VISUALIZING INFORMATION NETS IN THREE DIMENSIONS

by

Colin Ware and Glenn Franck

TR94-082 February 1994

Faculty of Computer Science
University of New Brunswick
P.O. Box 4400
Fredericton, N.B.
Canada E3B 5A3

Phone: (506) 453-4566
Fax: (506) 453-3566

Visualizing Information Nets in Three Dimensions

Colin Ware and Glenn Franck
Faculty of Computer Science
University of New Brunswick
P.O. Box 4400, Fredericton, NB.
E3B 5A3
cware@UNB.ca

Abstract

At present it is not clear if it is worth presenting abstract data in 3D, although there is a small body of evidence to support the idea. This paper reviews the evidence and reports on two new experiments that provide some compelling data showing that head coupled perspective transformations combined with stereo viewing can dramatically increase the size of the information structures that can be perceived. The first experiment was designed to provide the first quantitative measurements of how much more (or less) can be understood in 3D than in 2D. The 3D display used was configured so that the image on the monitor was coupled to the user's actual eye positions (and it was updated in real-time as the user moved) as well as being in stereo. Thus the effect was like a local "virtual reality" display located in the vicinity of the computer monitor. The results from this study show that true 3D viewing can increase the size of the graph that can be understood by a factor of three; using stereo or head coupling alone produce lesser advantages. The second experiment looked at a variety of viewing methods, both 2D and 3D, with and without stereo, and with and without motion. The results show that structured 3D motion and stereo viewing both help in understanding, but that the kind of motion is not particularly important; hand guided motion, head coupling (as in virtual reality displays) and automatic rotation all improve performance. These results provide strong reasons for using advanced 3D graphics for interacting with a large variety of information structures.

1. INTRODUCTION

A useful and interesting method for examining three dimensional structures is to couple a perspective stereoscopic view of a 3D scene to the user's eye positions and update the view in real-time as the user moves. The key elements of this are a high resolution monitor capable of running at a high frame rate, stereo glasses and some method for tracking the user's head position (see Figure 1) [5]. The position of the user's two eyes are computed from the head position and separate images are generated showing the correct perspective view of a set of virtual objects somewhere in the vicinity of the monitor screen. The result is a localized "Virtual Reality" (VR) environment which has a number of advantages over the much talked about immersion virtual reality [1, 5], not the least of these being that the everyday workspace of desk, filing cabinet, co-workers and coffee mug are not excluded. In our previous work we have called this "Fish Tank VR" to characterize its localized nature and distinguish it from the full immersion kind [19].

The advantages of Fish Tank VR representations of conventional objects are fairly obvious. For example, an industrial designer can see her design of telephone handset as a truly three dimensional object, or a trainee mechanic can see the engine part and how it is inserted as if he were looking at the real thing. But what about abstract information, such as data base schemas, networks of human relationships, the structure of object-oriented code or hypertext links? It is by no means clear that 3D diagrams offer any advantages over 2D ones in representing this kind of information. The present paper provides some strong evidence that 3D diagrams can offer significant advantages under the right viewing conditions. But before presenting this evidence we review some of the evidence which currently bears on the issue of whether advanced 3D viewing techniques should be used for viewing information networks.

1.1 How many dimensions are there in visual space?

In the following discussion we use the term visual space in the everyday sense of perceived area or volume, ignoring the dimensions of, for example, color and texture. According to a naive view, moving from a 2D to a 3D display should vastly increase the amount of information that can be represented. Consider a 1000x1000 computer display. On a line we can perceive 1000 distinct pixels, on the plane we can perceive 1000^2 distinct pixels. Extending this logic we should be able to display 1000^3 distinct voxels in a 3D volume. These relations can be succinctly expressed by the following equations

$$I_{2D} = I_{1D}^2 \qquad I_{3D} = I_{2D}^{3/2}$$

where I_{nD} represent the information that can be perceived given an n-dimensional display.

Clearly there is a flaw in this logic; in general we do not perceive volumes of data; we do not perceive details of the insides of solids, only the layout of surfaces in space. We only have two eyes which in the best case merely doubles the amount of information available, and we know that such effects as binocular rivalry mean that truly independent images from the two eyes cannot be perceived [11]. In fact it is only possible to extract depth information from highly correlated information presented to the two eyes. This alternative pessimistic view can be expressed by the equation

$$I_{3D} = C*I_{2D}$$

where C is some constant < 2.0; possibly a value of 1.05 might be reasonable.

This second view suggests that moving from 2D to 3D will only yield a 5% benefit in visualization, at least for the understanding of abstract data. If this view is correct it will be hardly worth while using 3D representations of abstract multidimensional networks of information; a 2D view is likely to be equally effective and much simpler to produce.

There is a third view based on the ecological argument that because we have evolved in a 3D world information presented in 3D will be processed more easily by the visual system. Networks of information do not have an inherent dimensionality in the Cartesian sense but if our brains prefer 3D layouts then a 3D layout may be more effective in conveying the information [15]. We should also allow that the brain is capable of integrating information from the sequence of views of the word that are obtained as we move about and that this considerable increases the amount of information obtainable, although as with stereopsis these views must be highly correlated.

A truly 3D display is not simply a perspective picture, or a stereo perspective picture. It requires that the image be accurate from the actual viewpoint of the observer, especially if the observer is moving. A large part of space perception comes from the relative movement of objects; in fact, this seems to be considerably more important than stereopsis in helping us interpret spatial layout, whether it be from head motion, resulting in motion parallax [9], or from object rotation giving the kinetic depth effect [2,3]. This movement allows the brain to integrate spatial information over time, and whereas a stereo display only gives two views to help understand a scene, a scene in which there is relative movement of the head and objects provides a whole continuum.

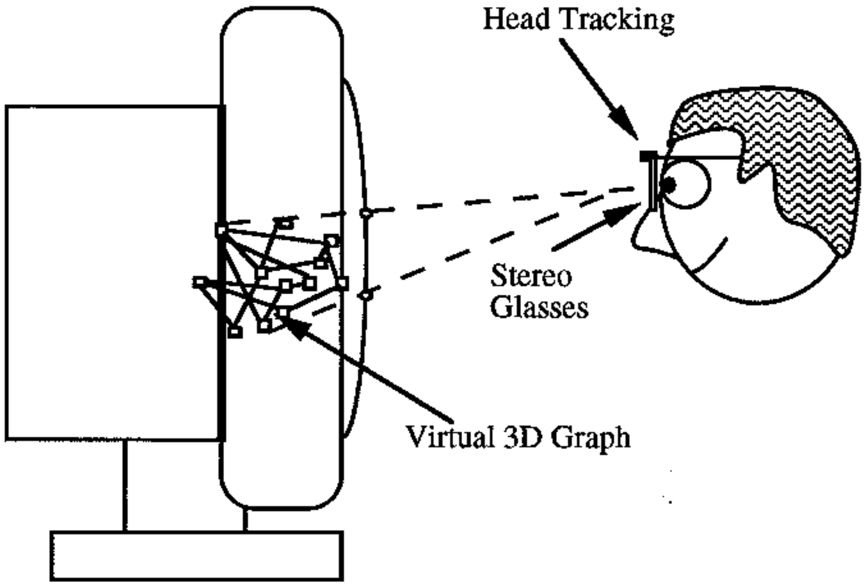

Figure 1. The preferred method for visualizing complex networks of information may be using head coupled stereo views. This results in a virtual 3D view of the graph placed in the vicinity of the monitor.

1.2 Anecdotal information

There is growing evidence that representing diagrams in 3D can allow more complex information to be comprehended. Perhaps the most influential work has been the SemNet project [7]. This used 3D representation to allow users to visualize large knowledge bases as nodes and arcs in a three dimensional space. No attempt was made to quantify the advantages of this mode of viewing over 2D layouts but much of the current interest in 3D abstract data visualization can be traced to this seminal study.

The Cone Tree technique developed by Robertson et al. [16] appears to have the capacity to show considerably more tree structured data than do more conventional 2D layouts. The Cone Tree system allows for the display of tree structured graphs displaying all the children of a node in the form of a cone of information. The authors claim that as many as one thousand nodes may be displayable using Cone Trees without visual clutter - this is clearly more than could be contained in a 2D layout, although the Cone Trees require certain user manipulations to access some of the information.

1.3 Empirical evidence

Of particular relevance to the display of information networks is work that has shown that the number of errors in detecting paths through tree structures is substantially reduced if a 3D display method is used [18,19]. Sollenberger and Milgram showed that both scene rotation and stereopsis helped reduce errors in a path tracing task [17,18]. In their motion conditions the stimulus pattern rocked back and forth about a vertical axis (the perspective imagery was not coupled to head position). They found that motion was more valuable than stereopsis in reducing errors. Ware et al's experiment used a similar task with a head coupled stereo display in which the perspective view was coupled to the measured eye position of the observer [19]. Although in Ware et al's display the motion was caused by head movement, the results were similar to those obtained previously by Sollenberger and Milgram. Since the geometric transformation of the image which may occur under head coupling or scene rotation may be very similar there is every reason to suppose that the same perceptual mechanisms are involved in both cases in constructing a 3D interpretation of the world.

A simple model for the integration of different depth cues is a weighted additive model [3], according to which the brain computes a weighted sum of different information in determining the depth that will be perceived. While this model is undoubtedly simplistic and the link to percent correct in a path tracing task is not explicit, Sollenberger and Milgram found their data suggested a superadditive model, however the relationship between depth perception and errors on their path tracing task is not clear. Conversely, the data of Ware et al is sub additive, although the same reservations about the link between the model and the data apply. Also, their experiment resulted in much lower error rates, and ceiling effects were undoubtedly present.

1.4 Motivation for new studies

There are a number of unanswered questions posed by the above studies. The first and most fundamental question is how much is gained by moving from a 2D to a 3D representation. While a completely general answer to this question can never be expected because the answer must to some extent be task specific, any answer derived from an abstract task such as path tracing in networks of nodes and arcs will generalize to the large set of problems that can be represented in this way. There is also the more profound question related to the issue of how depth information from different cues is integrated in the brain. By measuring task performance with a series of graphs of different sizes under different viewing conditions it should be possible to determine the function relating error rate to graph size for each condition. If these functions appear to belong to the same family we can model their relationship without making assumptions about such things as the additivity of error rates.

A second question is whether the results of Sollenberger and Milgram and Ware et al generalize to arbitrary graphs other than trees. This is important because tree layout is a relatively simple, well understood process and it is trivial to lay out a pair of trees in a plane so that they do not overlap. Hence the visualization problem posed by the previous studies could easily be solved without resorting to a 3D display. This is not the case for an arbitrary graph where the layout problem is more difficult and the advantages of 3D visualization may be more pronounced. While much of the work on 2D layout of directed graphs has been directed at minimizing arc crossings, the algorithms are often complex and in some cases are NP hard [6,8]. We hypothesize that 3D visualization will to some extent reduce the graph crossing problem because arcs will no longer appear in the plane of the screen.

A third question is what kind of motion is better for perceiving structure in information networks:

a) motion induced by perspective coupled to eye position,
b) automatic rotation of the object, or
c) motion caused by linking the user's hand movements to the object.

Both Sollenberger and Milgram and Ware et al. showed advantages of motion in information perception but the kind of motion was different in the two cases: scene rotation in the former case (causing the kinetic depth effect) and head motion in the latter (causing motion parallax). It should be noted that lateral motion of the viewpoint produces a very similar transformation to rotation of the scene about an axis vertically through the center of the scene. This leads to the hypothesis that both kinds of structured transformation may contribute to 3D space perception in similar ways to enhance network comprehension. However, it might be that the active search allowed by viewing modes a) and c) are critical to understanding since under these conditions the user has control over the view of the scene.

Experiment 1a was designed to address the first and second of the above questions, while Experiment 1b was designed to supplement 1a which yielded only an incomplete answer. Experiment 2 was designed to address the third question.

2. EXPERIMENT 1a: HOW MUCH BETTER IS A 3D/STEREO DISPLAY?

The purpose of this experiment was to provide a quantitative estimate of the improvements in graph understanding offered by adding head coupling and stereo disparity cues to the display. The answer we were looking for was of the form "adding head coupled stereo increases the size of a graph that can be understood by a factor of X with respect to error rates". In order to make such an evaluation we designed an experiment with a range of graph sizes. Two viewing conditions were used.

1) 2D: no stereo, no rotation; the 3D graph was projected onto a 2D plane using an orthographic (parallel) projection by removing Z axis information, hence no overlap information was presented.

2) Stereo, head coupled perspective: The correct view perspective was generated for each eye position (continuously updated) using an apparatus as shown in Figure 1. The perspective was continuously updated based on the measured head position (and derived eye position) of the subject.

2.1 Hardware

The equipment used in this experiment consisted of StereoGraphics Corporation's CrystalEyes 3D LCD shutter glasses to provide the stereo, with Logitech Corporation's ultrasonic head-tracking built into the frame of the glasses. The computer used to produce the display was a Silicon Graphics Crimson VGX. In stereo mode this system uses the top and bottom half of the frame buffer for the left and right images respectively. These images are vertically expanded by repeating lines at display time which results in half the vertical resolution and also causes lines to be drawn double thickness in a vertical direction.

2.2 Experimental Procedure

The subject's task was to decide whether there was a path of length two connecting two nodes which were highlighted in a randomly laid out graph. On each experimental trial there was either a path of length two (with an intervening node between the two highlighted nodes) or no path, with a 50% probability of each occurring.

The computer generated a random, 3D graph consisting of different numbers of nodes, arranged in a simulated 17 cm^3 volume. An example is shown in Figure 2. The nodes were divided into three equal-sized groups. Two of these groups were leaf nodes, while the third was a group of intermediate nodes. Each node in one of the leaf groups was connected via arcs to exactly two different nodes in the intermediate group. For *n* nodes, this produced a total of (4/3*n) connecting arcs. All nodes were placed randomly within the working volume.

Unhighlighted nodes were drawn in a dark gray color, while the highlighted ones were bright red; lighting was applied to all nodes to emphasize the three-dimensional effect. The size of each node was set to be 0.4 centimeters on each side. The arcs were white and drawn with double-width (two pixel) lines. Since the arcs were drawn as lines and not polygons, perspective sizing did not affect the apparent size of the lines and, due to the artifact associated with stereo display, the vertical thickness (1.1mm) was double the horizontal thickness (0.53 mm). The background was a flat light gray. These colors were chosen to minimize ghosting effects associated with stereo.

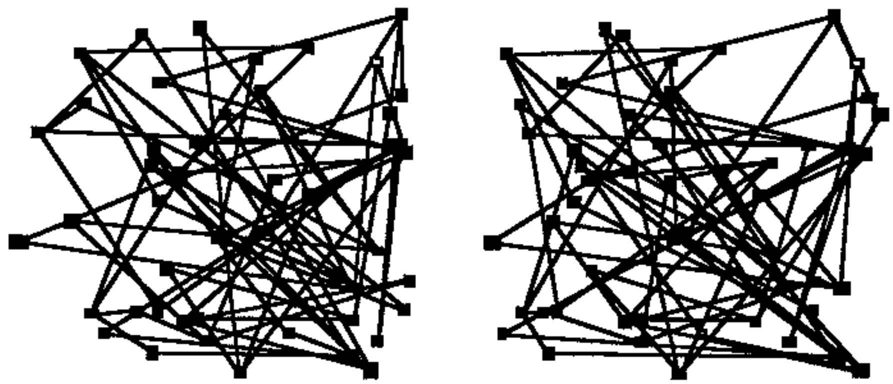

Figure 2. The stimulus patterns consisted of a set of nodes and arcs laid out in a volume and rendered according to a variety of different methods, either with or without stereo and with or without head coupled perspective. This stereo pair is not a screen image (which lost too much detail because of halftoning), instead it is a line drawing simulation of the screen display with 51 nodes and 68 arcs. The images are set up for crossed eye viewing. If they are inverted they can be viewed wall eyed.

Five graph sizes were used for each of the two viewing conditions.

Graph	#nodes	#arcs
1	24	32
2	51	68
3	78	104
4	105	140
5	132	176

This resulted in 10 size/condition combinations. Each of these was examined using two separate blocks of 20 trials for each subject giving a total of 40 trials per condition. The graph remained the same throughout a trial block, but for each trial a different pair of nodes were highlighted. The highlighted nodes were actually connected on 50% of the trials, randomly determined.

For each subject the experiment was divided into two experimental sessions administered on separated days with a complete set of conditions given on a pseudo-random order on each day. Before beginning the experiment each day, the subject was given a short warm-up which presented two trials under each of the conditions, but only using three of the five complexities (low, middle and high) for each condition.

Prior to each block of trials subjects were told which experimental condition to expect. On each trial subjects were given as much time as required to respond which they did by pressing a specified mouse button. The response time and the response validity was recorded.

Due to the different graph sizes, and the relatively high numbers of polygons to draw in the more complex scenes, the update rate was reduced to 30 updates per second for all conditions. That is, the scene was re drawn into the frame buffer 30 times per second. However, the video refresh rate was 120 Hz (60 Hz for each eye). The subject wore the stereo glasses regardless of whether the condition required them in order to avoid effects related solely to the reduced luminance caused by the glasses.

There were a total of eleven subjects who took part in this experiment, eight of whom had used similar apparatus and/or graphics systems before.

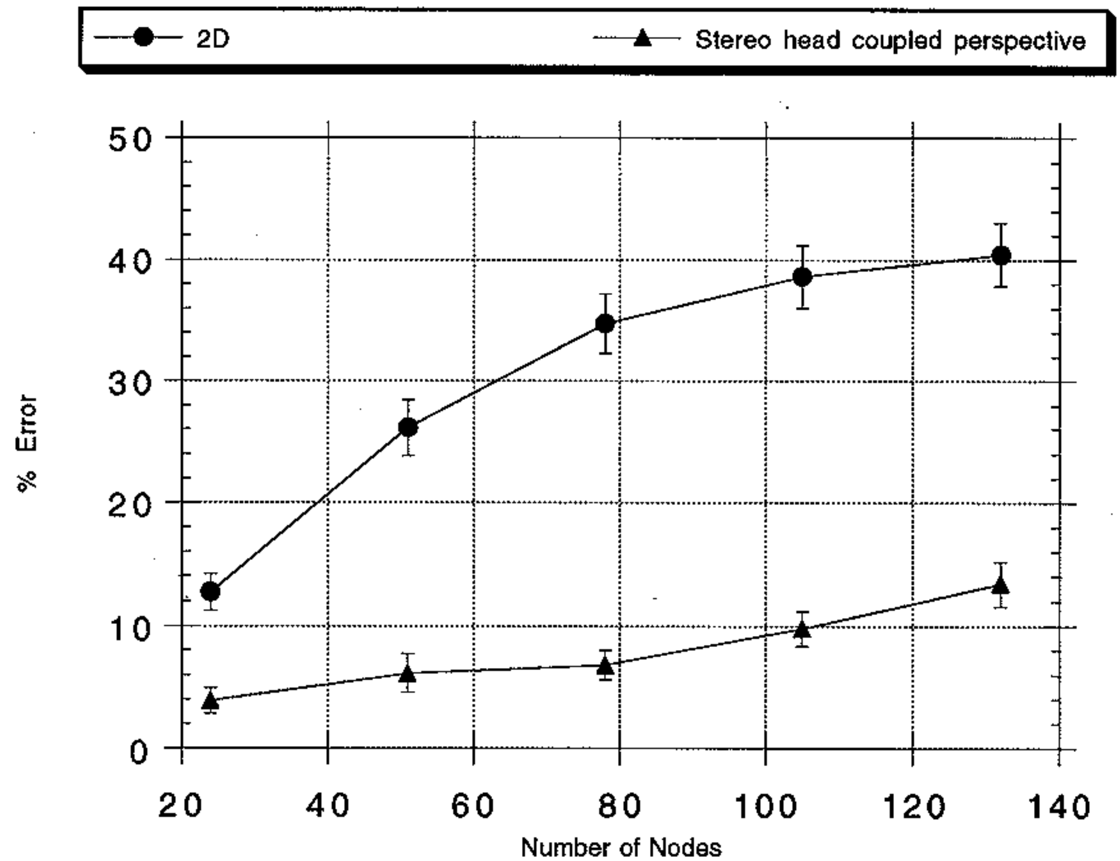

Figure 3. Error rate data from Experiment 1a.

2.3 Results and Discussion of Experiment 1a

The results from Experiment 1a are summarized in Figure 3. As can be seen, the errors in the head coupled stereo condition were dramatically reduced compared to those obtained with the 2D condition. In fact, the results were considerably better than anticipated; we had assumed that the condition with 132 nodes would yield large error rates in either condition. However, the error rate for 132 nodes in 3D was only marginally greater than the error rate for 24 nodes in 2D. Clearly, in order to make a meaningful estimate of the relative advantages of 3D over 2D, a considerably more complex graph was needed in the 3D condition.

3. EXPERIMENT 1b:

Based on the results of Experiment 1a we redesigned the experiment with different ranges of graph complexities for the different conditions. We also decided to add two conditions to help us answer the question of the relative benefits of head coupling and stereo.

3.1 Conditions

1) 2D: no stereo, no rotation; the 3D graph was projected onto a 2D plane using an orthographic (parallel) projection by removing Z axis information, hence no overlap information was available.

2) Stereo perspective: no rotation; this condition made use of a pair of StereoGraphics CrystalEyes LCD shutter glasses to provide the disparity depth cues.

3) Head coupled perspective: the scene's perspective projection changed continuously according to the subject's measured head position; the perspective projection was defined by a single viewpoint centered between the eyes.

4) Stereo, head coupled perspective: same as above, except with stereo. The correct view was generated for each eye position (continuously updated).

The numbers of nodes used in the different conditions were as follows.
1) 21, 42, 63, 84, 105
2) 51, 81, 111, 141, 171
3) 81, 117, 153, 189, 225
4) 111, 156, 201, 249, 291
To obtain the number of arcs multiply by 4/3.

The procedure for Experiment 1b was in most respects identical to Experiment 1a except that the number or trials per condition/size combination was reduced to 12 to keep the session time to a manageable size. This experiment involved 11 participants, six of whom had been exposed to a similar environment before. There was one person from Experiment 1a that took part in Experiment 1b as well.

3.2 Results and Discussion of Experiment 1b

Figure 4 summarizes the error data from this experiment, with data from Experiment 1a included for comparison. This figure shows a sequence of curves with varying gradients which appear to be roughly multipliers of each other with respect to the graph size. That is, error rate appears to be directly proportional to the number of nodes, with a different gradient for the different conditions. To test this model we fitted a set of straight lines through the data with a zero intercept. These are shown as the broad lines running through the sets of points in Figure 4. Note that the vertical bars represent one standard error and that the true mean should lie outside of the range of two standard errors approximately five percent of the time. This very simple model appears to be a reasonable first approximation to the data, although as the errors approach 40% there appears to be some flattening of the 2D curves.

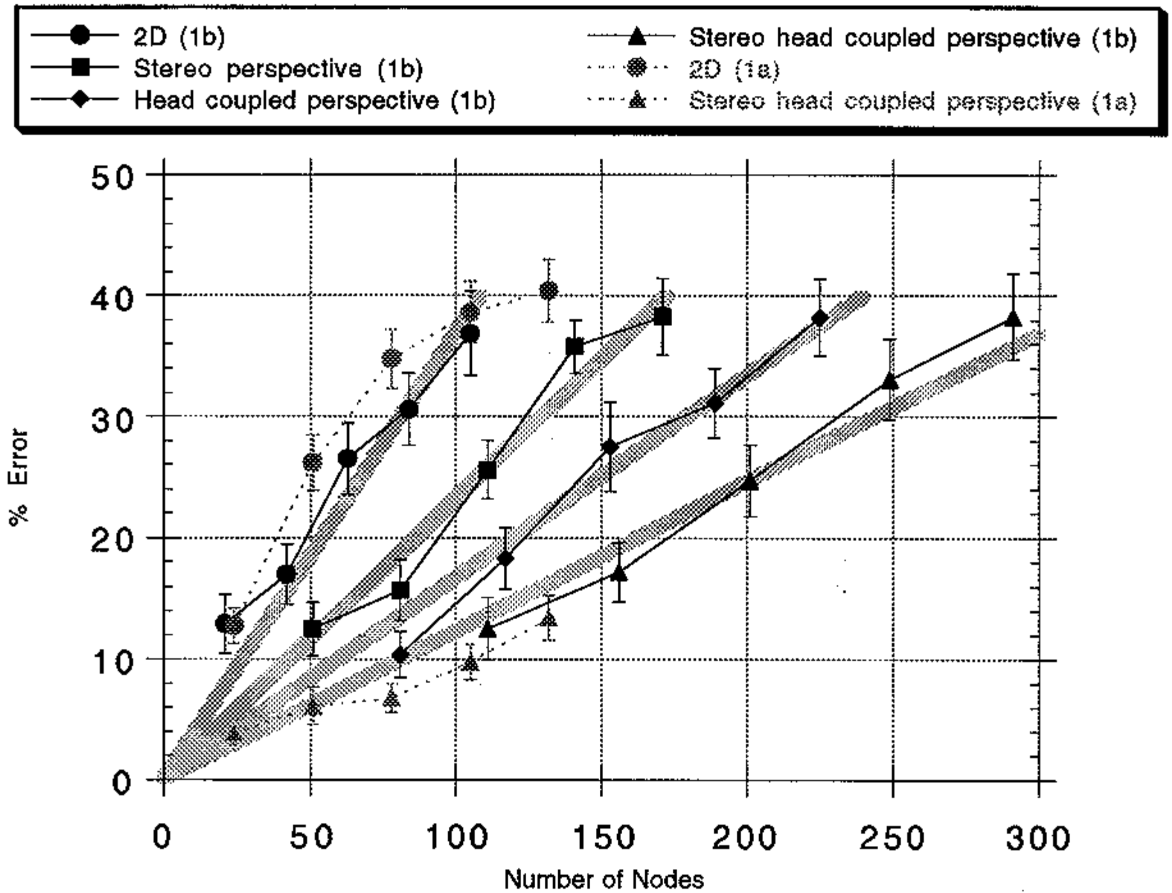

Figure 4. Error data from Experiments 1a and 1b. Vertical bars represent one standard error of the mean. The straight lines represent the simple model described in the text.

On this basis we conclude that the graph that can be understood with head coupled stereo is about 3.0 times as large as the 2D graph for any given error rate (taking the ratios of the gradients). Adding stereo alone appears to increase the comprehensible graph size by approximately a factor of 1.6 and adding head coupling alone appears to increase the comprehensible graph size by a factor of 2.2.

The average time for completion data is summarized in Figure 5. Completion time appears to depend more on the number of nodes and arcs in the graph than on the viewing mode, although there is evidence for different asymptotes for the different viewing modes. The times increase to about 13 seconds with approximately 100 nodes after which the curve levels off. The increasing part of the curve suggests that the time to process the data depends on the number of nodes, whereas the asymptote presumably reflects the amount of time it takes for a subject to feel that they have extracted all they usefully can from the data. This asymptote is lower for 2D viewing presumably because the graph appears as a hopeless jumble sooner than it does for the 3D viewing modes.

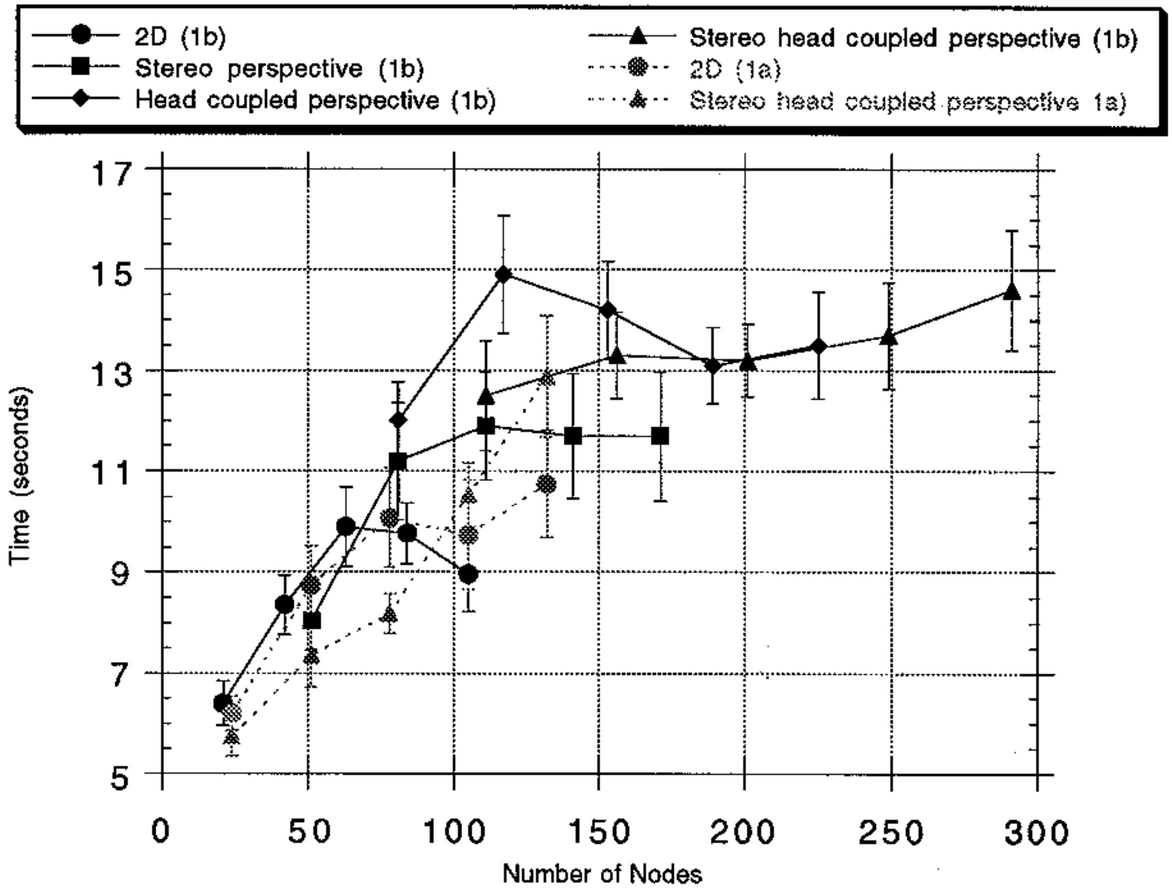

Figure 5. Time to completion data from Experiments 1a and 1b. Vertical bars represent one standard error of the mean.

4. EXPERIMENT 2: WHAT IS THE BEST VIEWING MODE?

Experiments 1a and 1b provide compelling evidence that the application of head coupling and stereo allows for greater information comprehension. Yet we know from the previous work described in the introduction that different kinds of rotation can be used to enhance comprehension of a tree structure. This raises the issue of the relative merits of different viewing rotation modes for data visualization. The nine conditions for this experiment included the four from Experiment 1b and in addition two conditions with rotation added by hand motion (with and without stereo), and two other conditions with the rotation happening automatically (with and without stereo). To complete the set we included a condition which was a static perspective view that contained overlap and size perspective information. In order to keep this experiment manageable in length we used only a single graph size for all conditions.

4.1 Method

There were 18 trials conducted under each of nine conditions (described below). There were always 75 nodes and 100 arcs in every trial. Otherwise the method was the same as for Experiments 1a and 1b. Once again there were eleven subjects involved in this experiment.

Five of these subjects had had exposure to a similar set-up before, and there were two participants who were involved in both experiments 1b and 2.

4.2 Conditions

1) 2D: no stereo, no rotation; the 3D graph was projected onto a 2D plane using an orthographic (parallel) projection by removing Z axis information, hence no overlap information was available.

2) Static Perspective: no stereo, no rotation; essentially the same task as in 1) above, except that the graph is displayed using a perspective projection with the depth cues of relative size and overlap/occlusion. Conditions 2 through 9 all used a perspective projection.

3) Stereo: no rotation; this condition made use of a pair of StereoGraphics CrystalEyes LCD shutter glasses to provide disparity depth cues.

4) Passive rotation: no stereo; the scene rotated at a constant angular velocity of 20 degrees/sec about a vertical axis.

5) Stereo, passive rotation: same as above except with stereo.

6) Hand coupled: no stereo, hand coupled; lateral movement of the mouse caused rotation of the scene about a vertical axis; mouse movement towards and away from the subject caused rotation of the scene about a horizontal axis. Movement was restricted to $\pm 128°$ about the vertical axis and $\pm 49.1°$ tilt (about the horizontal axis).

7) Stereo, hand coupled: same as above, except with stereo.

8) Head coupled perspective: the scene's projection changed continuously according to the subject's head position; the perspective projection was defined by a single viewpoint centered between the eyes.

9) Stereo, head coupled perspective: same as above, except with stereo. The correct view was generated for each eye position.

4.3 Results from Experiment 2

The results are summarized in Figure 6. The main difference between conditions is found in error rates. These ranged from a high of 26% in the 2D condition down to a low of 6.1% in the stereo hand coupled condition. These data confirm previous studies which show that motion is more important than stereo in reducing errors (the average for the motion alone conditions was 11.4% whereas the average for the stereo alone condition was 15.4%) However they also show that the combination of stereo and motion is the most effective (average error 7.5%), and interestingly, they suggest that the method for producing the motion is not particularly important. Overall, the data do appear to support the hypothesis that 3D visualization reduces the errors due to arc crossings.

The response times were relatively uniform across conditions and the fact that the times are approximately 10-12 seconds is entirely consistent with the data from Experiment 1b, given that the number of nodes was fixed at 75.

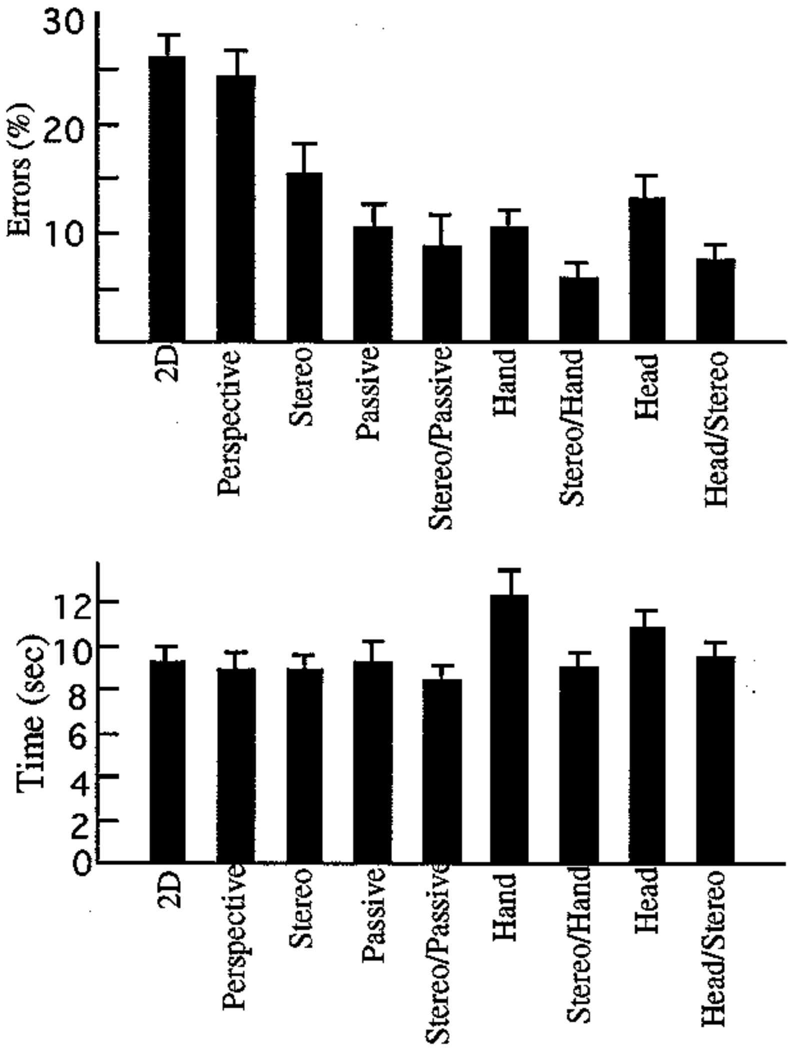

Figure 6. The results show small differences between times for the different conditions but large differences in error rates.

5. CONCLUSION

We believe that our major contribution in this paper has been to provide a first quantitative estimate of the benefits of stereo viewing with and without motion parallax. Our results surprised us, as we were prepared to believe that the second of the three theoretical views outlined in the introduction was the correct one and that adding 3D viewing would only have minimal benefits for an abstract data visualization task. The factor of three improvement that we actually measured (for the head coupled stereo condition) should provide some strong encouragement to those who are developing experimental applications which use 3D viewing in information management systems [10,12,14,20]. Our results suggest that this effort makes sense only if real-time rotation and/or stereo viewing is included in the data presentation. A static perspective image may add little in comparison with a 2D diagram and adding real time rotation is considerably more important than adding stereo. Unfortunately, structured motion is far more computationally expensive than stereo since it requires a minimum of ten screen updates per second for a duration of at least 30 seconds, whereas a stereo pair only requires the computation of two images.

The practical conclusion from Experiment 2 is that although a variety of methods for introducing structured motion can help comprehension, the type of motion used should depend on the application. For example, if the selection of objects is important then automatic rotation is not desirable because selecting moving objects is difficult. On the other hand, if head coupling is available then this would probably not interfere with 3D selection. This is presumably because the motor control systems used for visually guided hand placement have evolved to work in conjunction with simultaneous head motion.

On the theoretical issue of the additivity of depth cues we find a slight superadditivity in the cues of stereo and motion (60% improvement for stereo and 120% improvement for head coupling does not quite add up to the 200% improvement obtained with both). This agrees with the previous results of Sollenberger and Milgram [18]. We do not wish to make much of this observation since we find it hard to believe that such a simple model can account for something as complex as space perception. There is a possible alternative interpretation of our result which does not rely on space perception at all. In our displays, more 2D patterns were inherently ambiguous than were the 3D patterns in so far as they contained paths that could not be resolved due to the densely overlapping network of nodes and arcs. It might be argued that the main advantage of the 3D view was simply the fact that motion allows the resolution of ambiguous arc destinations. Perhaps a mechanism that made all the nodes in a 2D display move at random would work just as well. Certainly simple motion can be used to resolve patterns in abstract data [13]. This is an interesting possibility which needs further investigation, however it does not detract from the practical utility of our findings.

There does appear to be some penalty associated with using head coupled viewing. A number of subjects stated that they found viewing the graph in the head coupled stereo mode to be somewhat stressful. Part of this may be due to the difficulty of the task. Trying to perceive a path in a tangle of almost 300 nodes and 400 arcs is not easy. We also suspect that improvements in technology will help here. Noise in the head tracking system and the occasional briefly frozen image because of UNIX system functions are especially irritating in the head coupled viewing modes. It is as if having adopted the real-world virtual reality metaphor the brain expects the visual world to be stable, and if it is not we get a queasy feeling. This is a problem which improvements in head tracking and display technology will help solve.

As a final comment we wish to state that although the path tracing task we used is reasonably general to a large number of possible applications, more studies are clearly needed to establish the possible benefits of advanced 3D viewing using stereo and 3D motion for other tasks.

ACKNOWLEDGEMENTS

The primary source of support for this research is an NSERC Canada Strategic Grant. We are grateful to Arthur Ryman at IBM Toronto Labs for his encouragement and support and for his assistance in providing code analysis tools. We also thank Tim Dudley at Bell Northern Research for providing the initial impetus to start this project and for ongoing enthusiastic support.

REFERENCES

1. Arthur, K, Booth, K.S. and Ware, C. (1993) Evaluating Human Performance for Fishtank Virtual Reality. *ACM Transactions on Information Systems*, 11(3), 216-266
2. Braunstein, M.L. (1976) *Depth Perception Though Motion*. New York, Academic Press.

3. Bruno, N. and Cutting, J.E. (1988) Minimodularity and the Perception of Layout. *Journal of Experimental Psychology: General,* 117, 161-170.
4. Card, S.K, Robertson, G.G. and Mackinlay, J.D. (1991) The Information Visualizer, an Information Workspace. *CHI'91 Proceedings.* 181-188.
5. Deering, M. (1992) High resolution virtual reality. *Computer Graphics*, 26,2,195-202.
6. Eades and Xuemin. How to draw a directed graph. In *IEEE Workshop on Visual Languages.* 13-17, 1989.
7. Fairchild, K.M., Poltrock, S.E. and Furnas, G.W. (1988) SemNet: Three-Dimensional Graphic Representations of Large Knowledge Bases. In *Cognitive Science and Its Applications for Human-Computer Interaction.* Ed Raymond Guindon Lawrence Erlbaum. 201-233.
8. Garey, M.R. and Johnson, D.S. (1979) *Computers and Intractability - A Guide to the Theory of NP-Completeness,* Freeman, 1979.
9. Gibson, E.J., Gibson, J.J. Smith, O.W. and Flock H. (1959) Motional parallax as a determinant of perceived depth. *Journal of Experimental Psychology.* 58. 40-51.
10. Koike, H. (1993) The role of another spatial dimensions in software visualization. *ACM Transactions on Information Systems*, 11(3) 266-286.
11. Lack, L.C. (1974) Selective attention and the control of binocular rivalry. *Perception and Psychophysics*, 15, 193-200.
12. Lieberman, H. (1989) A Three-Dimensional Representation for Program Execution, *IEEE Workshop on Visual Languages Proceedings*, 111-116.
13. Limoges, S., Ware, C. and Knight, W. (1989). Displaying Correlations using Position, Motion, Point Size or Point Color. *Graphics Interface Proceedings*, 262-265.
14. Mariani, J.A., and Lougher, R. TripleSpace: an Experiment in a 3D Graphical Interface to a Binary Relational Database, *Interacting with Computers*, 4(2) 1992 147-162
15. Purcell, D.G. and Stewart, A.L. (1991) The object detection effect: Configuration enhances perception, *Perception and Psychophysics*, 50(3) 215-224.
16. Robertson, G.G., Mackinlay, J.D and Card, S. K. (1991) Cone Trees: Animated 3D Visualizations of Hierarchical Information. *CHI'91 Proceedings*. 189-194.
17. Sollenberger, R.L. and Milgram, P. (1991) A comparative Study of Rotational and Stereoscopic Computer Graphic Depth Cues. *Proceedings of the Human Factors Society Annual Meeting*, 1452-1456.
18. Sollenberger, R.L and Milgram, P. (1993) The effects of Stereoscopic and Rotational Displays in a Three-Dimensional Path-Tracing Task. Human Factors, 35(3) 483-500.
19. Ware, C. Arthur, K. and Booth, K.S. (1993) Fishtank Virtual Reality. INTERCHI'93 Technical Paper. Proceedings 37-42.
20. Ware, C., Hui, D. and Franck, G. Visualizing object oriented software in three dimensions. Proceedings, CASCON'93. IBM Canada Ltd. and NRC. Toronto, October 1993. 612-620.