Motion to Support Rapid Interactive Queries on Node-Link Diagrams

COLIN WARE
Data Visualization Research Lab
Center for Coastal and Ocean Mapping
University of New Hampshire
Durham, NH 03824
colinw@cisunix.unh.edu
and
ROBERT BOBROW
BBN Technologies
10 Moulton St.
Cambridge, MA
rbobrow@bbn.com

Many different problems can be represented as graphs displayed in the form of node-link diagrams. However, when a graph is large it becomes visually uninterpretable because of the tangle of links. We describe a set of techniques that use motion in an interactive interface to provide effective access to larger graphs. Touching a node with the mouse cursor causes that node and the subgraph of closely connected nodes to oscillate. We argue from perceptual principles that this should be a more effective way of interactively highlighting a subgraph than more conventional static methods. The MEGraph system was developed to gain experience with different forms of motion highlighting. Based on positive feedback, three experiments were carried out to evaluate the effectiveness of motion highlighting for specific tasks. All three showed motion to be more effective than static highlighting, both in increasing the speed of response for a variety of visual queries, and in reducing errors. We argue that motion highlighting can be a valuable technique in applications that require users to understand large graphs.

Categories and Subject Descriptors: H1.2 [**Models and Principles**]: User/Machine Systems – *Human information processing; Human Factors*; H5.2 [**Information Interfaces and Presentation**]: User Interfaces - *Interaction styles*; *Theory and methods*.
General Terms: Design, Experimentation, Human Factors, Design
Additional Key Words and Phrases: Human-computer interaction, motion highlighting, visual queries. visualization, data visualization, information visualization, interactive visualization

---

1. INTRODUCTION

Many problem domains can be mapped into a graph and visually displayed in the form of a node-link diagram. Typically, nodes represent entities while the links represent relationships between the nodes. Many variants on node-link diagrams allow for different types of nodes to be represented by shapes or colors and different relationships to be represented by varied graphical styles for the links. Good examples are the entity relationship diagram and the entire family of Unified Modeling Language (UML) diagrams. Such diagrams are used for software management, networks analysis and management, the mapping of social relationships and many other applications.

Node-link diagrams allow for the rapid comprehension of the inter-relationships between entities and for certain patterns of links to be immediately evident. Unfortunately, a major drawback is that they do not scale well to larger applications. A graph with more that twenty or thirty nodes and twice as many links between them becomes visually incomprehensible, although many applications require the viewing of much larger graphs.

Various attempts have been made to increase the size of the graph that can be readily comprehended. One proposed solution has been to represent the information structure in 3D (Fairchild et al, 1988; Parker et al, 1998; Chen et al, 2001). Unfortunately, research suggests that merely showing a static perspective picture of a graph laid out in 3D is not better than a 2D layout (Ware and Franck, 1996). It is important to include non-pictorial depth cues to actually see more. To visualize such structures effectively in 3D the most important cues are stereoscopic viewing and kinetic depth. Stereoscopic views can be produced on a computer screen by means of shutter glasses that alternately block each eye's view. When synchronized with the refresh of the computer monitor this allows for different stereo views to be shown to the left and right eyes of the viewer. Kinetic depth is produced by having the node link structure rotate while being viewed. Studies have shown that these methods do indeed allow for larger graphs to be perceived and that kinetic depth is more important than stereoscopic depth (Sollenberger and Milgram, 1993; Ware and Franck, 1996). When both stereo depth and kinetic depth cues are provided it may be possible to see a structure approximately three times as large (Ware and Franck, 1996).

However, neither stereoscopic viewing nor kinetic depth are practical techniques for many applications. Users do not like putting on the special glasses that are usually required to get stereoscopic depth and the extra cost can be an issue. Because depth of focus information is not provided in computer displays, eye-strain may result from prolonged stereoscopic viewing (Miyashita and Uchida, 1990). Also, having the graph rotate in 3D to provide kinetic depth can be a problem since it makes it more difficult to interactively select nodes using the mouse. Rotation also makes it more difficult to read text labels associated with nodes.

Another way of dealing with large graphs is to use interactive techniques to reveal or hide information as needed. If this is done, the interface must be simple, otherwise the cognitive burden of interaction will cause the user to lose his or her train of thought.
A good example of a system that used simple, low cognitive cost interactions is Constellation (Munzner et al, 1999). Constellation has been applied as a user interface to MindNet, a natural language semantic network that stores dictionary or encyclopedia entries in the form of definition graphs. The nodes represent conceptual entities. The links represent relations such as "is-a", "part-of" or other kinds of relationships between the entities. The Constellation user simply passes the cursor over a node to perform rapid "hover" queries, this causes all the links attached to that node to be highlighted and by brushing the cursor over the nodes the user can get information about the graph structure as needed. The result is that a graph having too many links to be comprehensible in a static diagram, can be interactively understood.

We believe that this kind of rapid interaction, using hover queries, can dramatically increase the size of graph that can be usefully displayed. However, the method used for highlighting the results of the hover queries is a difficult and critical design choice. For complex information structures, color and shape may have already been used to denote link type and node type; therefore we cannot use these visual properties as a highlighting technique.

1.1 Motion Coding

There are a number of reasons for thinking that using *motion* to highlight subsets of nodes and links in a graph may be especially effective. Motion can be used as a visual coding attribute in much that same way that color, shape or texture can be used (Bartram and Ware, 2002). Data glyphs can be motion coded by making them oscillate about a fixed center and different patterns of motion can be used to signal different data attributes. Bartram et al (2003) showed motion to be much more effective than color or shape change in signaling a change to a data glyph presented on the screen. Motion has also been shown to be an effective visual grouping method allowing disparate graphical objects to be visually perceived as "foreground" and distinct from other groups of objects that are either static or have different motion. Michotte (1963) found that objects moving together at the same speed were perceptually grouped and he called this effect kinematic integration.

Another aspect of motion that makes it especially attractive as a coding method is that it may allow for visual pre-attentive conjunction searches. Pre-attentive coding is the name given for a visual stimulus that "pops out" visually from a display. For example, a red dot will be immediately visible in a field of black dots. Orientation is also pre-attentive. A horizontal bar stands out in a field of vertical bars. But more complex patterns are typically not pre-attentive. For example, a green horizontal bar will not popout from a field containing both horizontal red bars and green vertical bars. This is called a conjunction search because it involves the combination of orientation and color (Treisman and Gormican, 1991).

There are, however, some examples where visual conjunctions do produce popout effects and these may be especially useful for use in interacting with data because they allow for a rapid visual search for data based on *two* attributes. In particular several conjunctions involving motion have been shown to be pre-attentive. The conjunction of motion and simple shape can be pre-attentive (McLeod et al., 1988; Driver et al 1992), although the effect is asymmetric, it is easier to fined the conjunction target if it is one of the moving items than if it is one of the static items (McLeod et al., 1988). Also, Nakayama (reported in Treisman, 1988) found that conjunctions of movement and color to be preattentive. As an underlying mechanism, Driver and McLeod, 1992) argued that cortical area MT may provide a kind of "movement filter" for separating out moving from stationary patterns. This suggests that using motion as a highlighting technique highlighting will be good for separating a group of objects enabling a rapid visual search on the group (Bartram et al, 2003). If a set of items are highlighted using motion it should be possible to do a rapid visual conjunctive search for the things that are in motion and are also red (or coded with a simple shape). It is this reasoning that lead to the development of the techniques presented and evaluated here. In this paper we present techniques for using motion in conjunction with rapid hover queries.

1.2 Design Philosophy

There is an assumption inherent in this work that we wish to make explicit at the outset. Our assumption is that that when we look at a node-link diagram we are mostly concerned with perceiving relationships between nodes that are close to one another (i.e. in topological proximity). For example, we might pick a particular node of interest that represents a known criminal. We will be interested in other nodes that are within one or two links. These might represent various other suspect individuals in a social network. Being able to see those nearby links clearly will support a whole range of visual queries concerning the nature of the other individuals and the nature of the links between them. To support this the attributes of the individuals and the links may be visually represented by color or shape.

The purpose of visualization, in general, is not to support any particular visual query, but rather to support a range of queries that are rapidly constructed on an ad hoc basis as part of visual thinking. Therefore, we need techniques that can enhance a large range of the most frequent and useful visual queries. We believe that topological range highlighting is such a general technique. This highlighting supports a wide range of more specific queries that are conducted by visual search. A search for any pattern of nodes and/or links within proximity to the selected node will be enhanced.

However, in order to *evaluate* the value of a particular method for representing data, it is necessary to make that evaluation for specific, well defined tasks. The experiments that follow test three different visual queries. These should be considered as examples of a very large class of visual queries that would be enhanced by our method.

2 THE MEGRAPH SYSTEM

In order to understand better how effective motion-based interaction could be in helping us understand large graphs we first developed a simple research prototype (MEGraph) and then carried out a series of formal experiments that assessed subject's ability to carry our simple visual search tasks with the system. MEGgraph stands for Motion Enhanced Graph. MEGraph contains algorithms to generate graphs of different sizes with different degrees of interconnectivity of the nodes. Once the graph is generated, a simple spring layout algorithm is applied to lay the nodes out on a 2D plane. Spring layout is widely used and is generally regarded as the best way of laying out non-directed graphs (Battista et al, 1999).

Highlighting methods: We carried out a period of iterative design where we informally tried out a number of motion and non-motion highlighting methods. We ended up with the following set for demonstration purposes (detailed parameter values for the different motions are provided in the method sections of the experiments described later in this paper).

1. *Static highlighting.* The nodes within the query radius are shown ringed in black and the highlighted links are also shown as black. Non-highlighted nodes and links are grayed out.
2. *Motion highlighting – circular.* In this mode the selected node takes on a circular motion around the center position, obtained from the spring layout. When connected are highlighted, the link between them also moves, creating a moving subgraph. In dual query mode the two subgraphs move counterphase.
3. *Motion highlighting – jolt.* In this mode the selected nodes and links move in pulses using the function (sine(t)/t^2). The result is similar to the effect of an object being struck briefly oscillating from the blow. Pulses are repeated about every second.
4. *Motion highlighting – crawl.* In this mode the selected links show smoothly animated sawtooth patterns radiating out from the selected node
5. *Motion highlighting – expanding nodes.* In this mode the selected nodes grow larger and smaller with a period of approximately one second.

From informal evaluations we concluded that all of these might have some application except for the expanding nodes method that we found less effective. Figure 1 shows an example of the kind of random graph produced by MEGraph and illustrates the static highlighting of nodded and edges as well as the pattern used in crawl highlighting.

In MEGraph we also implement some capabilities for rapid interaction. These include the following.

Click Queries: Clicking on a node causes a breadth first search of the graph beginning at that node and highlights nodes and links within a designated topological radius of the selected node. The radius of the query can be changed by entering an integer value using the keyboard.

Clickless Queries: These are the same as click queries but they do not require a mouse click. To elicit the range search and highlighting it is merely necessary to place the cursor over a node to elicit a query. In this mode the mouse can be moved over the nodes of the graph causing a series of queries to be rapidly triggered, perhaps as fast as 2 or 3 per second.

Dual Queries: In dual query mode, two regions of a highlighted graph can be simultaneously highlighted. The two most recently clicked on nodes are highlighted. If a motion highlighting method is used the two groups of nodes are highlighted in such a way that the groups are out of phase with one another.

Our experience showing MEGraph to a number of potential users of graph displays suggested to us that the motion highlighting techniques were effective in giving users visual access to a graph that was considerably larger than a static graph. The initial response was very favorable. To more formally evaluate the method we undertook a series of experiments to measure the potential benefits of motion highlighting. However, to repeat a point we made earlier, it is not that the particular queries in these experiments are especially important, although they may be. They are intended merely as examples of the hundreds of different visual queries that might be executed in using a graph as a problem solving tool.

In order to investigate the effectiveness of using motion as a visual highlighting technique to support visual queries, we constructed a version of MEGraph capable of running controlled experiments. We conducted a series of three experiments, each with a different visual query as the task. All three tasks required the user to find a particular pattern in the nodes and links lying within a topological radius of two links from a particular designated node. Our general hypothesis was that because motion affords conjunction search between color and motion, the responses would be faster and have fewer errors with motion highlighting than with static highlighting.

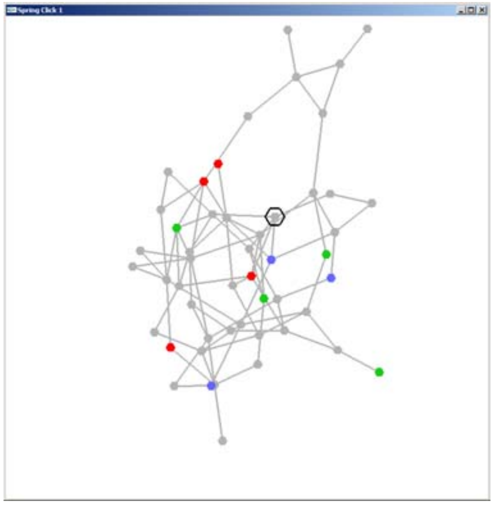

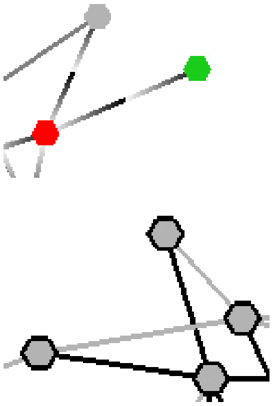

Figure 1. On the left is an example of the graph used in the experiment. This is how it looked in the no highlighting condition. On the right are fragments illustrating the crawl highlighting and the static highlighting.

3 EXPERIMENT 1: SIMPLE VISUAL QUERIES

For the first experiment the task was to determine whether or not there were two red-nodes within a radius of two links from a node identified by having a circle drawn around it.

3.1 The Graph

A graph was algorithmically generated for each trial. Each graph had exactly 50 nodes and 82 links on average. The generating algorithm produced on average 6 nodes with one link, 12 nodes with two links, 12 nodes with 3 links, 8 nodes with 4 links and 12 nodes with 5 links. The actual numbers varied from trial to trial. On average each node was connected to slightly more than 3 others.

The spring layout produced a mean link length of approximately 2.25 cm on average. The link width was 0.083 cm. The nodes were drawn with circles approximately 0.5 cm in diameter. The display area was 22 cm square and was viewed from approximately 56 cm (at this viewing distance 1 cm = approx 1 degree of visual angle). The background on which the graph was drawn was white with a luminance of 56 cd/m^2, the luminance of the gray values was 30 cd/m^2 and the luminance of the black background was 3.1 cd/m^2 under the low illumination levels used. The following rgb values were used for red green and blue respectively (1.0,0.0,0.0), (0.1,0.8,0.1), (0.4, 0.4,1.0). These produced colors roughly equated in brightness- with the gray value (judged by eye).

3.2 Conditions

In each case where highlighting occurred it was always within a toplogical radius of two links from a node selected at random. This selected node was always connected to the rest of the graph by at least two links. There were 6 different highlighting conditions.

1. *No highlighting.* All nodes and links were shown in gray except for 20% of the nodes which were colored red, green and blue.
2. *Static highlighting.* In the selected subgraph links were black and nodes were circled in black. In the non-selected subgraph the links and nodes were gray.
3. *Circular motion.* Nodes and links in the selected subgraph moved in a circular path. The amplitude of the motion was 0.4 cm and the frequency was 2 Hz. This resulted in an angular velocity at the circumference of the motion path of 2.5 deg/sec.
4. *Jolt motion.* Nodes and links in the selected subgraph moved with the jolt function. $ampl = 0.5*\sin(8.0\pi t)/(60t*t + 0.5)$ cm.
5. *Crawl motion.* A sawtooth pattern was used as shown in Figure 1. Two cycles were present on each link and the pattern was animated so as to repeat at 2 Hz. With an average angular velocity of 2.25 deg/sec.
6. *Motion and Brightness.* A combination of moving nodes and links as in condition 3 with static highlighting as in condition 2.

3.3 Trials

Subjects were asked, “Are there at least two red nodes within two links of the highlighted node?”   They responded “yes” or “no” by pressing the right mouse button or the left mouse button respectively.

A new graph was automatically generated and laid out for each trial. On half of the trials there were fewer than two nodes in the highlighted subgraph (requiring a “no” response). On the other half of the trials there were two or more (requiring a “yes” response). There was an inter-trial interval of one second when the screen was blank white.

This was a within subject’s design, each subject received all conditions. Trials were given in blocks of 12 graphs consisting of one for each condition requiring a “yes” response and one for each condition requiring a “no” response. Each block of 12 was repeated 9 times in a different random order. Subjects took short breaks after each set of three blocks.

At the start of the experiment subjects was given a practice session consisting of single block containing all of the conditions. These data were not included in the analysis.

3.4 Subjects

There were 13 subjects. They were a mix of students, research assistants, and one faculty member. Most were paid for participating but a few declined payment.

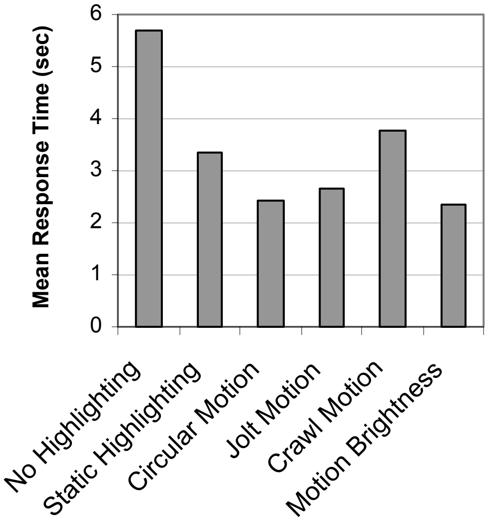

Figure 2. Summary response times from Experiment 1.

3.5 Results from Experiment 1

The response time results from Experiment 1 are summarized in Figure 2. An analysis of variance was run on the following factors: condition, target present vs target absent, and subject. The experimental condition was significant ($F(5,60) = 23.1$; $p < 0.001$). Tukey’s post hoc Honestly Significant Difference (HSD) comparisons revealed 3 groups. The *no highlighting* condition is significantly slower than all the others ($p<0.05$). The *crawl motion* and the *static highlighting* conditions are in a second intermediate group. The *circular motion, jolt motion* and *circular motion+brightness* highlighting conditions are in a third group with the fastest responses.

The analysis of variance also revealed the "yes" responses (2.97 sec) to be faster on average than the "no" responses (3.78 sec) ($F(1,12) = 38.7$; $p < 0.001$). Subjects also differed significantly (F12,30)=11.27; $p < 0.001$). In addition there was a condition/subject interaction (F60,60) = 2.5; $p < 0.001$).

The error results are summarized in Figure 3. As can be seen, the no-highlighting conditions had almost a 30% error rate, the crawl motion condition had a 12 % error rate whereas the other conditions had error rates of 4% or lower. Static highlighting had the lowest error rate (2.1%).

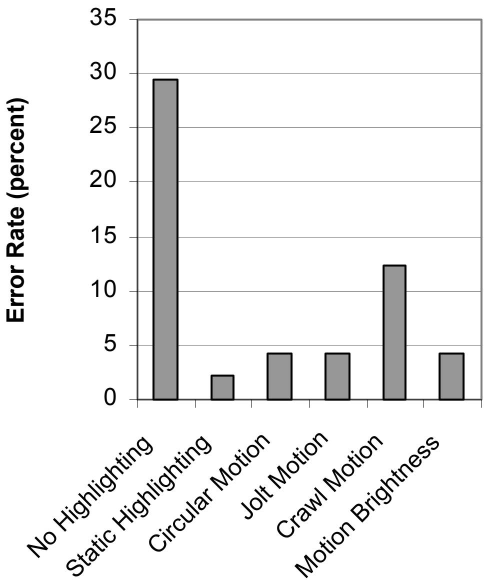

Figure 3. A summary of the error rates in Experiment 1.

3.5 Discussion of Experiment 1

The timing results from Experiment 1 supported our hypotheses that motion highlighting would be more effective in this conjunction search task. Motion highlighting cut the response times in half and drastically reduced the errors to a small fraction of those obtained with the no highlighting condition although static highlighting produced the lowest error rates. Motion highlighting was faster than static highlighting although static highlighting condition was much better than no highlighting. The error rate of close to 30% with no highlighting is so large as to suggest that a graph of this size is not likely to be useful in the absence of some interactive highlighting method. However, the error rates with all of the highlighting methods, except crawl motion, are probably acceptable given that this was a rapid response task.

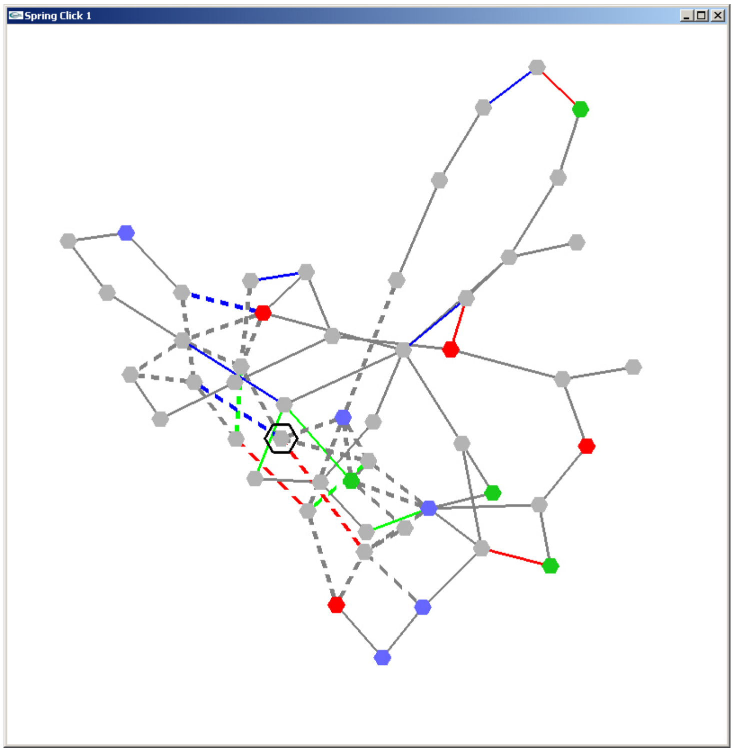

Figure 4. An example of edges only highlighting from Experiment 2. The fact that two red dashed edges are present indicates that this is a “yes” condition.

4 EXPERIMENT 2: SELECTIVE HIGHLIGHTING

With Experiment 2 we were interested in examining the selectivity of motion highlighting. Could we use the technique to selectively highlight nodes or links or both? To carry out this evaluation we constructed an experiment with two parts. In Experiment 2a subjects were required to look for at least two red *node* targets within a two link topological radius of the circled node. In Experiment 2b subjects were required to look for at least two red *link* targets within two links of the designated node.

4.1 Method

To look at the effects of selective highlighting we implemented conditions in which nodes or links or both were highlighted using motion or static highlighting methods.

All of the layout parameters and the highlighting techniques patterns were the same for both parts of the experiment. A difference from Experiment 1 was that we used more color-coding of nodes and links to simulate a graph visualization of greater complexity. Twenty five percent of the nodes and links were color coded with the colors red, green and blue randomly assigned. The remainder were colored gray.

A different method was adopted for static highlighting. Because of the colored links the method of using black lines instead of gray was not usable. To statically highlight links we made them wider (0.138 cm compared with .083cm for the other links) and dashed when highlighted. We changed the motion highlighting so that either the nodes, or the links, or both could move. We also reduced the amplitude of motion (to 0.3 cm) so that when a node was in motion independently from links, the components were still visually connected. The viewing distance was again 56 cm at which 1 cm subtends 1 degree of visual angle.

There were 8 conditions:

1. No highlighting
2. Static node highlighting (ringed in black)
3. Static link highlighting (wide dashed lines). This is illustrated in Figure 4..
4. Static node and link highlighting (ringed nodes and dashed links).
5. Circular motion of nodes
6. Circular motion of links
7. Circular motion of nodes and links
8. Crawl motion links.

Notice that conditions 5, 6 and 7 parallel conditions 2, 3 and 4 with the former using static highlighting methods and the latter using motion highlighting methods.

The two parts of the experiment differed in the task the user performed.

Task A: Are there at least two red *nodes* within two links of the highlighted subgraph?
Task B: Are there at least two red *links* within two links of the highlighted subgraph?

We predicted that for finding node targets (Task A), the moving node and static node highlighting would be most effective, whereas for finding link targets (Task B) the moving link and static link highlighting would be most effective. We also predicted that the crawl link motion would be effective in highlighting links.

The 13 subjects were the same as in the previous experiment.

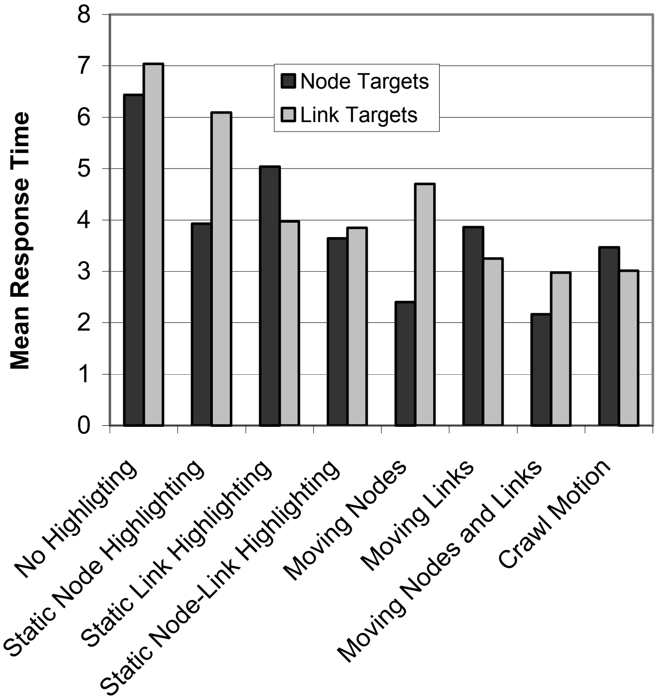

Figure 5. Summary of response time results from Experiment 2.

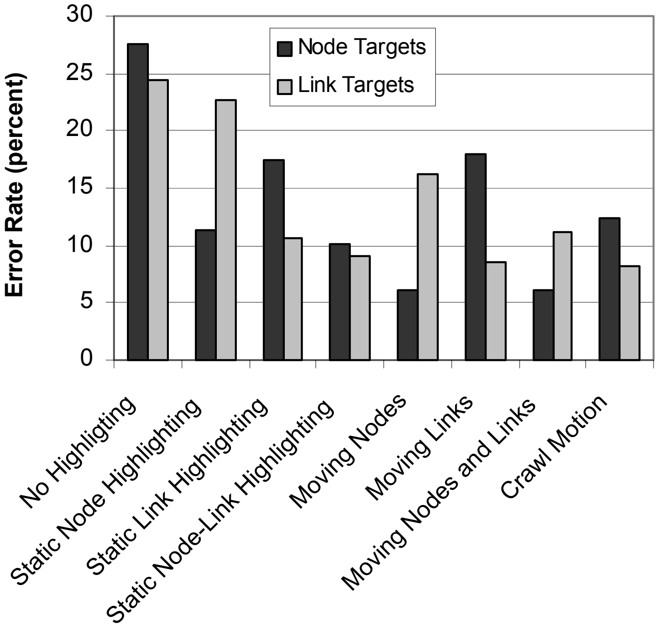

Figure 6. Summary of error results from Experiment 2.

4.2 Results from Experiment 2

The results from Experiment 2 are summarized in Figures 5 and 6. An analysis of variance was run on the following factors: condition, target type, and subject. There was a main effect of highlighting condition ($F(7,84)= 22.1$, $p < 0.0001$) and as predicted there was a strong interaction between the kind of highlighting and the task ($F(7,84)= 10.9$, $p< 0.0001$). When the nodes were highlighted, red node targets were responded to fastest with fewer errors. When the links were highlighted, the red link targets were responded to fastest and with fewer errors. It is notable that the combined node and link highlighting methods performed as well as the individual highlighting techniques for both nodes and links. The motion techniques (mean time = 3.25 sec) overall were faster than the static techniques (4.42 sec). As in Experiment 1 motion highlighting resulted in response times that were more than twice as fast as no highlighting. The analysis of variance also revealed the “yes” responses (3.9 sec) to be faster than the “no” responses (4.3 sec) ($F(1,12) = 6.2$; $p < 0.05$). There were significant individual differences ($F(12,15)= 5.16$, $p < 0.002$) and an interaction between subject and condition $F(84,77)= 2.079$; $p < 0.001$).

Separate ANOVAs were run on the node target and the link target conditions to look at the differences between highlighting methods using Tukey HSD tests for each target type. This revealed the differences shown in Tables 1 and 2. For the node targets the *moving modes and links* and the *moving nodes* conditions were the fastest and no highlighting was the slowest with intermediate subsets as shown. For the link targets the *moving nodes and link*s and the *crawling links* conditions were faster than all other conditions. The *moving links* condition was also in the fastest subset but did not differ significantly from the next subset. Our hypothesis that the *crawl motion* condition would be effective when the task is identifying links was confirmed. The *crawl motion* response times were in the fastest subset when the targets were links.

Table: I. Homogeneous subsets according to Tukey HSD test for node targets.

Moving Nodes and Links	2.12			
Moving Nodes	2.35			
Crawl Motion Links		3.41		
Static Node Link Highlighting		3.58		
Static Node Highlighting		3.84		
Moving Links		3.87		
Static Link Highlighting			5.0	
No Highlighting				6.55

Table: II. Homogeneous subsets according to Tukey HSD test for link targets.

Moving Nodes and Links	3.00				
Crawl Motion Links	3.01				
Moving Links	3.26	3.26			
Static Node Link Highlighting		3.85			
Static Link Highlighting		3.97	3.97		
Moving Nodes			4.70		
Static Node Highlighting				6.08	
No Highlighting					7.04

The pattern of errors followed the pattern of response times. Generally the conditions that had short response times also had the lowest errors as is shown in Figure 6. The lowest error rates were in the Moving Nodes and Links and the Moving Nodes conditions.

4.3 Discussion of Experiment 2

Although the results showed that separate highlighting of nodes and links could be used to pick out nodes and links respectively, they also showed that having both nodes and links move at the same time, was just as good or better. This supports our contention that the topological radius highlighting using motion is a general purpose technique. It facilitates a large variety of perceptual queries on the nodes and links within the highlighted topological radius of the target node. The static highlighting method was also quite good. It was the equal of the motion techniques for revealing links, although not as good as circular motion for revealing nodes.

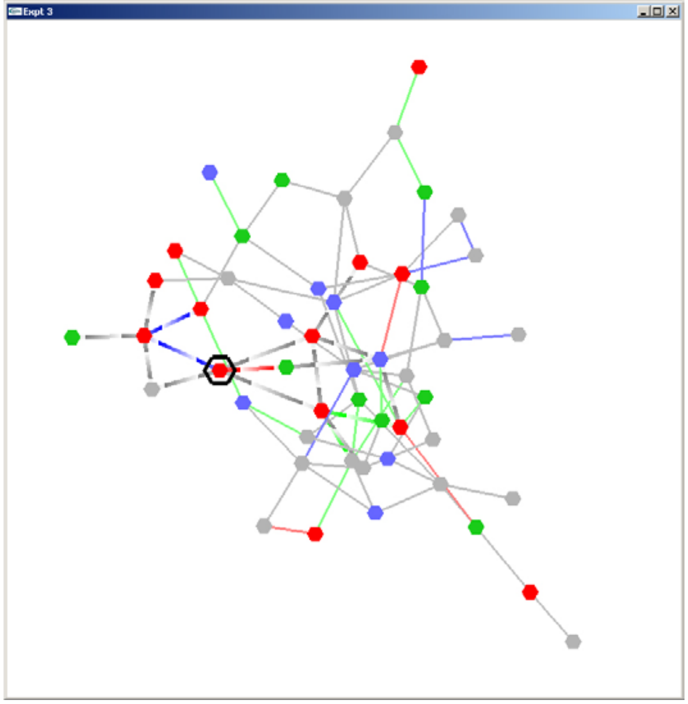

Figure 7. The target pattern consists of three red nodes connected by two blue links.

5 EXPERIMENT 3: COMPLEX PATTERNS

For the third experiment we designed a task involving the identification of more complex patterns. The target pattern consisted of a short path, consisting of three red colored nodes with blue colored links connecting them in a chain (see Figure 7). One of the red nodes was the circled start node. The links in the entire graph are pseudo-randomly colored. We conjectured that the crawl motion might be especially useful for this task since it seemed to emphasize paths best.

For this experiment there were only four conditions. We dropped the control condition with no highlighting.

The 13 subjects were the same as in the previous two experiments.

Conditions:

1) Static highlighting (Wide dashed lines and circled nodes)
2) Circular motion of nodes and links
3) Crawl motion of links
4) Circular motion of nodes and links combined with crawl motion.

All for the stimulus parameters (color, motion, line widths, etc. were the same as for Experiment 2).

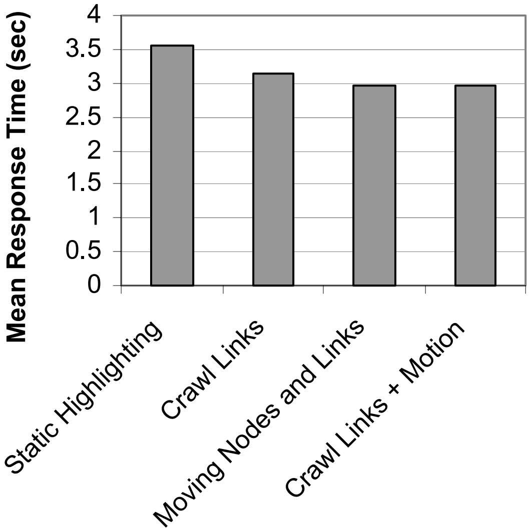

Figure 8. A summary of the response time results from Experiment 3.

5.1 Results and Discussion of Experiment 3

An analysis of variance was run on the following factors: condition, target present vs target absent, and subject. There were significant main effect for condition (F(3,65) = 24.9 $p < 0.0001$); for target present vs absent (F(1,12) = 42.1; $p < 0.0001$), and for subject (F(12,11)= 8.15; $p < 0.001$). A post hoc Tukey (HSD) test showed only the static highlighting condition to be significantly slower than the others with no differences between the three motion conditions.

The pattern of errors was very consistent across the four conditions, varying only between 10 and 11 percent. However, there were almost twice as many errors in “yes” responses (14.3%) than with the “no” responses (7.6%), indicating a bias towards false positives.

The results failed to confirm our hypothesis that the crawling motion on the links would be superior for the task of identifying this particular pattern on nodes and links. There was less of an overall difference between the conditions than in Experiments 1 and 2.

The results from Experiment 3 showed less of an advantage of motion highlighting over static highlighting relative to Experiments 1 and 2. This may have been because the pattern was relatively easy to identify, although more complex. The fact that the pattern always started at the ringed node may also have made the visual search easier.

6 DISCUSSION

Our results strongly support our original insight that quick clickless queries, together with motion highlighting will be effective in allowing rapid and more accurate interactions with graphs. All three of the experiments showed motion highlighting to be faster than the static highlighting method we employed. Since we did not test all possible static highlighting methods we cannot claim that motion highlighting will always surpass static highlighting, although we do believe that the parameter values we chose for both motion and static highlighting are reasonable. The error rate data generally followed the response time data, conditions with fast responses also yielded lower error rates.

The third experiment yielded much smaller differences between conditions. In this case the target pattern was more complex consisting of a pattern of nodes and links radiating from the highlighted node. One explanation for the relatively small differences may be that the visual search task was relatively easy. The highlighted node was easy to spot and provided visual guidance to the target pattern.

We believe that with clickless motion highlighting, a graph several times as large can become usable as an information display. The error rates in particular support this claim. Error rates with the highlighting methods were a fraction of those without highlighting. Nevertheless, a reasonable objection to our experiments is that they dealt with diagrams only slightly larger than those commonly used in practice. We do have some further anecdotal evidence to offer. We have improved our MEGraph software to the point where it can show graphs containing up to 1000 nodes and we are able to report, subjectively at least, that the motion highlighting techniques are even more effective with larger graphs.

The problem domain we are specifically interested in is that of social networks (Brandes et al, 2001). However, we can imagine the motion techniques we describe as being applicable to a broad class of complex node-link diagrams, including circuit diagrams, software diagrams and concept maps. In general, any application where the data can be represented with a node-link diagram and where it is important to find patterns of nodes and links in topological proximity may benefit.

As we stated in the introduction, our technique is based on the assumption that when we look as a node-link diagram, what we are primarily interested in are the items that are closely coupled. However, this does not mean that topological proximity to the selected node is the only possible basis for highlighting. We are beginning an investigation of other kinds of general purpose interactive queries including a method that interactively highlights the set of shortest paths and near-shortest paths between two selected nodes in a graph. In a transportation network visualization, this would highlight the most likely logistical pathways for the transport of goods.

As a final comment, using MEGraph in clickless mode is something like playing a musical instrument. Not in the sense that sound is produced, it is not, but rather in the immediacy of the feedback, and in the sense of being closely in touch with the instrument. As we pass the cursor over the nodes they vibrate, revealing their associations and we feel an immediacy of connection with the data that is unusual.

ACKNOWLEDGMENTS

This study was supported and monitored by the Advanced Research and Development

Activity (ARDA) and the National Imagery and Mapping Agency (NIMA) under Contract Number NMA401-02-C-0019. We also gratefully acknowledge the support of NOAA in providing facilities at the Center for Coastal and Ocean Mapping for running the experiments.

REFERENCES

BARTRAM, L., WARE, C. and CALVERT, T. 2003. Moticons: detection, distraction and task. International Journal of Human-Computer Studies, special issue on Notifications and Interruptions, 58(5), pp. 515-545.

BARTRAM, L. WARE, C. and CALVERT, T. 2001. Moving icons, detection and distraction, Interact 2001, Tokyo, July.

BARTRAM, L. and WARE, C. 2002 Filtering and brushing with motion. Information Visualization, vol 1 (1) 66-79.

BATTISTA, G.D. EADES, P., TAMASSIA, R., TOLLIS, I.G. 1999. Graph Drawing: Algorithms for the Visualization of Graphs. Prentice Hall.

BECKER, R.A., EICK, S.G., and WILKS, A.R. 1995. Visualizing network data, IEEE Transactions on Visualization and Computer Graphics. 1:1 March 16-28.

BRADDICK, O. and HOLLIDAY, I. 1987. Serial search for targets defined by divergence or deformation of optic flow. Perception, 20, 345-354.

BRANDES, Y., RAAB, J. and WAGNER, D. 2001. Exploratory network visualization: Simultaneous Display of Status and Connections. Journal of Social Structure 2(4) 1-28

CHEN, C., and PAUL, R.J. 2001. Visualization a knowledge domain's intellectual structure. IEEE Computer, 34, 65-71

DRIVER, J., MCLEOD P. and DIENES, Z. 1992. Motion coherence and conjunction search: Implications for guided search theory. Perception and Psychophysics. 51, 1, 79-85.

DRIVER, J., and MCLEOD, P. 1992. Reversing Visual Search Asymmetries with conjunctions of movement and orientation.

FAIRCHILD,.K.M., POLTROCK, S.E. and FURNAS, G.W. 1988. SemNet: three-dimensional graphical representations of large knowledge bases. In Cognitive Science and Its applications for Human-Computer Interaction. (R. Guindon. Ed) Lawrence Erlbaum, London, 201-233.

MCLEOD, P., DRIVER, J. and CRISP, J. 1988 Visual search for a conjunction of movement and form is parallel. Nature 332, 154-155.

MICHOTTE, C. 1963. The perception of causality, Methuen.

MIYASHITA, T. and UCHIDA, T. 1990. Cause of fatigue and its improvement in stereoscopic display. Proceedings of the Society for Information Display, 31,(3) 249-254.

MUNZNER, T., GUIMBRETIERE, F. and ROBERTSON, G. 1999. "Constellation: A visualization tool for Linguistic Queries from MindNet". Proceedings IEEE Information Visualization.

PARKER, G., FRANCK, G and WARE, C. 1998. Visualization of large nested graphs in 3D. Journal of Visual Languages and Computing. 9, 299-317

SOLLENBERGER, R.L. and MILGRAM, P. 1993. A comparative study of rotational and stereoscopic computer graphic depth cues. Proceedings of the Human Factors Society 35th annual meeting, San Francisco, CA, Sep. 2-6, pp. 1452-1456.

TREISMAN, A. 1988. Features and objects: The 14th Bartlett Memorial Lecture. Quarterly Journal of Experimental Psychology, 40A, 201-237.

TREISMAN, A. and GORMICAN, S. 1991. Search, similarity and integration of features between and within dimensions, Journal of Experimental Psychology, Human Perception and Performance. 17, 3, 652-676.

WARE, C., BONNER, J., CATER, R. and KNIGHT, W. 1992. Simple animation as a human Interrupt. International Journal of Human-Computer Interaction. 4(4) 341-348.

WARE, C. and FRANCK, G. 1996. Evaluating stereo and motion cues for visualizing Information Nets in Three Dimensions. ACM Transactions on Graphics. 15(2) 121-139.