Cross Display Mouse Movement in MDEs

Trina Desrosiers
Ian Livingston
Computer Science 481

David Noete
Nick Wourms
Human Computer Interaction

ABSTRACT
Multi-display environments are becoming more common in both homes and offices. Movement between displays is common in these settings and we should strive to optimize the action. The stitching technique causes the cursor to jump form one display to another, instantly, regardless of the displayless space between displays. In this paper we examine three techniques that attempt to provide benefits for the standard stitching technique used in most operating systems today. Two of the techniques apply a pseudo-haptic effect to the cursor, slowing the movement, when moving from one display to another. The last technique applies a visual cue to the cursor when crossing from one display to another. Each of these techniques is examined for different distances between displays. We carried out an experiment attempting to discover which technique is optimal for cross-display movement. We found that the visual cue technique was equivalent to stitching at all distance for task completion time, but both pseudo-haptic effects were significally slower. We also found that at medium and larger distances between displays all three techniques were significant in their reduction of target overshoot. Target overshoot occurs when the cursor is moved beyond the desired target location. These finding suggest that there is a practical advantage for providing a visual cue to users when passing from one display to another.

Author Keywords
Multi-Display environment, Stitching, pseudo-haptic, control/display gain, visual cursor feedback, targeting.

ACM Classification Keywords
HCI.

INTRODUCTION
The use of multi-display environments (MDEs) is increasing as the cost of computer displays decreases and users seek to maximize their productivity. Often new displays with different properties are added to a workspace. This results in discontinuity as size, resolution, location, and orientation of these displays change. As the use of MDEs increases and the additional display space is utilized, minor performance issues will become more apparent. Providing a performance increase in a MDEs will result in a benefit in the productivity of their users. As MDEs become more commonplace their optimization will become more important.

The stakeholders for MDE interaction optimization are a diverse group. Currently Multiple displays are commonly used for multimedia editing. The added workspace that is gained with MDEs, or larger displays, provides productivity gain [5]. MDEs appear as a result of connecting mobile devices, such as PDAs or tablet PCs, to a system that has an existing display. The speed and easy with which the user is able to interact with these additional displays will become more important as these devices and their interactions with each other become more common.

Our research extends work that looked exclusively at targeting across displayless space [13]. It was found that there is a positive correlation between the time required to complete a targeting task and the distance between the displays; essentially as the distance between displays increased so did the time required to complete the task. Another factor that was discovered was that as the distance was increased the number of overshoots, movement past the target location, also increased [13]. There has been work done which looks at alternative methods of representing MDEs. One such method, Mouse Ether, has been shown to provide a performance benefit [1]. Mouse Ether simulates the displayless space digitally between displays allowing the user to move through the space as if it were a continuous digital space. This technique has been shown to have decreasing performance as the distance between displays increases [13]. Stitching, where the system disregards the physical space between displays, is a common approach for handling multiple displays. As the space between displays increase, stitching provides better performance than other methods [13].

It should be noted that the display space has not changed; only the displayless space has been modified. The increase in task completion time occurred when the displayless space was both simulated (i.e. Mouse Ether) and not simulated (i.e. Stitching) [13]. This suggests that the effect is due to an external factor. It has been suggested that the main issues associated with MDEs are: visual loss of the cursor [5], window focus [5] and target overshoot [15] [13]. Our research focuses on improving the stitching method for both targeting task completion time and overshoot reduction.

Our experiment tests three potential MDE performance methods; two are pseudo-haptic effects (Constant slowdown and Scaled slowdown) and the other is the addition of a visual cue to the cursor when the display boundary is crossed. The goal for this experiment is to improve user performance for task completion time and decrease the number of overshoots performed. We believe that pseudo-haptic effects will reduce the amount of overshoot that occurs as the distance between displays increases. Also, we believe that our visual cue method will allow the user to reacquire the cursor faster thus decreasing the time required to complete the targeting task. In order to test this we have the user perform a number of targeting tasks at different display distances with each of our different methods.

Our results for the experiment showed that the application of either pseudo-haptic effect significantly increased the time required to complete the targeting tasks. The stitching with a visual cue method preformed as well as regular stitching, but not significantly better for task completion time. We were able to reproduce the findings in [13] where task completion time increased as the distance between displays increased. Our results do show a significant decrease in the number of observed overshoots for all methods, pseudo-haptic or visual cue, over regular stitching.

These results are not conclusive and require additional research to be carried out in order to determine if our techniques can provide a significant positive performance effect over the standard stitching technique. In the next section we will discuss, in more detail, work that is related directly with our research.

RELATED WORK

Multi-display Environment and Cursor Movement

MDEs are becoming more common everyday as displays become cheaper and smaller. In 2003 Mary Czerwinski et al. [5] stated that up to 20% of Windows OS users were using multi-monitor display environments. Also, it was shown that additional monitors provided more work space for users which lead directly to a decrease in the time required to complete tasks. This decrease in task completion time amounted to an increase in productivity of over 9%. While the benefits of MDEs are clear, usability concerns still remain. Some of these concerns are:

- Loss of the cursor, where the user visually loses track of the cursor across multiple displays [5].
- Window focus, where the user forgets to click the window in order to change focus from one window to another [5].
- Target overshoot, where the user overestimates the required input device movement to reach a target [13]

Research has shown that an increase in the gap between monitors results in an increase in speed of the cursor during movement [13]. This suggests that an observed physical distance between monitors is perceived by the user as a traversable space even when it is known that the cursor will instantly move from one screen to the other when the boundary is reached [13]. This is consistent with the described dominance of visual feedback outlined by M. A. Srinivasan [16] and reaffirmed by A. Lécuyer [10].

Researchers have developed various methods with which to deal with these concerns. For this paper we are interested only in improving the stitching technique currently in wide use in many MDEs. Stitching and a few other techniques are described below.

Stitching

Stitching is the most common technique used in MDEs. The edge of one display is “stitched” to another display. When the cursor reaches one boundary it instantly appears on the other screen completely ignoring any space that might exist between the displays. This technique does not account for any display-less space and windows placed across displays will be split between them with the entire window remaining visible.

Haptics

Haptic feedback is an interface mechanism that has been developed to provide users with simulated textural feedback in addition to visual feedback [3, as cited in 9]. One way to provide a haptic response is via a force feedback device. Such devices are commonly used in the computer game industry and can be found for current generation game consoles (i.e. XBOX 360, and Nintendo Wii). Haptic feedback could be used to help a user pass through the display-less space, found between displays, by signaling the boundary of a monitor. The value of such feedback has not been tested and is beyond the scope of this experiment.

Pseudo-haptics

Assuming that the user has an understanding of how an input device works when using it, there is an expected visual feedback associated with the input devices use (i.e. when the mouse is moved to the left, the cursor moves to the left on screen). Pseudo-haptics modifies the visual feedback of the input device to simulate haptic properties [9]. This technique can be used to imitate the effects of force feedback [11].

Mouse Ether

Mouse Ether takes into account the display-less space present between displays [1]. The display-less space is simulated in a Mouse Ether environment allowing the user to pass through it in a more accurate representation of the space. This eliminates the effect found in stitching where the cursor will instantly move from one display to another. Unfortunately, the major limitation of Mouse Ether is that the cursor is not visible while moving through the display-less space. As the amount of display-less space is increased the lack of visual feedback from the cursor makes this limitation more striking. For this experiment we will not test Mouse Ether since it has been found to be inferior to stitching as distance between displays increases [13].

Visual Cue Feedback

The use of a Visual Cue to provide a user with feedback is not a new idea. When working with MDEs or large displays it has been noted that visual tracking of the cursors location is a usability concern for researchers [5]. This is especially important for the Mouse Ether technique where the cursor is not displayed in the display-less space. This has lead to the exploration of visualization techniques that might aid the user in visually acquiring the cursor. Halo [13] [2] is a visualization technique that has been developed to aid the user in tracking object location in a display-less space. Another example is the use of arrows in wayfinding, which has been shown to be "sufficient to guide the user to his destination" [6]. In all cases the use of a visual cue has been shown to benefit the user by reducing the time required to complete task, though it has been noted that the visualization can be distracting [13].

Fitts' Law

The underlying model for our experiment relies heavily on the motor control and targeting movements of the user. Paul M. Fitts et al. quantified the tradeoff between speed and accuracy in his theorem Fitts' Law [8 as cited in 13].

$MT = a + b\log_2(2A/W)$ (1) Fitts' Law

$I_d = \log_2(2A/W)$ (2) Index of Difficulty

Fitts defined MT (movement time) as the average time required for completion of the movement with width W at a distance (or amplitude) A from the starting point to the center of the target. The value 'a' is the start/stop time respectively to the trials and b inherent speed of the device (i.e. Seconds/ bit). Lastly t is the average time per movements [7]. The Index of Difficulty specifies the minimum average information required for controlling or organizing the movements. This shows that the difficulty of a movement is related to the narrowness of a target space and the amplitude of the distance traveled by the input device [7]. MacKenzie and Buxton extended the models when it was found that Fitts' Law would often result in unrealistic values for the index of difficulty in two-dimensional tasks [12]. MacKenzie proposed:

$MT = a + b\log_2(A/W + 1)$ (4) Shannon Formulation

$I_d = \log_2(A/W + 1)$ (5) Index of Difficulty adjusted by the Shannon Formulation

This formulation was proposed because it fits better with observed results, maintains the quantified tradeoff between speed and accuracy of Fitts' Law and always provides a positive index of task difficulty [12].

Targeting and Movement

Rapid aimed movements can be divided into two phases of movements. The first phase consists of the initial aiming and movement towards the target area. The second phase repeats the first assuming the initial movement was unsuccessful in completing the task [14]. Woodworth was the first to define this two-phased movement model. [17 as cited in 13].

Control/Display Gain

There are two methods that are used in order to adjust control/display (CD) gain. The first method is constant gain (CG) where the input devices movement is adjusted by a constant multiplier. The second method is pointer acceleration (PA) where CD is adjusted based on the current movement characteristics such as movement velocity [4]. In our experiment we will only be examining CG, disabling PA for all trials. In order for the results of our experiment to be consistent we must set the CD ratio to an optimal value. This value must be set in order to avoid clutching, the action of lifting the input device from the surface and repositioning it so as to be able to continue a movement. Casiez et al. [4] determined that performance was only effect when the CG multiplier was very low, or when the CG multiplier was very high in conjunction with a small target region.

Smart Sticky Widgets

Smart sticky widgets are a concept where the movement velocity of input device is used to predict if the cursor should move freely over the boundary of a MDE or whether a pseudo-haptic should be applied to widgets close to the boundary [15]. This pseudo-haptic is applied to a widget, such as a scroll bar, if the cursor is determined to be moving over the widget at a specific velocity. The concept is similar to our experiment in that a pseudo-haptic [11] is applied close to the border of the screen. The smart sticky widget technique is a special case based on cursor velocity. In our experiment we hope to provide a more general solution that will aid not just boundary target acquisition [15] but target acquisition for all display areas of MDEs.

EXPERIMENT

Methods

Along with Stitching, we tested the following methods:

Stitching with a Visual Cue

This would use typical stitching with two lines (in a cross or X formation) extending in two constant angles following the cursor onto the upcoming monitor. The cursor would always remain at the intersection of these two lines. These lines would only appear when the cursor was a preset distance from the edge of the screen, as can be seen in Figure 1.

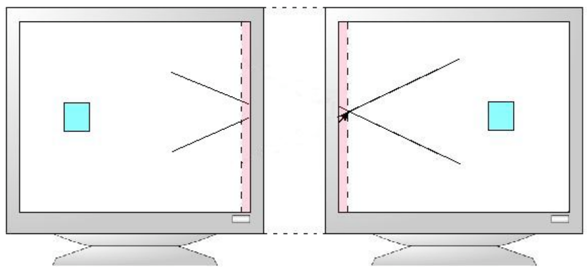

Figure 1. The red area represents the hotspot where the visual cue is displayed. If the cursor is within area the X appears. During the experiment the hotspot is not visible to the participant.

Stitching with a Fixed Slowdown on the Second Screen

This method will use regular stitching with a modified control-display ratio that slows down the effect of physical movement on the cursor. This speed decrease would occur for a preset amount of space after the cursor exits one screen and enters another, as can be seen in Figure 2. To determine how big the space should be and how much to slow the cursor down we performed a pilot experiment with 3 participants. We picked three widths (20, 45, 60) and three slowdown ratios (.75, .8, .95) based on what felt like reasonable values, and tried all combinations of them. Based on the results from this pilot we used a region width of 43 and a slowdown effect of .75 for the full experiment.

Stitching with a Variable Slowdown on the Second Screen

This method is similar to the previous except that the amount the cursor is slowed when entering the second screen is not constant; the factor by which cursor movement is slowed down increases as movement through the region occurs. When the user has moved through the whole slow down region the control display gain will have reached one. For the pilot experiment we tried widths of 20, 60, and 100 combined with slowdown ratios of .1, .25, and .4. From these results we chose a width of 55 and a slowdown ratio of .25 for the full experiment.

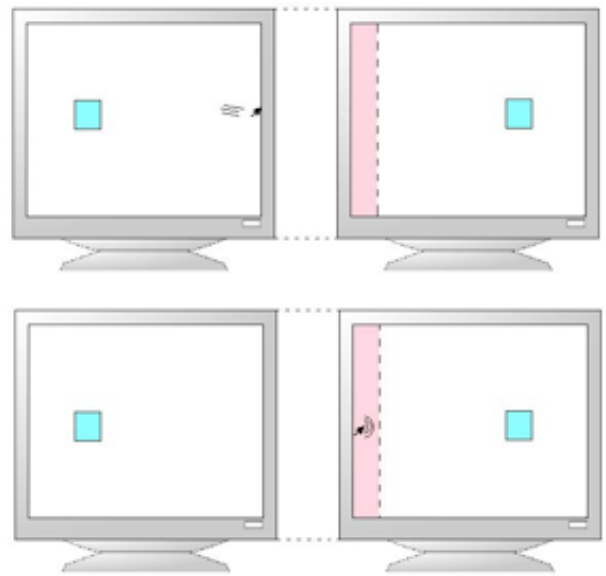

Figure 2. The red area represents the pseudo-haptic hotspot where the slowdown affect is applied. In order to reach the target the user must pass through this region. During the experiment the hotspot is not visible to the participant.

Hypotheses

Display-less space maps directly to physical space between displays. The space is not simulated in this experiment and we look solely at the MDE technique of stitching. We believe that the user, even when aware of the stitching technique, subconsciously tries to account for the display-less space during the ballistic phase of movement [17 as cited in 13]. We attribute this effect to the Gestalt principle of closure where by the user subconsciously views the display-less space between on displays as continuous and visually tangible. (H1) We believe that the pseudo-haptic effects, slow down and scaled slow down, as well as our visual cue will reduce the number of overshoots; and (H2) will reduce the amount of time it takes a user to complete the targeting task.

Movement of the cursor from one display to another along with increases in the size of display-less spaces leads to the user visually losing track of the cursor [5]. (H3) The use of our techniques (visual cue, pseudo-haptic slowdown) will a) cause task completion time and b) number of overshoots to decrease compared with regular stitching as distance between displays increases.

Apparatus

For this experiment we used a Logitech™ G5 2000dpi laser mouse. To ensure that participants would not be confused by the effects of the operating system and the effects of the experiment, we disabled mouse acceleration in the Windows™ Registry. Doing so also gave us more control over the cursors speed and thus more control over the experiment environment.

We used two ViewSonic™ VA912b 19" monitors with each having a 2cm wide bezel. Each monitor had the resolution set to the maximum 1280x1024. We chose these devices based on availability as well as the small bezel size.

The experiment software was written in Java 1.5. The software provided both audio and visual feedback when a participant clicked on the square to start a trial. When the participant clicked on the square to end a trial or missed the target, audio feedback was provided.

Participants

The age of the 8 participants in the experiment ranged from 23-34, with 5 being male and 3 female. All but one of the participants were computer science students.

Design & Tasks

The monitors were placed on an arc with a radius of 149.4cm. The distances between the monitors are in fact the number of degrees separating the monitors on the arc. For the shortest distance the monitors were placed next to each other with only the bezel of each separating the displayable space, which amounted to 1.5°. The other two distances correspond to 14.4° and 24.4° gaps. The monitors were set to the same height, as this would be a common case in real situations.

The paths used for the experiment were all left-to-right and were selected such that they comprised a wide range of movements between displays, as can be seen in Figure 3. Path 3 had both the start and the end targets on the same screen to avoid participants becoming accustomed to the left-to-right pattern.

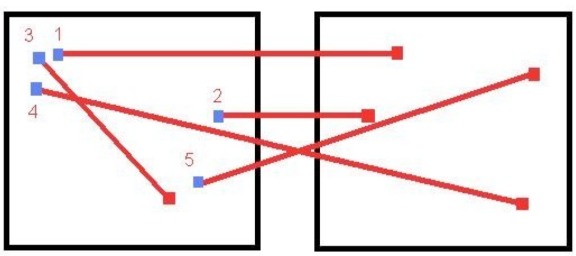

Figure 3. The five paths used for the experiment. Path 3 is a control path and is omitted from data analysis.

Each participant had a brief training session where they performed 10 trials (each path repeated twice) using each method at every distance. This gave them a chance to become acquainted with the environment and ask any questions they had regarding what was expected of them.

Once a participant had received the necessary training they were asked to perform a full round of trials. The participants used one method at all distances before moving on to the next method. The order the methods were introduced was based on participant number, with four possible choices; in our case we had two participants use each of the possible orderings. The order of distance was either ascending or descending, based on participant number. Each path was repeated 7 times for each method and distance. Within each method the five possible paths were ordered randomly such that each path would occur once over five trials. This was done to avoid the participants developing a pattern while performing the experiment. Based on the distance, techniques, paths, and repetition, each experiment participant performed 420 trials in total. The time to complete training and a full experiment was approximately 40 minutes.

RESULTS

To determine that results of our experiment we performed a repeated measures ANOVA on both overshoots and completion time. All trials of path 3 as well as all trials where the end target was missed were removed from the data. We will discuss the results as they relate to our hypotheses.

H1: Reducing overshoot

Figure 4 plots the average number of overshoots for each condition against the distance between the displays. At first glance it can be seen that there is variation in overshoots between conditions and over distance.

The results of the ANOVA revealed that there was a main effect of condition on overshoot ($F_{3,21}$=4.646, p=.012). Pairwise comparisons show that both Visual Cue (p=.022) and Scaled Slowdown(p=.009) produced significantly less overshoot than Stitching, while Slowdown had no significant effect (p=.077).

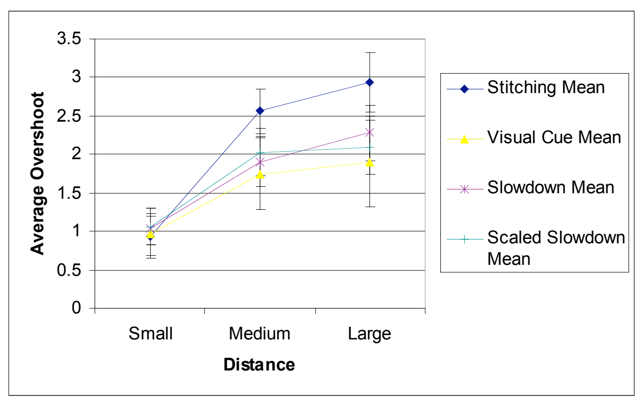

Figure 4. Target overshoot by display gap for each technique.

H2: Reducing completion time

The ANOVA revealed that there was a main effect of condition on completion time ($F_{3,21}$=6.508, p=.003). Pairwise comparisons show the Visual Cue (p=.841) is not significantly different from Stitching while Slowdown (p=.024) and Scaled Slowdown (p=.031) are significantly slower.

The average time in milliseconds for each method is plotted against distance in Figure 5. It can be seen that the average times of Slowdown and Visual Cue are similar and that Slowdown and Scaled Slowdown are similar at all distances.

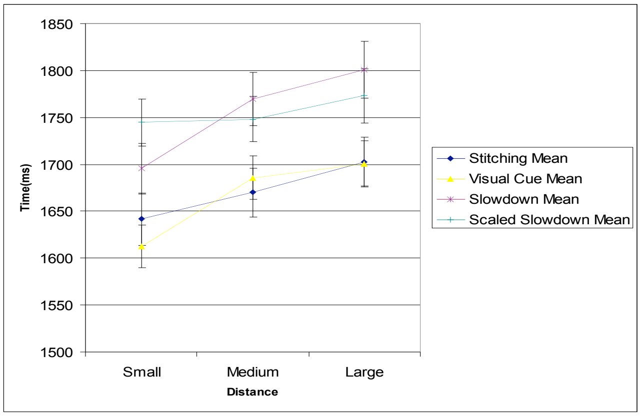

Figure 5. Completion time by display gap for each technique used. Error bars display standard error.

H3a: Reducing completion time over distance

Similar to previous work [13], there was a significant main effect of distance on completion time ($F_{2,14}$=17.482, p=.006). At the shortest distance Visual Cue (p=.307) and Slowdown (p=.173) were not significantly different than Stitching and Scaled Slowdown was significantly slower (p=.028). At the medium and large distances there were no significant differences between the new methods and Stitching.

H3b: Reducing overshoot over distance

Reaffirming previous work [13] we discovered there was a main effect of distance on overshoot ($F_{2,14}$=17.482, p =.000). At the shortest distance there were no significant differences between Stitching and the new methods. At the medium distance Visual Cue (p=.037), Slowdown (p=.036), and Scaled Slowdown (p=.038) all produced less overshoot than Stitching. At the largest distance, Slowdown (p=.111) was not significantly different than Stitching, while Visual Cue (p=.05) and Scaled Slowdown (p=.011) both produced less overshoot.

Other results

There is a significant two-way interaction between condition and path ($F_{9,63}$=17.482, p =.025) in the context of overshoot . Visual Cue significantly reduced overshoot for path 1 (p=.005), but for no others. Slowdown (p=.021) and Scaled Slowdown (p=.046) had significantly less overshoot on path 2, but had no effect on the others. It is interesting that these interactions occur on the paths where the target is furthest left on the target display. The two-way interaction between distance and path is also significant ($F_{6,42}$=15.021, p=.000) with respect to overshoot.

The two-way interactions between condition and path ($F_{9,63}$=1.180, p =.323) and between distance and path ($F_{6,42}$=15.021, p=.162) were not found to have any effect on completion time.

DISCUSSION

The discussion of our results will focus on the advantages of stitching over the pseudo-haptic techniques that were examined. We will also discuss the effects of decreasing overshoot.

Advantages of Stitching and Visual Cue

Based on our results Stitching and Stitching with a visual cue provide an advantage in task completion time over our pseudo-haptic effects. This is in conflict with our initial predicted results that the pseudo-haptic and visual cue techniques would decrease the task completion time. The pseudo-haptic slowdown and scaled slow effects resulted in an increase in task completion time. We believe that this increase in time was caused by the slowdown effect itself. Since the cursor is slowed while traveling through the hotspot the total time required to reach the target area is increased. This resulted in basic stitching and stitching with a visual cue to out perform the other techniques.

Overshoot reduction

Each of our new techniques resulted in a decrease in overshoot when compared to regular stitching as the distance between displays increased the visual cue and pseudo-haptic techniques always, with one exception, significantly reduced the number of overshoots that occurred (see figure 4). This result was expected though our prediction that overshoot was significant factor for the increase in task completion time at different distances appears to be false. We still believe that these two factors are connected, although the correlation seems to be less significant than initially predicted.

Distance and Time

Our results reconfirm the strong correlation between task completion time and the distance between displays. Our techniques did not result in the elimination or reduction of this correlation, which suggests that there are other key factors involved. Our experiment did show that there was a significant relationship between distance and the number of target overshoots but we didn't find a significant relationship between overshoot and time. This suggests that the additional time required to correct for overshoot is not significant and thus not a key cause of the increase in task completion time.

Recommendations

It is our belief that our techniques still hold promise in their application as a method to reduce the time required for task completion. We have two key suggestions to for future experiments.

Firstly, our techniques relied on a hotspot (see figure 1 & 2) to apply the pseudo-haptic affect to the cursor based on the number of mouse events received. This caused an inconsistency in the amount of slowdown the user perceived based on the velocity of the cursor as it passed though the hotspot. We would recommend applying a slowdown effect based on a period of time, rather than the time spent in a region.

Secondly, in our experiment our pseudo-haptic effect was applied at the same level for all distances. Previous research [13] has show, and confirmed in this experiment, that overshoot and task completion time both increase with the distance between displays. This suggests that the optimal CD ratio at one distance will most likely be different from another. It is possible that there is a linear relationship between the two and future research should consider this effect.

Future Work

In our cases it seems a good direction for future work would be to redo our experiment. For the second iteration we would recommend using experiment software that is closer to the operating system to allow for better mouse event control. We would also recommend that the second iteration have more control over the environment and participants.

Our results have shown that these new techniques have some credibility and hopefully a second iteration would reproduce this. Perhaps one direction future work could take is an examination of other visual cues that could be provided while transitioning between monitors.

It has been suggested that different areas of the screen provide different effects based on common usage. For example, the bottom of the monitors could provide a different effect than that of the middle. This might be beneficial for navigating to the controls of specific software, such as the Windows taskbar.

LESSONS LEARNED

This project has been a significant learning experience for all of us, and we will all take something away from it, whether we decide to go on to do research in the future or not. None of us appreciated how much time and thought goes in to successfully designing, conducting, and analyzing a research experiment. Gaining this experience now will certainly be beneficial to those of us going on to graduate work in computer science.

Doing a literature review of related work was a significant undertaking, especially since we came in to this project 'cold' with very little background knowledge. It is one thing to read a handful of research papers and quite another to understand everything that the authors are trying to get across. It takes a certain amount of skill and experience to analyze other's work to the point you are confident that you can utilize it and extend it; this experience should help us with that.

Setting up the experiment and the software for it were not as simple as we had expected it to be. Handling fringe cases of mouse movement, logging the correct information, and solving all of the minor details took longer than expected. Handling order effects, setting up the environment, and finding the time to have participants perform the experiment was also cumbersome. Part of the problem is that we did not reach this point until late in the term where people start getting busy with numerous deadlines looming.

One problem we had on the technical side of things was issues moving from Mac OS X to Windows with the experiment software. In hindsight we should have done all development on Windows because it was the platform we were expecting to use for our experiment. We did not keep track of how much time was spent on these issues, but it was significant.

By far the most difficult part of the project was analyzing the results of the experiment. Although most of us had learned about it in past courses, none of us had a firm grasp of how to run an ANOVA. We are very thankful for Dr. Mandryk's assistance with this portion of the project, as it is likely that we would have done it completely wrong without her. Even once we had run the data through SPSS we were still a not entirely sure what portion of the results were relevant for our purposes. We spent a significant amount of time trying to make sense of it all to determine what effects our techniques had. It probably would have been beneficial for us to spend some time reading up on these statistics before it came time to analyze our results.

CONCLUSION

Multi-display environments are becoming commonplace in both home and office settings. For this reason it is becoming important for designers and users to consider how they interact with these displays. The displayless space that is present between displays has been shown to affect users. Therefore we must strive to optimize user performance with movement techniques.

In our research we proposed and examined three new techniques that were applied to standard stitching. Two of these techniques were pseudo-haptic techniques, slowdown and scale-slowdown. The last technique applied a visual cue to the cursor. Our goal was to show that these techniques would decrease task completion time and overshoot as distance increased. We found that the pseudo-haptic techniques resulted in an increase in task completion time but generally decreased the number of overshoot when the distance between displays increased. The visual cue technique was comparable to the standard stitching technique for task completion time and decreased the number of overshoots as the distance between displays increased.

Our results suggest that a visual cue applied to the stitching technique is superior to regular stitching. Users were able to reacquire the cursor faster after crossing the displayless space. This allows the user to begin movement corrections faster resulting in less overshoot.

ACKNOWLEDGMENTS
We thank Professor Mandryk and Miguel Nacenta for their help with the various phases of this project.

REFERENCES
1. Baudisch, P., Cutrell, E., Hinckley, K. & Gruen, R. Mouse Ether: Accelerating the Acquisition of Targets Across Multi-Monitor Displays. *Proc. CHI'04*, 1379-1382.
2. Baudisch, P. & Rosenholtz, R. Halo: A Technique for Visualizing Off-Screen Locations. *Proc. CHI'03*, 481-488.
3. Burdea, G. *Force and Touch Feedback for Virtual Reality*, John Wiley & Sons, New York, (1996).
4. Casiez, G., Vogel, D., Balakrishnan, R. & Cockburn, A. The Impact of Control-Display Gain on User Performance in Pointing Tasks.
5. Czerwinski, M., Smith, G., Regan, T., Meyers, B., Robertson, G. & Starkweather, G. Toward characterizing the productivity benefits of very large displays. *Interact 2003*, 9-16.
6. Ekman, I. & Lankoski, P. What should it do? Key issues in navigation interface design for small screen devices. *CHI'02 Extended Abstracts*, 622–623.
7. Fitts, P.M. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47, (1954), 381-391.
8. Fitts, P.M. & Peterson, J.R. Information capacity of discrete motor responses. *Journal of Experimental Psychology*, 67, (1964), 103-112.
9. Lécuyer, A., Burkhardt, J.M. & Etienne, L. Feeling bumps and holes without a haptic interface: the perception of pseudo-haptic textures. *Proc. CHI'04*, 239-246.
10. Lécuyer, A., Coquillart, S., Burkhardt, J.M. & Coiffet, P. Boundary of Illusion: an Experiment of Sensory Integration with a Pseudo-Haptic System. *IEEE Virtual Reality Conference 2001*, 115.
11. Lécuyer, A., Coquillart, S., & Kheddar, A. Pseudo-Haptic Feedback: Can Isometric Input Devices Simulate Force Feedback? *IEEE Virtual Reality Conference 2000*, 18–22.
12. MacKenzie, I.S. & Buxton, W. Extending Fitts' law to two-dimensional tasks. *Proc. CHI'92*, 219-226.
13. Mandryk, R., Nacenta, M. & Gutwin, C. Targeting across Displayless Space.
14. Meyer, D., Abrams, R., Kornblum, S., Wright, C. & Smith, J. Optimality in human motor performance: Ideal control of rapid aiming movements. *Psych. Review*, 95, (1988), 340-370.
15. Rodgers, M.E., Mandryk, R.L. & Inkpen, K.M., Smart Sticky Widgets: Pseudo-haptic Enhancements for Multi-Monitor Displays. *Proc. of Smart Graphics'06*, 194-205.
16. Srinivasan, M.A., Beauregard G.L. & Brock D.L. The Impact of Visual Information on Haptic Perception of Stiffness in Virtual Environments. *Proc. of ASME Winter Annual Meeting*, (1996).
17. Woodworth, R.S. The accuracy of voluntary movements, *Psych. Review*, 3, (1899), 1-114.