pick. When the light pen initially sees an object, the display is drawing the picture too rapidly (fast-mode) for the hardware to identify just which line or character it was that triggered the light pen. The question is resolved on the very next refresh, which is performed more slowly (slow-mode) so that there will be no problem finding out what is being drawn when the light pen sees it (hopefully) again. Unfortunately, the inertia of the rotating shutter makes it impractical to track the rapidly changing speed of the display. As a result, the user may point the light pen at something that is visible in fast-mode, but that the display does not have time to draw in slow-mode within the fixed refresh time accorded each eye by the rotating shutter. The programmer cannot find out that this is happening.

Another problem with the light pen is that parallax problems can arise, because the user sees a stereo image while the light pen does not. If the user points the light pen up at an object floating in 3-D just in front of his nose, he may wonder why the light pen does not see it when, in fact, the light pen's field of view now covers only the ceiling tiles and nothing on the graphic screen.

I was able to solve these problems with the light pen well enough that users who liked the rotating-shutter stereo picture could get along with the light pen. Whenever the light pen sees something in fast-mode that it can also see in slow-mode, or whenever it does not and the user taps the light pen switch, two things happen. First, the stereo disparity is set to zero which squashes the picture flat against the surface of the screen. This helps to match the user's perception of the scene to that of the light pen. It is acceptable because users intersperse dynamic motions (which require stereo perception) with use of the light pen (which usually does not). Second, the refresh rate is decoupled from the rotating shutter to ensure that none of the picture is clipped by it. Unfortunately, the user is still looking through the viewing device which is now unsynchronized with the refresh. The resulting flicker is made tolerable only by being brief. At least with the stereo disparity set to zero the picture does not appear also to oscillate.

The change back to rotating-shutter stereo is delayed about a second after the most recent light pen detect so that the light pen may be moved about some without excessive snapping back and forth from stereo to monocular display.

7.4 SPACE-MULTIPLEXED STEREO

The stereo image may be transmitted to both eyes simultaneously over the same spatial path to a decoding device just in front of the eyes, or over different spatial paths separately to each eye. We call these forms of display space-multiplexed stereo.

Two forms of space-multiplexed stereo are familiar to chemists. Side-by-side stereo (Figure 7.7) is the standard form of stereo picture publication in chemical journals. Polarized-mirror stereo (Figure 7.8) [Ortony, 1971a] (page 25 above) is employed by Feldmann's [1976] (page 34 above) AMSOM microfiche atlas of large molecules.

In a molecular graphic system these techniques allow a richer picture with less flicker, by using a slow phosphor, than does rotating-shutter stereo which needs a fast phosphor. Full-screen stereo with no problems introduced by the curvature of the screen may be had at the expense of a second video tube.

In GRIP-75 we have one display tube and, for rotating-shutter stereo, fast phosphor. With only one display tube, our two space-multiplexed images must fit in the same space that a single monocular or time-multiplexed image could occupy. Our space-multiplexed pictures are thus limited to about half the linear dimensions of the rotating-shutter stereo images (we do not have dynamic clipping) and have comparable flicker.

Substructure translation in these smaller images involves smaller movements across the screen for a given movement of

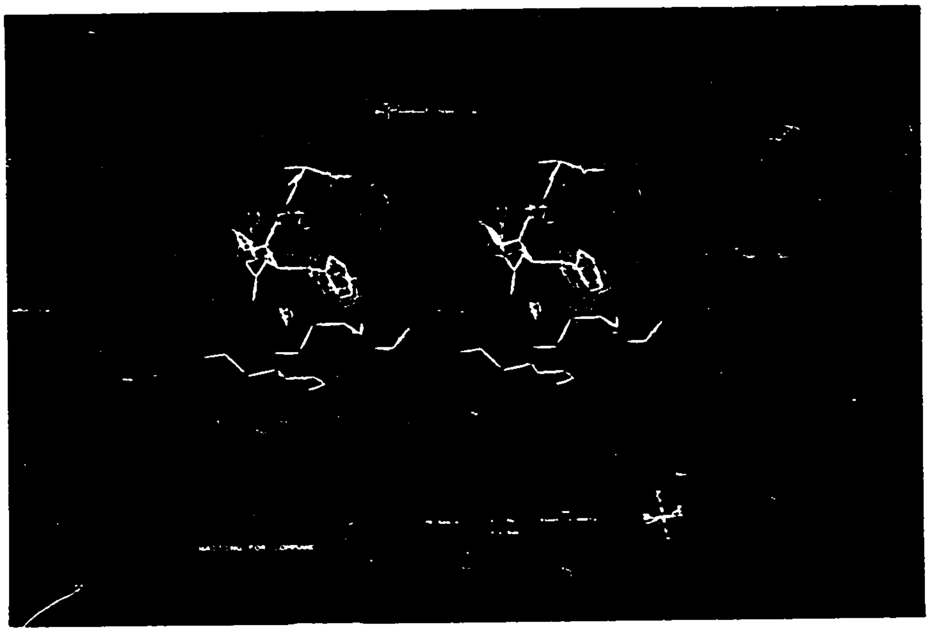

Figure 7.7: Side-by-side Stereo

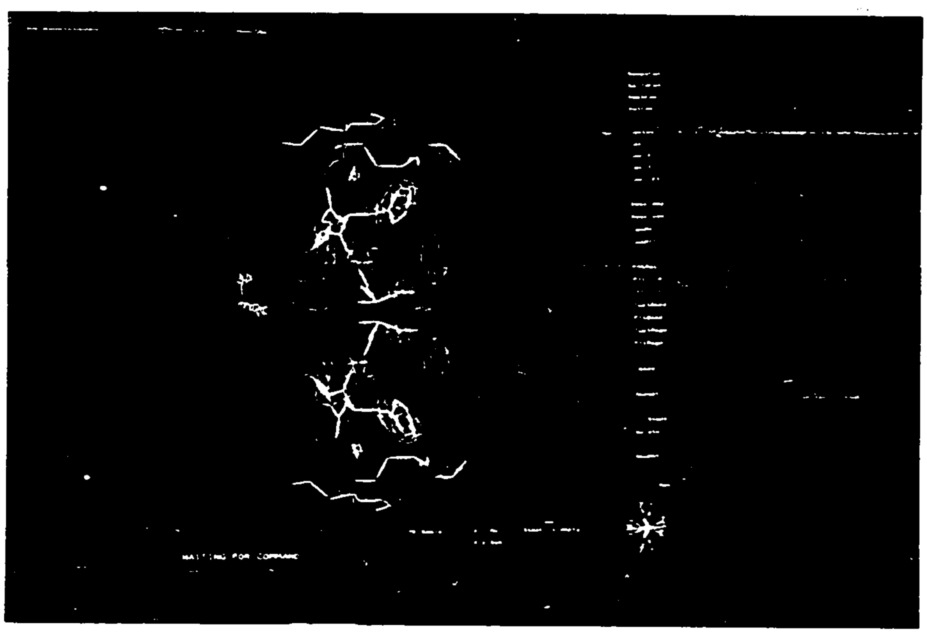

Figure 7.8: Polarized-mirror Stereo

the hand control. Users do not notice this change in the ratio of the size of the kinesthetic space to the size of the viewing space unless it is pointed out to them. This is consistent with Kilpatrick's [1976] experience (page 31 above).

With one display tube the two images should be equally distant from the center of the screen (Figure 7.9). The screen curves back equally from both sides of the vertical and horizontal midlines, and straight lines do not follow "great circle" routes across the screen's surface. Thus, placing the images equidistant from the midline of the screen minimizes (in side-by-side stereo) and cancels (in polarized-mirror stereo) the tendency of the curvature of the screen to deliver differently shaped pictures to each eye.

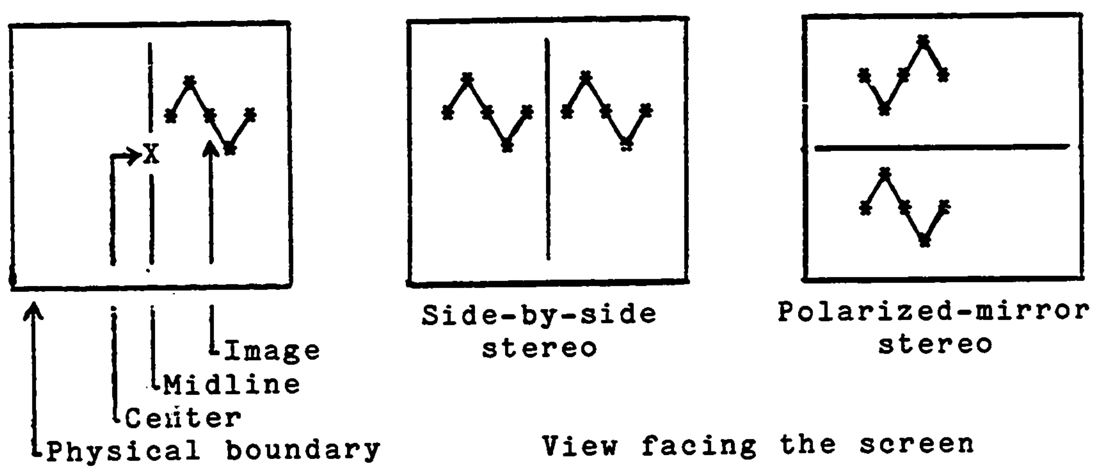

Figure 7.9: Images Equidistant from Midline

A slider sets the size of the two images with the amount by which they overlap each other (in the plane of the screen) held constant. The high and low ranges of the slider set the overlap to one of two values. The percent overlap should have been on a separate slider, but there were not enough to go around. Actually, when the images get large enough to bump against the edge of the screen and the menu, they are allowed to overlap each other and the menu in a two to one ratio. In this way they are allowed to grow to their full monocular size (figure 7.10).

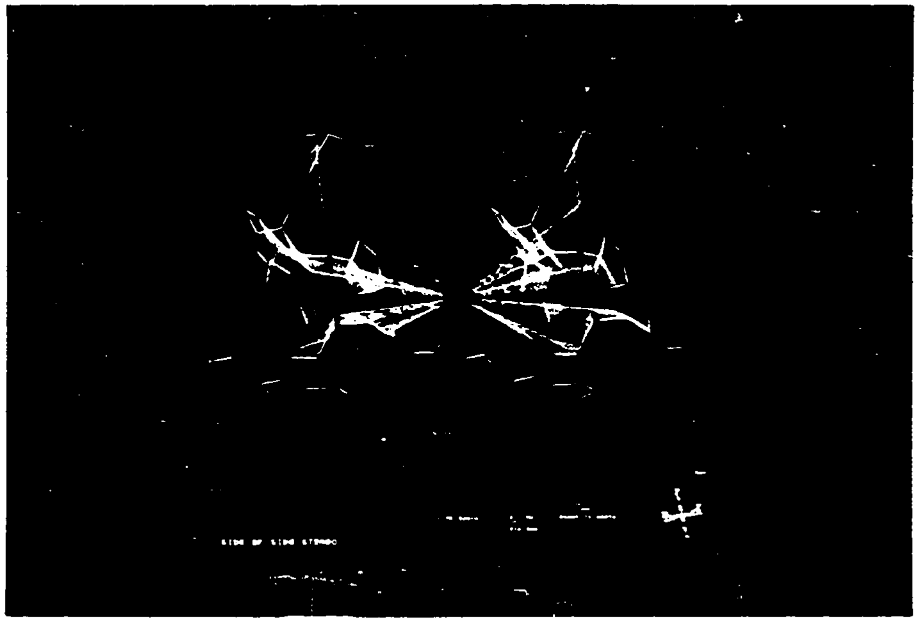

Figure 7.10: Side-by-side Stereo Size Adjustment

No control within reach of the client can throw the picture out of register with a stereo viewing device. Knobs that adjust such parameters are supplied by the display unit manufacturer on the side of the display. They are safely unreachable in normal viewing posture, and I do not make their functions any more convenient.

The picture is shown alternately to each eye with the following refresh scheme:

1. Left eye view and menu of commands.

2. Right eye view and menu of commands.

In rotating-shutter stereo the picture was squashed flat during light pen work, partially to prevent the image from oscillating out of the light pen's view, as the stereo disparity shifts a given line slightly left-to-right (on the alternate refresh). Remember from page 127 that the light pen can pick something on the screen of the Vector General Series 3 display unit only if the light pen sees it for two successive refreshes, the first in fast-mode triggering the second in slow-mode. That is, the light pen only sees things that are refreshed twice in the same place. Here the alternate refresh is some distance away on the screen, so drastic action is needed. The light pen does not normally see the 3-D image in space-multiplexed stereo, and it does not report out that it is trying. Thus the user has to tap the light pen switch to signal that he wants to take a hit, and the application program double refreshes each eye's image. The flicker is horrible, but there is no alternative for this display unit. Ortony's [1971b] work (page 26 above) suggests that the picture be flattened, but his experiments did not include anything like our preview character which jumps about to instantaneously mark what the light pen will hit if tapped. So I leave on the stereo disparity to see how well it is received. The implementation needs more work to remove artifacts of the implementation before a conclusion can be drawn.

7.4.1 Side-by-side Stereo

We do not have an adequate viewing device (stereo map reader) for this stereo method (Figure 7.7). If we did, it would allow one user at a time to see a complex and bright, though small, image in stereo.

It is used frequently in production for photography. For this reason alone it is worth building.

One or two of our hardy crystallographers actually use this stereo mode without a viewing device by simply crossing or wall-eyeing their eyes. For them the positive-negative range of the stereo disparity slider lets them choose whichever they want.

7.4.2 Polarized-mirror Stereo

Only a few of our clients use this form of stereo in production (Figure 7.8). Despite its potential, the lack of a suitable viewing system has held it back. Polarized-mirror stereo allows several users to see a complex though dim 3-D image. The users must each wear a pair of polarizing glasses.

Our first polarized mirror was designed to be usable for both this project and Kilpatrick's [1976] (page 30 above). As one might have expected, it was suitable for neither. It is too long and wide, and its two images differ too much in brightness from too much silvering on the mirror. Other than that, the design was sound.

A horizontal rather than vertical mounting was selected for two reasons. First, the picture could be viewed by several seated clients, rather than by one seated and one standing. Second, the natural side-to-side motions of the user's head would not change the relative brightness of the images presented to each eye. The lower polarizing filter was not hung down from the mirror. Instead, it was laid along the mirror's underside (Figure 7.11) as suggested by D. Barry at Washington University in St. Louis. With the polarizing material aligned as shown, it, like the more conventional hanging arrangement, inhibits internal reflections in the glass and delivers the brightest possible images. But it also makes the lower image available to the light pen at the minor cost of not delivering precisely cross polarized images to a user standing to one side.

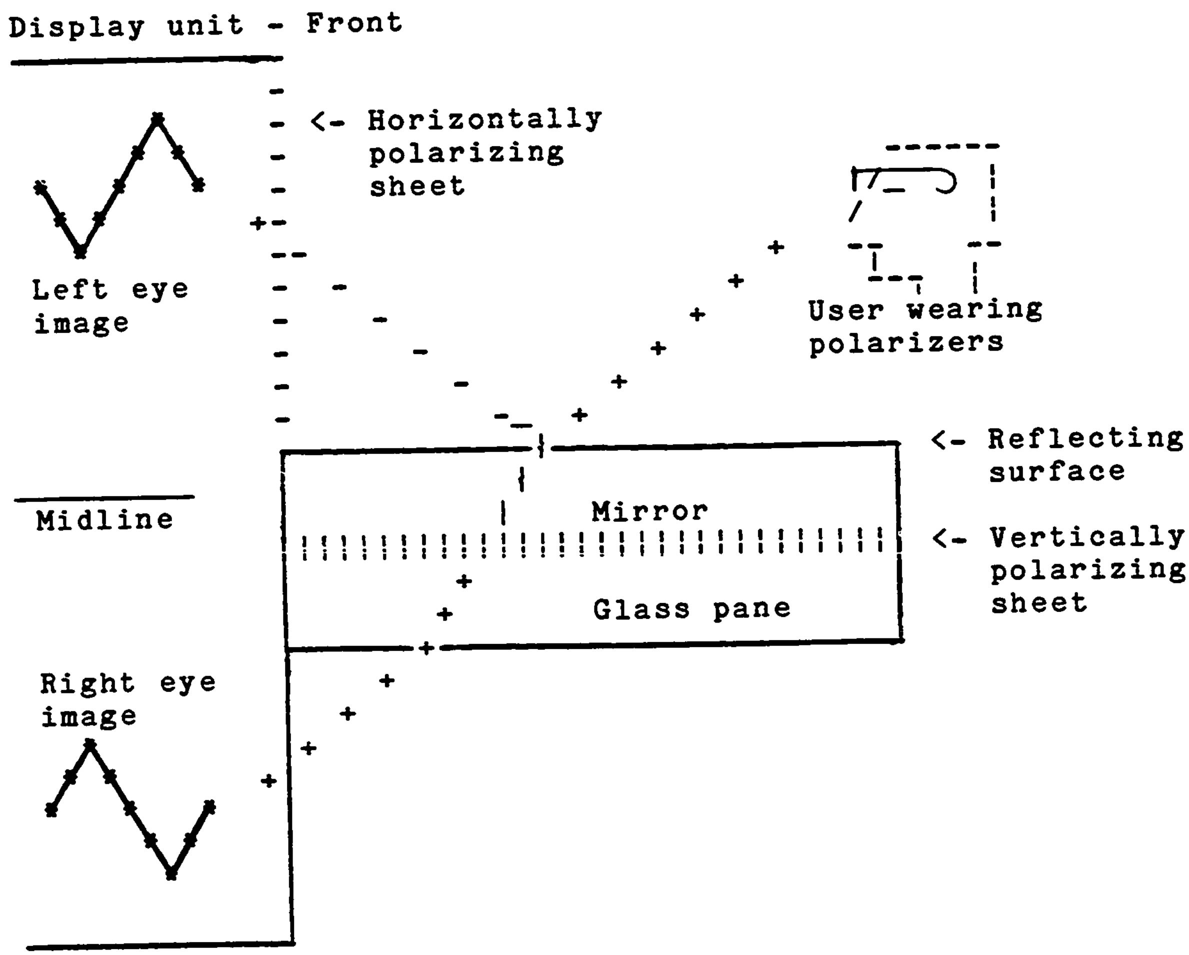

Figure 7.11: Scheme for Polarized-mirror Stereo

We have recently acquired an AMSOM viewer [Feldmann, 1976] (page 34 above) shown in Figure 7.12.

Figure 7.12: AMSOM Polarized-mirror Viewer

Remembering that our display shows vectors v' at location

$$v' = M v$$

where

$$M = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix}$$

then flipping one image upside down is easily done by modifying matrix M.

$$v'' = M' v$$

where

$$M' = \begin{bmatrix} a & -b & c \\ d & -e & f \\ g & -h & i \end{bmatrix}$$

A 180 degree rotation combined with a negative scale does the same thing.

Our display unit cannot flip characters upside-down, so they are automatically removed from the stereo portions of the polarized-mirror stereo picture.

Only the positive half of the stereo disparity slider's travel is useful here. Only half the stereo disparity slider's travel is meaningful at a time. The other half of its range inverts the sign of the stereo effect. The negative range is useful for cross-eyeing side-by-side stereo, and sometimes for rotating-shutter stereo, if the mechanism acts up.

7.5 LEFT VIEW AND RIGHT VIEW

None of the stereo techniques described above are compatible with the hardcopy plot feature, which will produce only one of the two alternating stereo refreshes. No warning is issued if the user commands a plot. This design shortcut has occasionally caused disappointment.

Another way is provided to get these 3-D plots. At any time in the 3-D mode the user may press the LEFT VIEW or the RIGHT VIEW button. This shows a full size picture corresponding to the left half or right half of a side-by-side stereo picture. He can then command a plot, press the other button, and command a second plot. The plots are annotated LEFT VIEW and RIGHT VIEW.

Photographs for publication are sometimes taken from these views on the screen rather than from side-by-side stereo, because each single image fills the screen and fills a full frame of film for best resolution on both the film and on the graphic screen. The alternate view is then selected without moving the camera.

Chapter 8

RESULTS

8.1 USE

Table 8.1 shows the usage of the system's 3-D techniques by about two-thirds of the GRIP-75 users. The selection is somewhat uneven; it is not contrived but neither is it random. The users are listed in approximate order of their first use of the system.

Chemist 2 declares that rotating-shutter stereo is necessary for him.

Chemist 6 stopped using intensity depth-cueing when she switched from monocular viewing to rotating-shutter stereo.

Chemist 7, when not actively fitting, would sometimes inspect his molecule using side-by-side stereo while crossing his eyes.

Chemist 10 has refused to use the system when the automatic rotation features were out-of-order. He likes to fit <u>while</u> the image is rocking or spinning, rather than alternating manipulation with viewing. Fitting while the image is rocking is unusual, but several other chemists do this. Fitting while spinning is very unusual. He much prefers rocking to spinning or to 90-degree turn.

Chemist 12 reports having dreams of rotating molecules after an intensive week of work. I do not question users about their dreams, and so do not know how common this phenomenon is, but E. Britton and I suffered similar experiences during the early months of intensive system development. Significantly, all victims of these unwanted visitations report seeing the molecules in motion. They cannot recall whether or not the image is intensity depth cued or in stereo.

Chemists 9 and 12 wanted horizontal motions of the Noll box to cause translations in the (vertical) plane of the screen. Chemist 13 had yet another idea. He rotated the device vertically by 90 degrees to achieve correspondence of

TABLE 8.1

Use of 3-D Features

Key

F: Frequently use
O: Occasionally use
*: Never use
-: Not working

Viewpoint-Rotation

MA: Manual
90: 90-degree turn
RK: Rocking
SP: Spinning

Manual Manipulation

UT: Unconstrained Translation
UR: Unconstrained Rotation

Structural Cues

RF: Relative Fading
DQ: Depth-cueing
ZC: Z-axis Clipping

Stereo Display

RS: Rotating Shutter
SS: Side by Side
PM: Polarized Mirror

	Rotation				Manip.		Structural			Stereo		
Chemist	MA	90	RK	SP	UT	UR	RF	DQ	ZC	RS	SS	PM
1	F	-	-	-	*	*	F	F	*	F	-	-
2	F	F	*	O	F	F	F	F	*	F	-	-
3	F	*	F	*	F	F	F	F	O	*	-	-
4	F	F	*	*	F	F	F	F	*	*	-	-
5	F	*	*	*	F	F	F	O	*	*	*	O
6	F	F	*	*	*	*	F	*	*	F	*	*
7	F	O	*	*	F	F	F	O	*	*	O	*
8	F	F	*	*	F	F	F	F	*	*	*	*
9	F	*	*	*	F	F	F	*	*	*	*	F
10	F	F	F	F	F	F	F	*	*	O	O	*
11	F	*	*	*	F	F	F	F	*	*	*	*
12	F	F	*	*	F	F	F	F	O	*	*	*
13	F	*	*	*	F	F	F	O	*	O	*	*
14	F	*	*	*	F	F	F	*	*	-	*	*
15	F	F	F	*	F	F	F	*	*	*	*	*
16	F	F	O	*	F	F	F	F	*	*	*	*
17	F	F	F	*	F	F	F	F	*	O	*	*
18	F	*	F	O	F	F	F	F	*	*	*	*
19	F	O	*	*	F	F	F	F	*	F	*	*
20	F	F	*	O	*	*	F	F	*	F	*	F
21	F	*	*	*	*	*	F	F	*	F	*	O
22	F	F	O	F	F	F	F	F	O	*	*	F
23	F	O	*	*	F	F	F	F	*	F	*	*
24	F	O	*	O	F	F	F	O	O	*	*	*
25	F	O	F	*	F	F	F	*	*	-	*	O
26	F	*	*	*	F	F	F	*	*	-	*	*

motion when he turned his head to look at it. The Noll box is both off to one side and below the screen.

Chemist 15, like chemist 10, translates and rotates the substructure during automatic rocking. He uses only one hand to operate the joysticks. He leans on his other elbow or rests his chin in his free hand. Automatic rocking and spinning were intended to be a "third" hand, but this chemist uses automatic rocking as a "second" hand.

Chemist 16 was trained by chemist 10 who would reach over and shut off intensity depth-cueing whenever chemist 16 turned it on. This training was effective for about 10 days, and then chemist 16 returned to using intensity depth-cueing. The use of 90-degree turn and rocking were more throughly inculcated by chemist 10 than the non-use of intensity depth-cueing. Chemist 16, like chemist 12, once had dreams, but in this case they were dreams of interchanging protein residues within a larger, undulating protein structure.

Chemist 18, like chemist 10, sometimes fits with the picture spinning continuously. But unlike chemist 10, she sets the spinning rate very low, just barely fast enough to preserve the 3-D illusion, but not so rapid as to disrupt her slow manual manipulations.

Chemist 19 likes rotating-shutter stereo, but uses it only part of the time because it tires her eyes. She is not very active with viewpoint rotation in stereo, because she feels she does not need its help (with 3-D perception) in stereo as much as in monocular display. She feels that intensity depth-cueing does not help as much in monocular display as it does in stereo. She sometimes spins the picture continuously, pushing the HALT button and taking a photograph when a revealing view comes around (this use not listed in table).

Chemist 20 likes rotating-shutter stereo. He uses it to view molecules, but because of the rotating shutter's limit on the number of lines it allows, he switches to polarized-mirror stereo to view electron density maps. He uses continuous spinning for periods of quiet contemplation.

Chemist 21 normally uses only a little intensity depth-cueing; its effect is just slightly visible. This is not unusual.

Chemist 22 finds polarized-mirror stereo helpful with complex pictures (many lines) and also when contouring with only one set of planes. He uses the "picture size" slider as a dynamic zoom since it can make the picture larger than the polarized-mirror viewing device. The viewing device then clips away the outer edges of the scene. He really prefers the larger, brighter stereo view offered by the rotating-shutter viewer, but it cannot display enough lines to meet his needs. 90-degree turn helps him to see the backside of the image, which the toothpick has trouble reaching. His says that his molecular structure has a definite front, side, and back, which can be inspected individually by successive 90-degree turns. Rocking was useful at first for quiet contemplation, but he stopped using it after a few days. Instead, he rocks the picture with the plus to minus range of the continuous spinning speed control. He also uses spinning, with one hand on its slider, to substitute for the vertical axis of the toothpick with polarized-mirror stereo. Otherwise the toothpick tends to bump into the polarized-mirror viewer. He had trouble adapting to the fist, and he agrees with my speculation that the two similar axes should be placed in the plane of the screen with the third axis (twist) coming out of the screen (and the device remounted to preserve the effect of the image moving in the same direction as the user's hand). Intensity depth cueing was useful at all times but especially with stereo. Z-clipping sometimes helped in polarized-mirror stereo to produce a satisfying balance between relative fading and intensity depth-cueing. He reports having dreams on two occasions. In his dreams he was translating large protein chains, very much like his real work. He was particularly struck by the clarity and the complexity of the dream images.

Chemist 23 likes working in stereo. Early in his visit, the rotating shutter did not work, and he used polarized-mirror stereo. Later, when the rotating shutter was fixed, he switched to using it exclusively.

Chemist 24 used spinning and 90-degree turn mainly to see the back of the molecule inaccessible by use of the toothpick alone. He used z-clipping only to inspect thin stacks of contoured electron density map planes rotated to face him.

Chemists 25 and 26 worked as a team. Chemist 25 liked polarized-mirror stereo a lot. He would have used it more, but the viewing device kept chemist 26 from getting a good view while chemist 25 was working.

Everyone used manual viewing rotation and relative fading. Everyone whose work required unconstrained manipulations used these features.

Note the frequent use of the 90-degree turn, intensity depth-cueing, and some form of stereo. Table 8.1 does not list the frequent use of side-by-side stereo and 90-degree turn for photography, nor does it list the use of other display techniques used for photography but not for fitting.

All users establish a preferred set of 3-D techniques in the first few days of work and then vary from it little or not at all. Local users appear to employ a wider variety of 3-D techniques than do visitors. Some 3-D features are so useful that the chemist really has no choice but to use them. These include manual viewpoint-rotation and relative fading, both of which are used by everyone, and also unconstrained manipulations. All other 3-D techniques are at the user's option. Local chemists use frequently an average of about 3 of these "optional" three-dimensional display techniques each, while visitors average about 2 each. (Occasional use is counted as half.)

A correlation is evident in table 8.1 between the orthogonality and the utility of the different three-dimensional display techniques. Continuous spinning disturbs the user's viewpoint more than does rocking and is less often used (page 71). Intensity depth-cueing interferes more with the light pen than does relative fading (page 105). Z-clipping (on our display unit) interferes with both relative fading and intensity depth-cueing (page 111) as well as with manual viewpoint-rotation (page 112). Sure enough, of the intensity cues, relative fading is used most, intensity depth-cueing the next most, and z-clipping the least. There are so many trade-offs among monocular display, rotating-shutter stereo, and polarized-mirror stereo that no ordering of orthogonality can safely be assigned among them. Unfortunately, no numerical value can be calculated for the general correlation between orthogonality and utility, because the amount of orthogonality among functions remains a subjective quantity, even when the degree of non-orthogonality (the number of conflicts) is clear. And, of course, the amount of orthogonality is not the only difference among these features.

The effect of individual user preference is evident as noise in the data presented in the table. It appears to be fruitless to attempt to assign a numerical value to this variance, because the interactions among features and among users are unknown and potentially complex.

8.2 CONCLUSIONS

Based on my experience with this system and my observation of other similar systems, I offer the following conclusions and recommendations.

1. Orthogonality:

   a) Three-dimensional display techniques that are orthogonal to other techniques appear to be more popular than those having side-effects (pages 3 and 140).

   b) Orthogonality of function returns the user to a "home state" even when the system cannot be returned to a "home state" (page 56).

   c) Orthogonality of function (combined with short commands) breeds a sentence-structured command language (page 54).

   d) Functions that are not naturally orthogonal should be made to act as independently of each other as possible (dynamic motions versus interactions, light pen versus stereo, rotation rate versus computation load, etc.) (pages 5, 54, 55, and 85).

   e) However, orthogonal functions should be mixed if that makes them appear more orthogonal. The only orthogonality that matters is that which the user believes he sees. Rocking rate and amplitude (and also intensity level and depth) perhaps should be linked in such a way as to give a better appearance of orthogonality than naturally occurs between them (pages 77 and 107).

2. Action Language:

   a) A rich menu of redundant techniques is needed for serving a large user group (pages 3 and 52).

   However, to some extent a person can be told what to like. When one chemist teaches another to use the system, he also teaches his individual style to which the student can become personally attached (pages 70, 125, and 138).

b) Users establish a preferred set of 3-D techniques in the first few days of work and then vary from it little or not at all (page 140).

c) Local users appear to employ a wider variety of 3-D techniques than do visitors (page 140).

d) Many of Kilpatrick's [1976] conclusions apply to this application as well as to his (page 30):

   i) users decompose their 3-D manipulations into successive 1-D and 2-D motions (pages 69, 89, and 116);

   ii) users adapt quickly to a difference between the size of their kinesthetic space and the magnitude of its effect on the image (pages 87 and 130);

   iii) when users make motions into the plane of the screen, they can see their effects only in their mind's eye, but they do depend on the magnitude of these effects to be equal to those of their motions in the plane of the screen (page 89);

   iv) users prefer to adjust the amount of stereo disparity (page 118);

   v) users often employ a larger stereo disparity than is theoretically correct (page 118).

e) Our users perform these acts most frequently during substructure fitting, (in decreasing order) continuing, on many levels, the strategy of motion decomposition in new and surprising ways (page 53):

   i) a series of similar 3-D motions (e.g. repeated translations or rotations) each decomposed into successive 1-D and 2-D motions;

   The user keeps his hand on the 3-D joystick and his eye on the molecule during these motions.

   ii) alternating one motion with another (e.g. viewing vs. translation, translation vs. rotation, etc.);

The user keeps one hand on each joystick and his eyes on the molecule. Note that when users have both hands on joysticks, they rarely move the two joysticks at once (pages 64, 69, and 71).

iii) re-positioning the hands on different dynamic-motion controls (tooothpick, Noll box, fist, and knobs);

The user generally looks down at the devices, places his hands, and then looks back at the molecule.

iv) interspersing dynamic motions with interactive commands.

f) Simultaneously active dynamic functions avoid delay for interactively commanding the change from one dynamic motion to another. Instead, the user need only move his hands from one device to another. Even this delay can be avoided by using few devices (or stations for the hands), which individually control many motions (page 54).

g) When users intersperse dynamic manipulations with interactions, simply leaving the dynamic devices active during and after these interactions avoids any delay in returning to dynamic manipulations (page 54).

h) Commands should not be sensitive to their context. This argues for a one-button--one-function design, orthogonality among functions, and short sentences (page 54).

i) One-button--one-function design, in turn, breeds short sentences and many buttons. This is good because the user then selects, rather than evokes, commands (page 54).

j) Short sentences are easier to use than longer sentences. Users often cannot cope with even moderately complex interactions. This appears to be because they are concentrating fully on their application, and so cannot and will not pay much attention to the computer system (page 55).

k) Short sentences breed simple syntax and reduce the importance of the choice between a prefix or a postfix command language (page 55).

l) The picture should move in the same direction as the user's hand, except where motion along or about application axes or structure has special significance. This aids 3-D perception of rotations and speeds manual manipulations (pages 65, 87, and 89).

However, if the physical device is too far to one side or too far below the screen, some users will want the device rotated 90 degrees to match the coordinate system of their turned heads or eyes (page 136).

m) N-key rollover can be useful and is easily programmed. I recommend that it be considered in future button-language designs (page 75).

n) A "grace period" is useful for "picking" devices in a reactive display. If the user's hand jumps while he makes the hit, the pick may still register properly (page 58).

3. <u>Dynamic Viewpoint-rotation</u>:

a) Smooth manual rotation is the single most effective aid to three-dimensional perception that I have tried (page 64).

b) The image should look and act like a physically real object (should be rigid, should have mass, etc.) for those aspects of realism that help the user achieve rapport with the image, quite apart from whether or not such realism is otherwise "useful" to the application (pages 64 and 72).

c) Users keep the image in nearly constant motion, through manual manipulation, for hours at a time (page 65).

d) Orthogonal views separated in space (side-by-side orthogonal views) or in time (90-degree turn) are quite useful in this application (pages 27, 29, 36, and 69).

e) Automatic rocking and spinning are helpful to only a small percentage of users. Continuous spinning is particularly unpopular and is therefore not worth much effort to implement (page 70).

4. Refresh versus Update:

   a) When the update rate cannot equal the refresh rate, it should be synchronized to an integer submultiple of the refresh rate, or else the motion will appear jerky (page 81).

   b) At three refreshes per update and beyond, the user perceives a rippling motion opposite to the direction of the image's movement (page 79).

5. Structural Cues:

   a) Intensity cues are popular, probably because they can simplify the picture (pages 98, 101, 105, and 110).

   b) To some extent the perceived brightness of an intensity-depth-cued object may be that of its near surface (page 108).

6. Stereo:

   a) Feelings run high, in the literature and among our users, both for and against stereo display. This is probably because stereo aids 3-D perception, yet the plane of the screen is often a useful guide for manipulations and is less evident in stereo (pages 21, 22, 25, 27, 29, 33, 38, 41, 42 and 116).

      For this reason, I recommend that graphic systems not make stereo display obligatory.

   b) A shear may be used in place of a rotation to achieve stereo disparity (pages 119-125).

8.3 SUMMARY

Users still often report insufficient 3-D perception, or exhibit it by motion decomposition. Three-dimensional perceptual problems with computer graphic displays have been solved only to the point that nearly all of a large group of people can surmount these problems enough to perform useful work. Three-dimensional perception in computer graphics remains a fruitful area for future research.

Chapter 9

FUTURE WORK

This thesis raises a number of questions whose answers could be pursued. However, major advances in 3-D perception in vector computer graphics will, I believe, require trying something completely different from the work done for this thesis.

The questions that are unresolved, or that need more research, are these:

1. Orthogonality:

   a) Orthogonality appears to be important. Those functions that do not interfere with each other are the more often used. How can the orthogonality versus utility of 3-D display techniques be measured accurately, and what is the value of this correlation (page 140)?

   b) Different people prefer different display techniques. This trend weakens the correlation between orthogonality and utility. Can a numerical measurement of this force be made (page 140)?

   c) When one chemist teaches another to use the system, he also teaches his individual style to which the student can become personally attached. To what extent, again numerically, can a person be told which techniques he will like (pages 70, 125, and 138)?

2. Action Language Design:

   a) A "grace period" is useful for "picking" with a light pen. If the user's hand jumps while he makes the hit, the pick may still register properly. Is it really true that this grace period is useful for picking devices in general, not just for the light pen (page 58)?

3. Dynamic Viewpoint-rotation:

a) Note that when users have both hands on joysticks, they rarely move the two joysticks at once. Why do users avoid making two motions at once? Is it because one motion hides another (pages 53, 64, 69, and 71)?

b) Rocking and spinning are sometimes good for contemplation, but they disturb manipulation. Should automatic rocking or spinning halt when the system senses a dynamic-motion input from the user (page 71)?

c) Should a rocking picture commanded to turn 90 degrees revert back to rocking? The general question is, should a command that takes several seconds to complete (90-degree turn) finish in a fixed state (halted) or revert to the previous state (rocking) (page 74)?

d) Giving the user independent control of the period and amplitude of rocking does not appear to be satisfactory. How, exactly, should the rocking rate vary to compensate for changes in amplitude, or is it better not to bother trying (pages 77)?

e) Orthogonal views separated in space (side-by-side orthogonal views) or in time (90-degree turn) are quite useful in this application. Can one evaluate how orthogonal views separated in time compare with those separated by space (pages 27, 29, 36, and 69)?

4. Refresh versus Update:

a) Are the multiple images of two or more refreshes per update caused by the image being painted on different places on the retina (page 79)? Most researchers believe this, but nobody seems to have done the experiment.

b) At low update rates the uneven spacing of the perceived multiplet of lines suggests that the eye physically turns and stops with each update. Is this really so (page 81)?

5. Manipulation:

a) I suspect that dropping the square and cube terms from the user-space rotation algorithm will speed it up without introducing significant error. Is this so (page 96)?

6. Structural Cues:

a) When an intensity-depth-cued object is made smaller, reducing the intensity drop-off with depth, my impression is that keeping the near surface at constant intensity fixes the perceived brightness. Is this true (page 108)?

b) Is the proposed "best" intensity depth-cueing really better than the one implemented (page 109)?

7. Stereo:

a) How and why does the good 3-D perception resulting from stereo display interfere with motion decomposition (page 116)?

b) To what extent does stereo r: Is this true (pa In near darkness the eye tends to focus close. Does this night myopia play a role in vector-graphics display in a dimly lit room (page 119)?

c) What are other techniques for "picking" in stereo and how do they compare (pages 127 and 131)?

In the near term, I believe the GRIP-75 system most needs better manipulation tools to ease the demands on the user's 3-D perception. Specifically, the system's on-line idealization should run while the user is manually fitting a residue. Idealization could then keep the molecular geometry, between the free-flying and the stationary structures, in reasonably good order at all times while the user concentrates on satisfying other constraints by eye. Obviously, the user should avoid making unreasonable demands on the molecular geometry when he has met those other constraints.

For the intermediate term, several advances in technology have been made which one would like to buy.

1. Dynamic picture size and better x,y,z clipping would be nice. Display units that support these features are well liked by their owners.

2. Color tubes are just now becoming available. Using color to distinguish between the molecule and the map, or for other uses, should make the picture easier to understand (page 114)?

3. Full-screen stereo could be had with one controller driving two display tubes at right angles to each other. A combining mirror at 45 degrees could superimpose the images, and large polarizing sheets could present one image to each eye. This could look quite good.

Also, in the intermediate term, we need better tools for manual manipulation. In this department, Dr. Brooks is directing several efforts which are exploring the wide range of alternatives. The furthest along on these is D. Voss [in preparation]. I personally favor a six degree of freedom velocity control for unconstrained translation and rotation (page 54).

In the long run, the 3-D perception problem needs to be solved. This will require one or more fresh approaches. In this department, Dr. Brooks is directing work on head-motion parallax, Dr. Fuchs is starting work on an advanced head-mounted display, and finally, Dr. Fuchs and Dr. Pizer are directing an ambitious varifocal mirror project whose first steps have been made by J. Cohen [1979] and P. Reintjes.

REFERENCES

Barry, C. D. 1976. COMPUTERIZED MOLECULAR GRAPHICS -- Some Notes and Observations. Technical Report No. 228, Computer Science Laboratory, Washington University, St. Louis, Missouri.

Barry, C. D., Bosshard, R. A., Ellis, R. A., and Marshall, G. R. 1974. Evolving Macromodular Modeling System. Federation of American Societies for Experimental Biology, Federation Proc., 33, #12 (December), 2368-2374.

Barry, C. D., Ellis, R. A., Graesser, S. M., and Marshall, G. R. 1971. CHEMAST: A Computer Program for Modelling Molecular Structures. Proc. 1971 IFIP, 1552-1558.

Barry, C. D., Ellis, R. A., Graesser, S. M., and Marshall, G. R. 1969. Display and Manipulation in Three Dimensions. Pertinent Concepts in Computer Graphics, Univ. of Ill. Press, 104-153.

Barry, C .D. and North, A. C. T. 1972. The Use of a Computer-Controlled Display System in the Study of Molecular Conformations. Cold Springs Harbor Symp. Quant. Biol., 36, 577.

Beem, K. M., Richardson, D. C., and Rajagopalan, K. V. 1977 Metal Site of Cu, Zn Superoxide Dismutase. Biochemistry, 16, #9 (3 May), 1930-1936.

Boring, E. G. 1942. Sensation and Perception in the History of Experimental Psychology. New York: Appleton-Century-Crafts, Inc.

Braunstein, M. L. 1976. Depth Perception Through Motion. New York: Academic Press, Chapter 6.

Britton, E. G. 1977. A Methodology for the Ergonomic Design of Interactive Computer Graphic Systems, and its Application to Crystallography. Ph. D. Dissertation, University of North Carolina, Chapel Hill, North Carolina.

Britton, E. G., Lipscomb, J. S., and Pique, M. E. 1978. Making Nested Rotations Convenient for the User. Proc. 1978 ACM SIGGRAPH Conf., Computer Graphics, 12, #3 (August), 222-227.

Brooks, F. P., Jr. 1977. The Computer "Scientist" as Toolsmith: Studies in Interactive Computer Graphics. Proc. 1977 IFIP, 625-634.

Brown, D. A. 1977. Military Uses Seen for Visual Simulators. Aviation Week and Space Technology, 107, #23 (5 December), 60-63.

Carter, C. W. 1977. X-ray Analysis of High Potential Iron-Sulfur Proteins and Ferrodoxins. Iron-Sulfur Proteins, Academic Press, Chapter 6.

Cohen, J. 1979. Three-dimensional Display of Parallel Sliced Images. M. S. Thesis, University of North Carolina, Chapel Hill, North Carolina.

Cohen, G. H. and Feldmann, R. J. 1974. MAP - An Interactive Graphics Computer Program for the Manipulation and Fitting of Protein Molecules to Electron Density Maps. Am. Crystallogr. Assoc. Spring D3, (Abstr.).

Collins, D. M., Cotton, F. A., Hazen, E. E., Jr., Meyer, E. F., Jr., and Morimoto, C. N. 1975. Protein Crystal Structures: Quicker, Cheaper Approaches. Science, 190, #4219 (12 December), 1047-1053.

Cornsweet, T. N. 1970. Visual Perception. New York: Academic Press.

Diamond, R. 1978. Bilder. A computer graphics program for bipolymers and its application to interpretation of structure of tobacco mosaic virus protein disks at 2-Å resolution. Proc. International Union of Pure and Applied Biochemistry: International Symposium on Biomolecular Structure, Conformation, Function, and Evolution. Madras, India, (4 January), Pergamon Press.

Eastman Kodak Company. 1950. Color as Seen and Photographed. Second Ed., p. 23.

Efron, R. 1957. Stereoscopic Vision. British Journal of Opthalmology, 41, 709-730.

Feldmann, R. J. 1976. AMSOM -- Atlas of Macromolecular Structure On Microfiche. Maryland: Tracor Jitco Inc.

Ferro, D. R., McQueen, J. E., Jr., McCown, J. T., and Hermans, J. (submitted). Energy Minimizations of Rubredoxin. Journal of Molecular Biology

Foley, J. D. and Wallace, V. L. 1974. The Art of Natural Graphic Man-Machine Conversation. Proc. IEEE, 62, #4 (April), 462-471.

Foley, J. D. and Wright, W. V. 1975. An Interactive Molecular Graphics System with a Satellite Terminal Closely Coupled to its Host. (Panel discussion position paper) Proc. ACM 1975 Annual Conference, 88-89.

Fritch, J. M., Ellis, R. A., Jacobi T. H., and Marshall, G. R. 1975. A Macromodular Graphics System for Protein Structure Research. Computers and Graphics, 1, #2/3 (September), 271-278.

Gibson, J. J. 1953. The Perception of the Visual World. Boston: Houghton Mifflin Co., Cambridge: The Riverside Press.

Girling, R. L., Houston, T. E., Schmidt, W. C., Jr., and Amma, E. L. (submitted). Macromolecular Structure Refinement by Restrained Least Squares and Interactive Graphics as Applied to Sickling Deer Type III Hemoglobin. (Abstract). ACTA Crystollographica.

Graham, C. H. (ed). 1965. Vision and Visual Perception. New York: John Wiley & Sons, Inc.

Halas, J. and Privett, B. 1958. How to Cartoon for Amateur Films. London and New York: Focal Press Ltd.

Hansen, W. J. 1971. User engineering principles for interactive systems. Proc. 1971 AFIPS FJCC, 39, 523-532.

Helmholtz, H. von. 1910. Treatise of Physiological Optics III. Southall, J. P. C. (ed). Optical Society of America, 1925, Menasha Wisconsin: George Banta Pub. Co.

Hobgood, W. S. 1969. A Three-Dimensional Computer Graphics Display Using a Varifocal Mirror. M. S. Thesis, University of North Carolina, Chapel Hill, North Carolina.

Hofstetter, H. W. 1956. Industrial Vision. Philidelphia: Chilton Company, 76-78.

Holbrook, S. R., Sussman, J. L., Warrant, R. W., Church, G. M., and Kim S. H. 1977. RNA-ligand Interactions: (I) Magnesiun Binding Sites in Yeast tRNA Phe. Nucleic Acids Research, 4, #8 (August), 2811-2820.

Johansson, G. 1950a. Configurations in Event Perception. Uppsala: Almqvist & Wiksells Boktryckeri Ab.

Johansson, G. 1950b. Configurations in the Perception of Velocity. Acta Psychologica, 7, 25-79.

Johnson, T. E. 1963. SKETCHPAD III: A Computer Program for Drawing in Three Dimensions. Proc. 1963 AFIPS SJCC, 23, 347-353.

Jones, T. A. 1978a. A Graphics Model Building and Refinement System for Macromolecules. Journal of Applied Crystallography, 11, 268-272.

Jones, T. A. 1978b. The RING [user manual]. Max-Planck-Institut fur Biochemie, 8033 Martinsried bei Muchen, Germany.

Joyce, J. D. and Ciancilo, M. J. 1967. REACTIVE DISPLAYS: Improving the Man-machine Graphical Communications. Proc. 1967 AFIPS FJCC, 31, 713-721.

Judd, D. B. 1943. Facts of Color-Blindness. Journal of the Optical Society of America, 33, 6 (June), 294-307.

Julesz, B. 1971. Foundations of Cyclopean Perception. Chicago and London: The University of Chicago Press, 270.

Kaufman, L. 1974. Sight and Mind, an Introduction to Visual Perception. New York: Oxford University Press, 67-76.

Kilpatrick, P. J. 1976. The Use of a Kinesthetic Supplement in an Interactive Graphics System. Ph. D. Dissertation, University of North Carolina, Chapel Hill, North Carolina.

Kim, S. H. and Sussman, J. L. 1976. Pi Turn is a Conformational Pattern in RNA Loops and Bends. Nature, 260, #5552 (15 April), 645-646.

Kimball, M. R., Sato, A., Richardson, J. S., Rosen, L. S., and Low, B. W. 1979. Molecular Conformation of Erabutoxin b; Atomic Coordinates at 2.5 Å Resolution. Biochemical and Biophysical Research Communications, 88, #3 (13 June), 950-959.

Kling, J. W. and Riggs, L. A., et al. 1971. Woodworth and Schlosberg's Experimental Psychology. New York: Holt, Rinehart, & Winston.

Kolers, P. A. 1972. Aspects of Motion Perception. Oxford, New York: Pergamon Press.

Kriloff, H. Z. 1976 Human Factor Considerations for Interactive Display Systems. User-oriented Design of Interactive Graphics Systems, ACM/SIGGRAPH Workshop, 14-15 October, Pittsburgh PA, 45-52.

Kuhn, H. S. 1950. Eyes and Industry. St. Louis: C. V. Mosby Company, 116, 124-125, 148.

Land, R. I. and Sutherland, I. E. 1969. Real-time Color, Stereo, Computer Displays. Applied Optics, 8, #3 (March), 721-723.

Leibowitz, H. W. and Owens, D. A. 1977. Nighttime Driving Accidents and Selective Visual Degradation. Science, 197, #4302 (29 July), 422-423.

Levelt, W. J. M. 1962. Motion Braking and the Perception of Causality. In: Michotte, A. et al. 1962. Causalité, Permanence, et Réalité, phénoménales. Louvain: Publications Universitaires, Paris: Béatrice-Nauwelaerts.

Levinthal, C. 1966. Molecular Model-building by Computer. Scientific American, 214, #6 (June), 42-52.

Levinthal, C., Macagno, E., and Tountas, C. 1974. Computer-Aided Reconstruction from Serial Sections. Federation of American Societies for Experimental Biology, Federation Proc., 33, #12 (December), 2336-2340.

Love, W. E., Fitzgerald, P. M. D., Hanson, J. C., and Royer W. E. (in press). Intermolecular Interactions in Crystals of Human Deoxy Hemoglobin A, C, F and S. Proc. The International Meeting on The Development of Therapeutic Agents for Sickle Cell Diseases. 19-20 July 1978.

Martin, J. 1973. Design of Man-Computer Dialogs. Englewood Cliffs, New Jersey. Prentice-Hall, Inc. Chapter 31.

Michotte, A. 1946. The Perception of Causality. Louvain: Publications Universitaires. Mace, C. A. (ed), Miles, T. R. and Miles, E. translation, London: Metheven & Co. Ltd. 1963.

Monaco, H. L., Crawford, J. L., and Lipscomb, W. N. 1978. Three-dimensional Structures of Aspartate Carbamoyltransferase from Escherichia coli and of its Complex with Cytidine Triphosphate. Proc. National Academy of Sciences USA, 75, #11 (November), 5276-5280.

Monty, R. A. and Senders, J. W. 1976. Eye Movements and Psychological Processes. New York, Toronto, London: John Wiley & Sons, Inc.

Newman, W. M. and Sproull, R. F. 1973. Principles of Interactive Computer Graphics. McGraw-Hill, Inc., 127, 235.

Ogle, K. N. 1963. Stereoscopic Depth Perception and Exposure Delay between Images to the Two Eyes. Journal of the Optical Society of America, 53, 11 (November), 1296-1304.

Ophir, D., Shepherd, B. J., and Spinard, R. J. 1969. Three-Dimensional Computer Display. CACM, 12, #6 (June), 309-310.

Ortony, A. 1971a. A System for Stereo Viewing. Computing Journal, 14, #2 (May), 140-144.

Also appears in: Conference on Displays, Institution of Electrical Engineers Conf. Pub. No. 80 (7-10 September 1971), C. Baldwin Ltd., 225-232.

Ortony, A. 1971b. Interactive Stereographics. Conference on Displays, Institution of Electrical Engineers Conf. Pub. No. 80 (7-10 September), C. Baldwin Ltd., 185-193.

Pearson, D. E. 1975. Transmission and Display of Pictorial Information. New York and Toronto: John Wiley & Sons, Inc., 56-60.

Perkins, W. J., Piper, E. A., Tattam, F. G., and White, J. C. 1971. Interactive Stereoscopic Computer Displays for Biomedical Research. Computers and Biomedical Research, 4, #3, 249-261.

Pique, M. E. (draft). Nested Dynamic Rotations for Computer Graphics. M. S. Thesis, University of North Carolina, Chapel Hill, North Carolina.

Plott, H. H., Jr., Irwin, J. D., and Pinson, L. J. 1975. A Real-Time Stereoscopic Small-Computer Graphics Display System. IEEE Transactions on Systems, Man, and Cybernetics, SMC-5, #5 (Summer), 527-533.

Richards F. M. 1968. The matching of physical models to three-dimensional electron density maps: a simple optical device. Journal of Molecular Biology, 37, #1 (14 October), 225-230.

Richardson, D. C. 1977. Three-dimensional Structure of Cu, Zn Superoxide Dismutase. Superoxide and Superoxide Dismutases. Edited by A. M. Michelson, J. M. McCord, and I. Fridovich. London, New York: Academic Press. 217-223.

Rosenberger, F. U., et al. 1976. Extracts from 1976 NIH Annual Report. Technical Memorandum No. 230, Computer Science Laboratory, Washington University, St. Louis, Missouri.

Ross, J. 1976. The Resources of Binocular Perception. Scientific American, 234, #3 (March), 80-86.

Rushton, W. A. H. 1975. Visual Pigments and Color Blindness. Scientific American, 232, #3 (March), 64-74.

Sackman, H. 1967. Computers, System Science, and Evolving Society, New York: John Wiley & Sons, Inc.

Schevitz, R. W., Podjarny, A. D., Krishnamachari, N., Hughes, J. J., Siglar, P. B., and Sussman, J. L. 1979. Crystal structure of a eukaryotic initiator tRNA. Nature, 278, #5700 (8 March), 188-190.

Schmerler, J. 1976. The Visual Perception of Accelerated Motion. Perception, 5, 167-185.

Sekuler, R. and Levinson, E. 1977. The Perception of Moving Targets. Scientific American, 236, #1 (January), 60-73.

Shephard, R. N. and Metzler, J. 1971. Mental Rotation of Three-Dimensional Objects. Science, 171, #3972 (19 February), 701-703.

Shurtleff, D. A. 1967. Studies in Television Legibility - A Review of the Literature. Information Display, January, 40-45.

Snell, A. C. 1945. Industrial Challenge to Opthamology. Proc. of a Conference on Industrial Opthamology. 7-11 May, Columbia University Press. 12-13.

Spigel, I. M. (ed). 1965. Readings in the Study of Visually Perceived Movement. New York: Harper & Row.

Stellman, S. D. 1975. Three-dimensional Interactive Graphics in X-ray Crystalographic Analysis. Computers and Graphics, #2/3 (September), 279-288.

Stenkamp, R. E., Sieker, L. C., Jensen, L. H., and McQueen, J. E., Jr. 1978. Structure of Methemerythrin at 2.8-Å Resolution: Computer Graphics Fit of an Averaged Electron Density Map. Biochemistry, 17, #13 (27 June), 2499-2504.

Stevens, S. S. 1957. On the psychophysical law. Psychol. Rev., 64, 153-181.

Strauss, C. M. and Poley, S. 1968. 3DPDP, a Three-Dimensional Piping Design Program. Proc. 1968 IFIP, 2, 1431-1440.

Sussman, J. L., Holbrook, S. R., Church, G. M., and Kim, S. H. 1977. A Structure-factor Least-squares Refinement Procedure for Macromolecular Structures using Constrained and Restrained Parameters. ACTA Crystollographica, A33, 800-804.

Sussman, J. L., Holbrook, S. R., Warrant, R. W., Church, G. M., and Kim, S. H. 1978. Crystal Structure of Yeast Phenylalanine Transfer RNA. I. Crystallographic Refinement. Journal of Molecular Biology, 123, #4 (25 August), 607-630.

Sussman, J. L. and Kim, S. H. 1976a. Idealized Atomic Coordinates of Yeast Phenylalanine Transfer RNA. Biochemical and Biophysical Research Communications, 68, #1, 89-96.

Sussman, J. L. and Kim, S. H. 1976b. Three-Dimensional Structure of a Transfer RNA Common in Two Crystal Forms. Science, 192, #4242 (28 May), 835-838.

Sutherland, I. E. 1963. SKETCHPAD: A Man-Machine Communication System. Proc. 1963 AFIPS SJCC, 23, 329-346.

Sutherland, I. E. 1968. A Head-Mounted Three-Dimensional Display. Proc. 1968 AFIPS FJCC, 757-764.

Szabo, N. S. 1978. Digital Image Anomalies: Static and Dynamic. Visual Simulation and Image Realism, Proc. 22nd Annual International Technical Symposium, Society of Photo-Optical Instrumentation Engineers. San Diego, Calif., (30-31 August), 162, 11-15.

Tiffin, J. 1942. Industrial Psychology. New York: Prentice Hall, Inc.

Tountas, C. and Katz, L. 1971. Interactive Graphics in Molecular Biology. Real-time Three-dimensional Rotations of Images and Image Fragments. Proc. Summer Computer Simulation Conf., 1, 241-247.

Traub, A. T. 1968. A New 3-Dimensional Display Technique, The Mitre Corp., Bedford, Mass.

Tsernoglou, D. and Petsko, G. A. 1977. Three-Dimensional Structure of Neurotoxin a from Venom of the Philippines Sea Snake. Proc. National Academy of Sciences USA, 74, #3 (March), 971-974.

Tsernoglou, D., Petsko, G. A., and Hudson, R. A. 1978. Structure and Function of Snake Venom Curarimimetric Neurotoxins. Molecular Pharmacology, 14, 710-716.

Tsernoglou, D., Petsko, G. A., McQueen, J. E., Jr., and Hermans, J. 1977. Molecular Graphics: Application to the Structure Determination of a Snake Venom Neurotoxin. Science, 197, #4311 (30 September), 1378-1381.

Tsernoglou, D., Petsko, G. A., and Tu, A. T. 1977. Protein Sequencing by Computer Graphics. Biochem. and Biophys. ACTA, 491, 605-608.

Vickers, D. L. 1974. Sorcerer's apprentice: Head mounted display and wand. Ph. D. Dissertation, University of Utah.

Voss, D. (in preparation). Three-dimensional Manipulation by a Two-dimensional Locator. M. S. Thesis, University of North Carolina, Chapel Hill, North Carolina.

Wallach, H. 1976. On Perception. New York: Quadrangle.

Wallach, H. and Karsh E. B. 1963a. Why the Modification of Stereoscopic Depth-Perception is so Rapid. American Journal of Psychology, 76, 413-420.

Also appears in: Wallach, H. 1976. On Perception. New York: Quadrangle, 278-288.

Wallach, H. and Karsh E. B. 1963b. The Modification of Stereoscopic Depth-Perception and the Kinetic Depth-Effect. American Journal of Psychology, 76, 429-435.

Also appears in: Wallach, H. 1976. On Perception. New York: Quadrangle, 289-297.

Wallach, H., Moore M. E., and Davidson L. 1963. Modification of Stereoscopic Depth-Perception. American Journal of Psychology, 76, 191-204.

Also appears in: Wallach, H. 1976. On Perception. New York: Quadrangle, 261-277.

Wallach, H. and O'Connell, D. N. 1953. The Kinetic Depth Effect. Journal of Experimental Psychology, 45, #4 (April), 205-217.

Also appears in: Wallach, H. 1976. On Perception. New York: Quadrangle, 35-56.

Wallach, H., O'Connell, D. N., and Neisser, U. 1953. The Memory Effect of Visual Perception of Three-dimensional Form. Journal of Experimental Psychology, 45, #5 (May), 360-368.

Also appears in: Wallach, H. 1976. On Perception. New York: Quadrangle, 170-183.

Warrant, R. W. and Kim, S. H. 1978. Alpha Helix-Double Helix Interaction: Structure of a Protamine-Transfer RNA Complex and Nucleoprotamine Model. Nature, 271, #5641 (12 January), 130-135.

Williams, L. Aliasing. 1979. Computer Animation Techniques. ACM/SIGGRAPH'79 Tutorial, 111-115.

Wipke, W. T. and Dyott, T. M. 1973. GIGL (General Interactive Graphics Library) User Manual. Department of Chemistry, Princeton University, (October), page 10-2.

Woodson, W. E. and Conover, D. W. 1964. Human Engineering Guide for Equipment Designers, Univ. of Calif. Press, Cambridge Univ. Press.

Wright, W. V. 1972a. An Interactive Computer Graphic System for Molecular Studies. Ph. D. Dissertation, University of North Carolina, Chapel Hill, North Carolina.

Wright, W. V. 1972b. The Two-Dimensional Interface of an Interactive System for Molecular Studies. ACM SIGPLAN notices, 7, #10 (October), 76-85.

INDEX

Appendix A

MAKING NESTED ROTATIONS CONVENIENT FOR THE USER

MAKING NESTED ROTATIONS
CONVENIENT FOR THE USER

Edward G. Britton
James S. Lipscomb
Michael E. Pique

University of North Carolina
at Chapel Hill

ABSTRACT

Subimage motion in a three-dimensional computer graphic system is much easier for the user to control if the subimage moves the same direction as his hand while he manipulates the control device. The implementation of such coordinated motion of hand and subimage implies modification of the normal procedure for calculating transformations. The convenience of such manipulation also depends on appropriate selection or design of the input device.

This paper reviews the relevant attributes of locator devices and presents an approach to selection. It presents the mathematics of transformation nesting and "compensation" to preserve motion-synchrony. Finally, it offers a case history of an interactive graphic system whose human factors were improved by these techniques.

Key Words and Phrases: three-dimensional computer graphics, coordinate transformations, nested rotation, ergonomics, human factors, analog input devices, man-machine interaction, kinesthetic correspondence, molecular graphics

CR Categories: 3.13, 8.2

1. INTRODUCTION

Many interactive computer graphic systems can present an object from different viewpoints and can move some parts of a picture with respect to others. The system designer decides which moving and viewing operations to provide the user. The designer should implement these operations in forms that facilitate the user's accomplishments. Since moving and viewing pictures *qua* pictures is the user's method, but rarely his goal, he should feel he is manipulating the objects rather than their images.

The user should command an interactive system by manipulating devices in ways related to his thoughts by kinesthesia or convention. Hansen (5) says to use "muscle memory -- make actions mechanical not thoughtful." The user should move images by moving knobs or joysticks, not pushing buttons or typing.

We suggest that when the user's hand moves a device, the image controlled by the device should move the same direction. We call this similarity of paths *kinesthetic correspondence*. In our laboratory we have observed that such correspondence increases user productivity.

2. MEANINGFUL SUBIMAGE MOTION

2.1 CONTROL DEVICE DESIGN

The application model's motions may be rotations or translations, and may be confined to certain directions or be unconstrained. Selecting a control device requires examination of available devices in light of the application's requirements.

Fig. 1 summarizes some important physical device properties that we now describe. Not all of the combinations of alternatives are possible. Woodson & Conover (14) discuss other properties important to good human factors.

The user operates a rotation device by turning it about one or more axes, and a translation device by moving or pushing it in a straight line. Each is best used to control its own type of motion, but when limited to small movements each may effectively masquerade as the other.

A bounded device has mechanical limits beyond which it may not move. All translation devices are bounded, but rotation devices may be bounded to less than 360 degrees, bounded to more than 360 degrees (as multiturn potentiometers are), or un-

----

This research was partially supported by NIH Biotechnology Research Resource grant #RR00898, NSF grant #GJ-34697, AEC contract #AT-(40-1)-3817, and the IBM Corporation.
Authors' present addresses: E.G. Britton, Defense Communications Engineering Center, 1860 Wiehle Avenue, Reston, VA 22090; J.S. Lipscomb and M.E. Pique, Department of Computer Science, University of North Carolina, Chapel Hill, NC 27514.

rotation	vs.	translation
bounded	vs.	unbounded
homogeneous	vs.	distinguished position
spring return	vs.	no spring return

Figure 1: Physical Device Properties

bounded. Bounded devices can effectively control unbounded motions if a clutch, when disengaged, logically disconnects the device and allows it to be moved back from its bound without moving the image.

A homogeneous device presents no indication of its current position to the user, and has no positions with special significance. Unbounded devices are homogeneous unless they have a position-indicating pointer. Bounded devices have distinguished positions at their bounds.

A spring-return device has a distinguished position to which it returns when released. Such a device is generally unsuitable for direct positioning, but works well to control a motion rate. Rate control benefits from super-linear response (e.g., x * ABS(x)), and a dead band that remedies mechanical inaccuracy in the return: values near zero are considered zero.

2.2 AXIS COMPENSATION

When nested transformations contain rotations, the control algorithm may need to perform axis compensation to allow kinesthetic correspondence. This compensation is a rotation transforming the movement from the user's reference space into the rotated space. For example, imagine the user has a spring-return two-axis joystick that controls the rate of translational motion of a small square in a picture, and a knob that rotates the picture in the plane of the screen. The square is nested within the rotation: when the knob is turned, the square follows as the picture turns. To achieve kinesthetic correspondence, the two-axis joystick should move the square in the direction the joystick is pushed, regardless of the picture's orientation.

Such kinesthetic correspondence is achieved by rotating the vector representing the joystick motion backwards by the over-all orientation angle and adding the result to the square's position (described in the rotating space).

3. ALGORITHM DESIGN

Accomplishing meaningful subimage motion in three dimensions requires the ability to construct a rotation about an arbitrary axis, modify that rotation properly, and nest the resulting transformations.

Every nesting of transformations may be described as a tree, with the outermost space as the root and nested spaces as the other nodes. We consider the outermost space an unmoving reference; every other space has a father space immediately above it in the tree. The arc between two nodes represents the relationship between the spaces; we describe this relation in the father space's coordinate system. All picture elements of a subimage are transformed by the same transformation.

3.1 ROTATIONS ABOUT ARBITRARY AXIS VECTORS

A 3x3 matrix can describe any rotation about an axis through the origin of the coordinate system. As in Newman & Sproull (9) we use a right-handed coordinate system, treating vectors (underlined) as row vectors and post-multiplying them by rotation matrices, i.e.,

VECTOR = (VECTOR) (MATRIX).

If scalars x, y, and z are components of a unit-vector along the specified rotation axis vector, VECTOR, whose length, θ, is the amount of right-hand rotation of the image (counterclockwise when the axis of rotation points at the observer) with respect to its father space, then the matrix function describing this rotation, ROTN_ABOUT_AXIS(VECTOR), is

$$\begin{bmatrix} x^2+c(1-x^2) & xy-cxy+sz & xz-cxz-sy \\ xy-cxy-sz & y^2+c(1-y^2) & yz-cyz+sx \\ xz-cxz+sy & yz-cyz-sx & z^2+c(1-z^2) \end{bmatrix}$$

where s = sin θ, and c = cos θ. Derivations of this may be found in Rogers & Adams (11) and Konopinski (7). To follow the convention that clockwise rotations are positive, transpose the matrix given here. The algebraic arrangement shown (Barry et al. (1)) is suited for fixed-point computing because it avoids intermediate terms exceeding unity in magnitude.

3.2 ROTATIONS AS OPERATORS

A rotation matrix can either specify the orientation of one coordinate space with respect to another, or describe a change in orientation of a single space. In the latter case the new orientation is produced by multiplying the old orientation matrix by the change matrix:

new orient. = (old orient.) (change matrix).

This allows the modification of an orientation by any amount around any axis through its origin.

For dynamic motion one can iterate the process over time. Computational error, however, will soon cause the image to skew into a pancake shape and eventually shrink out of sight. Tountas & Katz (12) have proposed a renormalization of the new orientation matrix to eliminate this problem. In our laboratory we control it by a Gram-Schmidt orthonormalization of one pair of columns with each iteration.

3.3 TRANSFORMATION NESTING

The following mnemonic notation for designing transformations identifies a transformation by a capital letter, and indicates the two spaces that it relates by small letters. For example, transformation 'S' between spaces 'c' and 'b' is cSb. This represents a single relationship comprising the orientation of space 'c' with respect to space 'b', cRb, and the origin of 'c' in 'b', <u>cTb</u>.

In the usual process of drawing the picture only one transformation is applied to each picture element, the transformation relating the picture element to the reference space. If a space 'c' is nested in another space 'b', then when 'b' is rotated, the origin and orientation of 'c' with respect to the reference space must be changed to maintain the nesting. In this discussion we assume the nested space pivots independently only about its own origin.

The following expressions nest one space, 'c', inside another, 'b'; that is, given the relationship cSb between 'c' and 'b', and bSa between 'b' and reference space 'a', they compute the relationship cSa between 'c' and 'a'.

1. The orientation of space 'c' with respect to space 'a' is a 3x3 matrix product:

   cRa = (cRb) (bRa).

2. The origin of 'c' with respect to 'a' is a 1x3 matrix product followed by a vector sum:

   <u>cTa</u> = (<u>cTb</u>) (bRa) + <u>bTa</u>.

We use the symbol '<>' for this nesting algorithm:

cSa = cSb <> bSa.

The small letters between the two transformations must match for the algorithm to nest the spaces.

Multiple levels of nesting require serial applications of the algorithm. For example, to nest a space 'd' in a space 'c', and 'c' in 'b', the algorithm would be applied twice:

dSa = dSc <> cSb <> bSa.

Associativity permits evaluation of this to produce an intermediate result of either space 'd' with respect to space 'b', or 'c' with respect to 'a'.

Using two small letters per transformation, any nesting may be designed keeping in mind that the innermost space appears on the left of the list and the outermost on the right.

If vectors are treated as column vectors and pre-multiplied by the matrices, the direction of nesting is reversed and the subscript letters for each transformation must be interchanged to preserve the adjacent-letter guideline.

3.4 TRANSFORMATION INVERSES

We can write the inverse of cSb, $(cSb)^{-1}$, as bSc, which describes the orientation and origin of 'b' with respect to 'c'. The rotation and translation components of $(cSb)^{-1}$ are:

$bRc = (cRb)^{-1}$ = transpose(cRb)

<u>bTc</u> = $(\underline{cTb})^{-1} = -(\underline{cTb})\ (bRc)$.

3.5 CONTROL IN ROTATED SPACES

To provide kinesthetic correspondence in unconstrained subimage rotation and translation the algorithm must compensate for the nesting when it changes the relation between a space and its father.

Suppose space 'c' is nested in space 'b' but rotates about an axis defined in a third space 'd'. We want to find the transformation relating 'c' to 'b' after the rotation. Let '<u>ROTNd</u>' be the axis vector defined in space 'd' about which 'c' is to be rotated and whose length is the amount of rotation. This vector can be rotated into space 'b' by

<u>ROTNb</u> = <u>ROTNd</u> dRa $(bRa)^{-1}$
= <u>ROTNd</u> dRa aRb
= <u>ROTNd</u> dRb.

The result is an equivalent vector defined with respect to space 'b', the father space of the one to be rotated.

The matrix describing the new orientation of 'c' is

(new cRb) = (old cRb) (ROTN_ABOUT_AXIS(<u>ROTNb</u>)).

Similarly, a translation defined in space 'd', <u>TRANd</u>, can translate space 'c' by:

<u>TRANb</u> = <u>TRANd</u> (dRb),

(new <u>cTb</u>) = (old <u>cTb</u>) + <u>TRANb</u>.

4. HUMAN FACTORS CASE HISTORY

We installed controls providing kinesthetic correspondence in an interactive graphic system for biochemical crystallography and made it much easier to use. The system, GRIP-75, is an extension of Wright's GRIP-71 (15) system.

The analysis of the crystallographic process that guided the GRIP-75 design is in Britton (2), and comments on the system's philosophy and results are in Brooks (3). More technical notes and implementation details will be in Lipscomb (8) and Pique (10). Tsernoglou et al. (13) include a short description of the system architecture written by four of its users.

Crystallographers use our system to determine the shapes of large molecules. They fit a three-dimensional stick-figure molecule into a contour map of its electron density, whose shape they have found through X-ray diffraction. They fit the molecule one small substructure at a time, by translating and rotating each as a rigid unit and twisting its internal bonds.

4.1 VIEWING CONTROL DESIGN

The user needs dynamic viewing rotation for three-dimensional perception and for selecting views that reveal how he should fit the substructure. The molecular structure and density map must be viewable from any direction, but not necessarily in any orientation, and this requires two degrees of freedom.

4.1.1 Viewing angle control

We require a mechanical device with two angular degrees of freedom that will hold its position when released. We provide kinesthetic correspondence by mounting the device's two rotation sensors so they are nested the same way as the image rotations they control.

Viewing rotation seems to be our single most effective depth cue, even better than stereoscopy. Fig. 2 shows our viewing control. The horizontal rotation angle is nested in the vertical by the nesting algorithm in section 3.3, causing the axis of the angle controlled by the horizontal potentiometer to be rotated according to the angle controlled by the vertical one. This works exceedingly well; the handle seems to be attached to the image itself, and viewing rotations are carried out easily and nonconsciously. Because the kinesthetic correspondence is strong, the perception of depth is considerably greater for the individual operating the device than for onlookers.

4.1.2 Viewing angle indication

Since our users occasionally want to know how much they have turned the picture, in a corner of the screen we display the application coordinate axes as a small gnomon, Fig. 3, which turns about its center as the picture rotates. Merely indicating the viewing angle to the user, however, would not substitute for kinesthetic correspondence between viewing angle and control device. The user's glancing over at the indicator would impede the work by disturbing visual continuity (Foley & Wallace (4)).

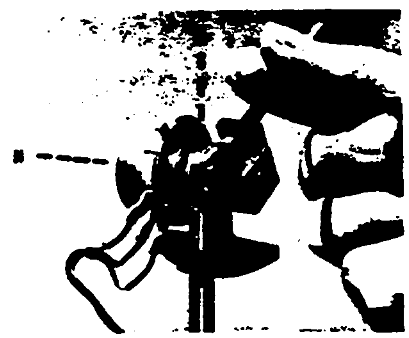

Figure 2: Viewing Rotation Control Device (Vertical potentiometer not shown)

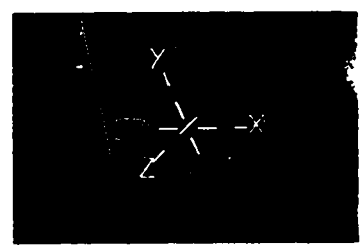

Figure 3: Viewing Angle Indicator Gnomon

4.2 UNCONSTRAINED MANIPULATION DESIGN

The unconstrained translation and rotation of the selected substructure requires six degrees of freedom; we assigned the control functions to separate translation and rotation devices. All manipulations are nested in the viewing rotation space, so the substructure moves with the viewing control. To provide kinesthetic correspondence, however, the unconstrained manipulation controls are compensated to act in the user's space.

We use a non-spring-return device to control translation of the substructure: a box with three orthogonal slide potentiometers attached to a single control handle that moves within a cube 0.2 meters on a side. This device was given to us by Bell Telephone Laboratories through the courtesy of A. M. Noll. Through software a clutch (section 2.1) permits movements beyond the box's bounds. The user simply slides the control handle in the direction he wishes the substructure to move; the kinesthetic correspondence between the translation control and the substructure's motion allows easy positioning of the substructure.

We now believe a spring-return rate control would be more convenient because it would not require a clutch and the handle would always be in the same place when the user reached for it. Foley & Wallace (4) mention the difficulties with devices whose locations change. The user's having to look for the handle of a large joystick distracts him as much as searching for a lightpen.

We use a spring-return device to control the substructure's rate of rotation. This has worked well. The joystick used has three degrees of freedom: push, lean, and twist. Through software we provide super-linear response and a dead band.

4.3 CONSTRAINED MANIPULATION DESIGN

The constrained manipulations (twists of chemical bonds) required by our users are multiple concurrent single-axis rotations nested in the unconstrained manipulations. Some of these rotations are nested in others, depending on the molecular model. Mechanical difficulties seem to preclude complete kinesthetic correspondence, because the constraint axes rotate with the viewing control. We use a set of unconnected single-axis knobs.

Our analysis indicated that the user would often want to leave some of the angles in the substructure unchanged during each fitting operation. We therefore provide self-zeroing knobs whose setting is taken to be zero each time a fitting is begun. We use knobs with a 300-degree range, and scale their response so they can rotate each subimage through somewhat more than 360 degrees, a reduction rarely noticed. This is consistent with Kilpatrick's experience (6) with an enlarged kinesthetic space for translations. We suggest that although motions on the screen should correspond in kind to those of the control device (e.g., provide rotary control for rotary response) they need not in magnitude.

4.4 IMPLEMENTATION NOTES

Our program traverses a linked list to update the transformation describing the orientation and location of each dynamically-controlled subimage. From control device values the relation of each space to its father is either calculated afresh or modified. Then its relation to the reference space is computed and encoded into transformation orders for the display unit. The list of transformation spaces is processed in the preorder traversal of the nesting tree to ensure that no space is calculated before its father.

A schema of a typical nesting is:

(b2a, c3a) <> a1m <> mMv <> vHh <> hVu

where 'V' represents viewing vertical rotation, 'H' represents viewing horizontal rotation, 'M' represents unconstrained manipulation, and '1', '2', and '3' are three constrained manipulations ('2' and '3' nested in '1'). Outermost space 'u' ('user') is the unmoving reference space.

We now consider calculation of the unconstrained manipulation node as an example.

Each time the user selects a substructure for fitting, the following operations are performed exactly once (':=' indicates assignment):

1. vTu := offset of origin of viewing space from center of screen, a constant.
2. mTv := initial relation between manipulated substructure origin (pivot) and viewing pivot. Initially the substructure is placed at its pre-fitting position.
3. mRv := Identity matrix. Initially there is no difference between the substructure's orientation and that of the surrounding structure.

At each update, the node uses information prepared by higher nodes:

1. Manipulation control vectors TRAN DEVICEu and ROTN DEVICEu defined in the user's space.
2. The viewing orientation space with respect to the user, vRu.

At each update, the unconstrained manipulation node is processed:

1. uRv := vRu^{-1} = transpose(vRu)
2. TRAN DEVICEv := TRAN DEVICEu uRv
   ROTN DEVICEv := ROTN DEVICEu uRv

   These rotations compensate the user's motions into the viewing space, the manipulation space's father (see section 3.5).
3. mTv := mTv + TRAN DEVICEv

   This adds the rotated translation device motion to the manipulation translation vector.
4. CHANGEv := ROTN_ABOUT_AXIS(ROTN DEVICEv)
   mRv := orthonormalize(mRv CHANGEv)

   This rotates the manipulation space about the user-specified axis.
5. mTu := mTv vRu + vTu
   mRu := mRv vRu

   These nest the manipulation in the viewing rotation, making the subimage follow the motion of the viewing rotation control device (see section 3.3).

Some steps are skipped if a device has not moved since the previous update.

4.5 EXPERIENCE

In the two and a half years that GRIP-75 has been in production 27 chemists from 15 institutions have logged over 1700 hours of work on 16 molecules, with 11 published papers so far.

Users generally agree that kinesthetic correspondence eases their work, and several of them have reported difficulty using dynamic manipulations in comparable systems that lacked kinesthetic correspondence.

5. CONCLUSIONS

We believe that careful control-device selection and kinesthetic correspondence increase user productivity. We base this upon personal observation rather than controlled experiment. Nevertheless, we maintain that the benefits of kinesthetic correspondence justify its low computational cost.

6. ACKNOWLEDGEMENTS

We appreciate the advice and guidance given us over the past several years by Frederick P. Brooks, Jr., and William V. Wright. We are grateful to James D. Foley and Victor L. Wallace for stimulating our work in human factors. We thank P. Jerome Kilpatrick for sharing with us his ideas about rotations. We thank S. Bellovin, P. Calingaert, G. Frank, and D. Tolle for their criticism of this paper in draft.

7. REFERENCES

1. Barry, C.D., Ellis, R.A., Graesser, S.M., Marshall, G.R. CHEMAST: a computer program for modelling molecular structures, Proc. IFIP 1971, 1552-1558.

2. Britton, E.G. A Methodology for the Ergonomic Design of Interactive Computer Graphic Systems, and its Application to Crystallography. Ph.D. dissertation, Department of Computer Science, University of North Carolina, Chapel Hill, N.C., 1977.

3. Brooks, F.P., Jr. The computer "scientist" as toolsmith: studies in interactive computer graphics. Proc. IFIP 1977, 625-634.

4. Foley, J.D., and Wallace, V.L. The art of natural graphic man-machine conversation. Proc. IEEE 62, 4 (April 1974), 462-471.

5. Hansen, W.J. User engineering principles for interactive systems. Proc. AFIPS FJCC, 39, 1971, 523-532.

6. Kilpatrick, P.J. The Use of a Kinesthetic Supplement in an Interactive Computer Graphics System. Ph.D. dissertation, Department of Computer Science, University of North Carolina, Chapel Hill, N.C., 1976.

7. Konopinski, E.J. Classical Descriptions of Motion. W.H. Freeman, San Francisco, 1969.

8. Lipscomb, J.S. Three-dimensional Display of Molecular Models. M.S. thesis in progress, Department of Computer Science, University of North Carolina, Chapel Hill, N.C.

9. Newman, W.M., and Sproull, R.F. Principles of Interactive Computer Graphics. McGraw-Hill, New York, 1973.

10. Pique, M.E. Nested Dynamic Rotations for Computer Graphics. M.S. thesis in progress, Department of Computer Science, University of North Carolina, Chapel Hill, N.C.

11. Rogers, D.F., and Adams, J.A. Mathematical Elements for Computer Graphics. McGraw-Hill, New York, 1976.

12. Tountas, C., and Katz, L. Interactive graphics in molecular biology: real-time three-dimensional rotations of images and image fragments. Proc. Summer Computer Simulation Conf. 1: session II-2.7 (1971), 241-247.

13. Tsernoglou, D., Petsko, G.A., McQueen, J.E., and Hermans, J. Molecular graphics: application to the structure determination of a snake venom neurotoxin. Science, 197, 4311 (30 September 1977), 1378-1381.

14. Woodson, W.E. and Conover, D.W. Human Engineering Guide for Equipment Designers, Univ. of Calif. Press, Cambridge Univ. Press, 1964.

15. Wright, W.V. An Interactive Computer Graphic System for Molecular Studies. Ph.D. dissertation, Department of Computer Science, University of North Carolina, Chapel Hill, North Carolina, 1972.

PLEASE NOTE:

Appendix B, "Computer Program Listing for Three-Dimensional Cues for a Molecular Computer Graphics System" not microfilmed due to size. Available for consultation at The University of North Carolina at Chapel Hill Library.