Targeting across Displayless Space

Miguel A. Nacenta, Regan L. Mandryk, and Carl Gutwin
Department of Computer Science, University of Saskatchewan
110 Science Place, Saskatoon, SK, S7N 5C9, Canada
{nacenta, regan, gutwin}@cs.usask.ca

ABSTRACT
Multi-monitor displays and multi-display environments are now common. Cross-display cursor movement, in which a user moves the pointer from one display to another, occurs frequently in these settings. There are several techniques for supporting this kind of movement, and these differ in the way that they deal with *displayless space* (the physical space between displays). Stitching is the method used by most operating systems; in this technique, the cursor jumps from the edge of one display directly into the next display. In contrast, Mouse Ether maps the motor space of the mouse exactly to the physical space of the displays, meaning that the cursor has to travel across displayless space until it reaches the next display. To determine which of these approaches is best for cross-display movement, we carried out a study comparing Stitching, Mouse Ether, and a variant of Mouse Ether with Halo for off-screen feedback. We found that Stitching is equivalent to or faster than any variant of Mouse Ether, and that Halo improves Ether's performance (but not enough to outperform Stitching). Results also indicate that the larger the gap between displays, the longer the targeting takes – even for Stitching. These findings provide valuable guidance for practitioners and raise new interesting questions for research.

Author Keywords
Multi-display environments, multi-monitor, Mouse Ether, Stitching, off-screen feedback, Halo, displayless space.

ACM Classification Keywords
H.5.2 Information Interfaces and Presentation: User Interfaces - *Interaction Styles; Input devices and strategies.*

INTRODUCTION
Multi-display environments (MDEs) are becoming common due to the reduced cost of new displays. Interfaces for MDEs are fundamentally different from their single-display counterparts and require a fundamentally different outlook on design; in particular, the fractured nature of display space in multi-display systems (display surfaces can be of different sizes, resolutions, and are separated from each other by physical gaps) calls for a specific look at how these discontinuities affect common interaction tasks.

Current multi-display interfaces are direct adaptations of single-display designs and therefore tend to ignore gaps between surfaces. For example, in current operating systems and multiple research prototypes [18, 16, 9] the cursor 'warps' directly into another display once it has crossed the boundary of the original display. In other words, the display spaces are "stitched" to each other in the motor space: mouse movement displaces the cursor from one display to another as if they were perfectly contiguous, completely ignoring any physical gap between them.

Baudisch et al. [5] pointed out that Stitching disrupts the relationship between motor and visual spaces (i.e., between how the mouse is moved and how the cursor changes position in the physical world). The mismatch between movements in motor space and visual feedback from the cursor movements on the display makes feedback more difficult to interpret during mouse motion, and may result in a reduction of targeting performance (see Figures 1 and 2). To reduce this effect, Baudisch and colleagues proposed Mouse Ether, a technique that accounts for the space between displays and corrects for resolution differences. With Mouse Ether, the cursor does not warp from the border of one display to the next; instead, the mouse has to be further moved through displayless space until the cursor crosses the gap and appears at the destination display.

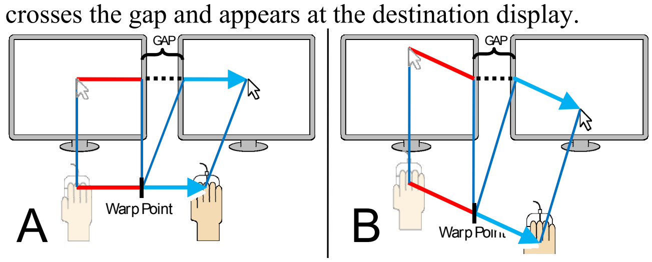

Figure 1. Inconsistency between motor space and visual feedback with Stitching. A) The gap is ignored in motor space (gap is compressed into the warp point). B) A diagonal motion is transformed into a multi-linear trajectory.

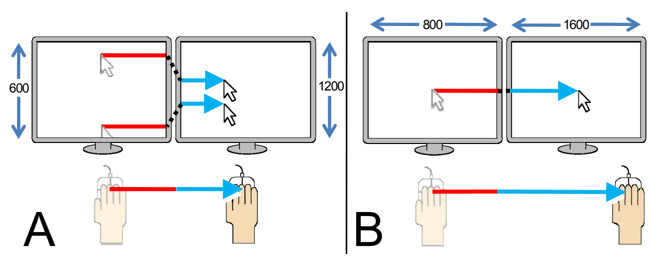

Figure 2. Motor-visual inconsistency due to resolution differences. A) Alignment of all trajectories is impossible B) Magnitude of the required movement depends on the display.

In Mouse Ether's initial evaluation study [5], performance improved up to 28% for a setting of two monitors of different resolutions with a small gap between them. However, Mouse Ether still has limitations that were not tested in the initial study – users might lose the cursor, and it is not clear whether the lack of feedback in the displayless space affects the targeting process in other situations such as displays separated by larger gaps, or when resolutions differences between displays are corrected for in motor space. Current targeting models (e.g., [13, 29, 1]) do not help much in determining the performance of different cross-display techniques because discontinuity or fidelity of feedback are not explicitly considered.

As multi-display PCs and smart rooms with heterogeneous arrangements of displays become ubiquitous, multi-display interface designers will be compelled to find the appropriate techniques for cross-display targeting in each situation. In order to help designers make informed decisions, we designed and executed a study that investigates which cross-display technique (Stitching or Mouse Ether) provides the best targeting performance for different gap sizes.

Our study takes the original Mouse Ether experiment as a starting point and looks separately at the effects of corners, motor-visual space relationship and off-screen feedback at difference gap distances. Through the study we arrived at four main conclusions, each of which is relevant for the design of multi-display systems:

- Mouse Ether did not improve performance in any configurations that we tested.
- Ordinary Stitching was significantly better than Mouse Ether at medium and large gaps.
- Halo (an off-screen cursor visualization [7]) led to significant improvements with Mouse Ether.
- There was a significant performance cost to moving monitors further apart – even for Stitching, which does not involve any additional movement in motor space.

These results provide valuable guidance to designers of MDEs and uncover a new set of important directions for future research such as the possibility of accounting for displayless space in current targeting models.

In the following section we present a review of existing options for multi-display pointer movement, including a detailed analysis of Mouse Ether and off-screen feedback mechanisms. Then we survey current targeting models and theories that might help understand the role of displayless space and continuity in the targeting process. A rationale for our study follows, which leads us to the description of the empirical design, its results, and a discussion thereof. Finally we end with a set of recommendations for practitioners and a set of directions for further research.

RELATED WORK

Our study is concerned with three main streams of research: cross-display cursor movement mechanisms, off-screen feedback, and targeting models.

Cross-display cursor movement

A mechanism for moving the cursor from one display to another is fundamental to multi-display interfaces (e.g., to move content or act on objects across displays). Accordingly, researchers have proposed a number of novel techniques. For this study we are concerned only with mechanisms that make use of indirect input – in particular, techniques that rely on the mouse for pointing – because indirect input is arguably the most common way to control the cursor. We divide existing techniques into four groups: Stitching and Cursor Warping, Mouse Ether, and Portals.

Stitching and Cursor Warping

A user's visual field can be divided into two regions: space occupied by displays and displayless space. Although displayless space is irrelevant in single-display systems, multi-display environments are fractured, and displayless space fills the visual space between displays.

Stitching is the default way of dealing with displayless space: the motor-space mapping simply ignores gaps between displays. Note that displayless space is not accounted for in motor space, but it still forms part of the visual space that lies in front of the user, causing a mismatch between user action and user perception. Stitching is the model used by current operating systems (e.g., Windows™, MacOS™, XFree86®), input redirection applications such as Synergy (synergy2.sourceforge.net), and some current research systems [9, 16, 18].

Cursor warping is a variant of Stitching in which the cursor can jump from any location in a display to another location in a different display. For example, Multi-Monitor Mouse [8] has several modes in which the cursor jumps to a specific location in another monitor with the click of a button or key. Other varieties of Cursor Warping provide implicit activation of the warping mechanism: Delphian Desktop [2] calculates the destination of the cursor using peak velocity data during pointing gestures; Object Pointing [14] allows the cursor to ignore display space between objects. Although these two techniques have been designed for single-display desktops, they could as well be applied to multi-display interaction due to their 'warping' nature.

Finally, some researchers have applied head tracking and eye tracking to multi-display cursor switching [12, 8, 3]. Although these approaches show promise, they require equipment that is still expensive or cumbersome.

Some of the strategies described above have been proven useful. We will, however, test only a simple version of Stitching because we are interested in the basic mechanisms and problems of warping; all the aforementioned techniques are either equivalent to Stitching in their discontinuity, or bring in other problems that fall outside the scope of this research (e.g., costs in performance of a mode switch).

Mouse Ether

Mouse Ether [5] was designed to solve some problems of current Stitching implementations. By taking into account the actual size of displays and the space between them, Mouse Ether provides a more accurate representation of the physical environment in motor space (i.e., visual space and motor space match better).

Mouse Ether has four advantages over ordinary Stitching:

- A better match between visual space and motor space avoids the potentially distracting jumps and trajectory inconsistencies shown in Figure 1.
- The visual-motor consistency results in cross-display movements that are consistent with movement within a display, allowing for more natural cross-display transitions (i.e., moving across displays is more like moving within displays than with Stitching).
- The correction in motor space of resolution differences between displays avoids Control-Display (CD) gain discontinuity problems (e.g., Figure 2).
- The existence of virtual space around some parts of the displays allows users to 'cut corners'. Some required movements become rectilinear, and therefore shorter than their Stitching counterparts (see Figure 3).

However, Mouse Ether has an evident drawback: the cursor is invisible when it is in displayless space, and the user lacks visual feedback on its position.

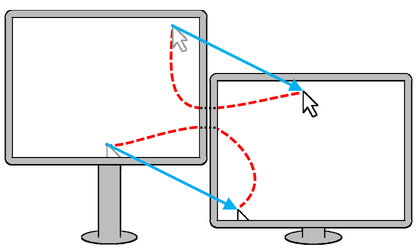

Figure 3. Two examples of cutting corners with Mouse Ether (blue arrows) and the usual trajectories for the same result with Stitching (red dotted lines).

So far, Mouse Ether has been evaluated twice. In its original evaluation [5], Mouse Ether was tested for a single (short) distance, and with two monitors set to different resolutions. The results from the study indicated a significant performance improvement over Stitching for most of their targeting tasks.

In [25], Nacenta and colleagues evaluated Perspective Cursor, a cross-display technique that generalizes Mouse Ether for environments with displays at different angles. Perspective Cursor maps the motor space to a spherical representation of the physical space around the users and allows the movement of the cursor across displayless space. Their evaluation showed an advantage of Perspective Cursor over Stitching and Laser Pointing, but it is not clear whether this is due to the existence of Mouse Ether between displays or to the novel motor-visual mapping.

Portals

Portals are a completely different way to provide access to other displays. Portals replicate the content of a different display and allow manipulation of the content through the replica, without having to displace the cursor to the original objects. Many versions of portals exist that can be more or less sophisticated (see, for example [28, 27, 19, 6]); we do not analyze portals because they present different issues than cross-display cursor movement (e.g., setting up the region of interest, collaborative awareness problems), although designers should take these techniques into account when designing multi-display systems.

Off-screen feedback

Several techniques have been proposed for showing the presence, direction, and distance of off-screen objects. The simplest technique, used often in video games, is an arrow on the boundary of the screen that points to the off-screen object. Additional cues can be added to the arrow to convey more than just direction; this is seen in the City Lights technique [22], which represents off-screen objects with blocks on the edge of the display. The size, shape, and colour of the blocks indicate different properties of the object. Arrows and City Lights, however, have difficulty showing the distance to off-screen objects. Halo [7] was designed to address this problem, by explicitly representing both distance and direction. Halo shows an arc on the edge of the view for each off-screen object; the arc is centred on the object, so the radius of the arc indicates the object's distance and the location on the display indicates direction.

Models for targeting

In 1899, Woodworth [29] performed a series of experiments to investigate the speed-accuracy tradeoff in aiming movements. Not only did he originally identify the speed-accuracy tradeoff, but he also proposed that aiming movements are comprised of two phases: the initial impulse towards the target, and a deceleration phase under current control to home in on the target. It was 65 years later that Fitts and Peterson [13] quantified the speed-accuracy tradeoff for discrete aiming movements in Fitts's Law. More recently, Fitts's Law was adapted to model virtual aiming in two-dimensional environments [20], and has been used extensively in HCI to evaluate input devices.

There are various underlying models of motor control for aiming movements, but the most successful model to date corresponds with Woodworth's two-phase model of aiming. The optimized initial impulse model [23] suggests that an initial ballistic movement is made towards the target, and if successful, then the task is accomplished. If not, a secondary movement is undertaken and this process repeats until the target is acquired. To optimize aiming movements under the constraints of the speed-accuracy tradeoff in this model, most movements consist of a high-velocity initial phase (ballistic phase), and a series of slower, visually-guided, feedback-corrected movements (homing-in phase). Prior to the optimized initial impulse model, other motor control models emphasized the ballistic phase (impulse variability model [26]) or the homing-in phase (iterative corrections model [11]), but these models are too narrow to account for all of the experimental results.

MacKenzie et al. [21] examined the 3D velocity profiles of discrete aiming movements and showed that for movements of the same index of difficulty (i.e., same movement time), the timing and magnitude of the peak velocity (ballistic phase) related to the movement amplitude, while the proportion of time spent decelerating (homing-in phase) depended on the target width.

Closed-loop and open-loop aiming

In the homing-in phase, we use vision to guide our movements in a closed-loop manner. Visual feedback of the end effector's movements allows us to adjust our trajectories. In open-loop pointing, no feedback is provided, which allows the motor sequence to be carried out uninfluenced by corrective movements. Although feedback can occur via other sensory channels, (e.g., proprioception, auditory, haptic), we will use open-loop throughout this paper to describe aiming without visual feedback. The open- and closed- loop descriptions have obvious correlates to the two phases of aiming; however, the two-phase model fails to describe what might happen when feedback is variable throughout an aiming movement.

CD gain

In motor control studies, the relationship between motor space and display space is one to one. When aiming on a computer display, the relationship between control and display (CD gain) can be manipulated. With static CD gain, the cursor can move at a different speed than the hand. With dynamic CD gain, (e.g., mouse acceleration), the relationship between movements in motor and display space is varied within an aiming movement. CD gain manipulation has been effectively used to enhance targeting on a computer (see [4]).

EMPIRICAL STUDY

As we have seen in the previous sections, there is uncertainty about which way of dealing with displayless space is better in which situations. To investigate the differences between the most common cross-display techniques, we designed an experiment that tests Stitching and Mouse Ether with different gaps between displays.

Early in the design process of the experiment, we realized that the main drawback of Mouse Ether (the lack of feedback when the cursor is in displayless space) can be remedied by using off-screen feedback; therefore we decided to include a condition in which Mouse Ether was assisted with Halo. We selected Halo [7] as off-screen feedback because it is simple to understand and implement, and requires very little interpretation by novice users.

Unlike in the initial Mouse Ether study, we deliberately excluded resolution differences from our study. The problems of changes in CD gain caused by differences of resolution are not intrinsic to Stitching or Mouse Ether; Stitching setups can easily be corrected to provide a homogeneous CD gain by taking the monitors' resolution (pixels per centimeter) into account.

Questions

The study was designed to answer four specific questions that were stated before the data analyses:

Q1: Which cross-display technique is best for a particular distance between displays?

Q2: Does Halo help performance in Mouse Ether?

Q3: What is the relation between gap size and performance when using Stitching?

Q4: Is the ability to cut corners an advantage of Mouse Ether over Stitching?

These questions correspond to specific statistical tests that were designed before data collection.

Cross-display techniques

The experiment tested three cross-display techniques: Stitching, Mouse Ether, and Ether+Halo.

Stitching

Our implementation of Stitching replicates that of most current operating systems. The cursor jumps from one display to another when crossing a display boundary that is contiguous to some other display. If the cursor reaches a display boundary that is not connected to any other display, the cursor stays in the same display. When hitting the boundary, components of the input movement that are parallel to the screen boundary are still applied to the movement of the mouse (i.e., the cursor can still "slide" through the boundary – see Figure 4).

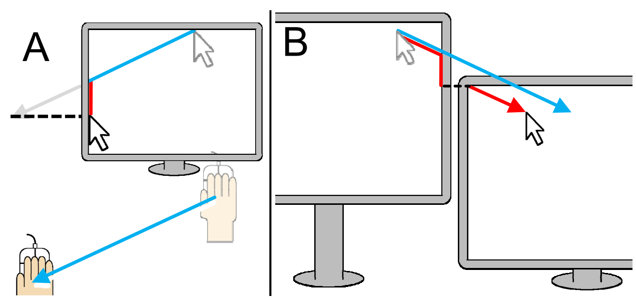

Figure 4. A) The outer boundaries of the display environment allow the cursor to "slide" (movement is right to left). B) The boundaries in Stitching prevent cutting corners and result in a partial loss of horizontal movement when sliding through the corner boundary (red arrow) with respect to the intended movement (blue arrow). Note that the vertical component of the movement is preserved because of the sliding mechanism.

In our implementation, the difference of alignment between screens is accounted for so that a perfectly horizontal movement of the mouse does not change the vertical position of the cursor when changing displays (Figure 5).

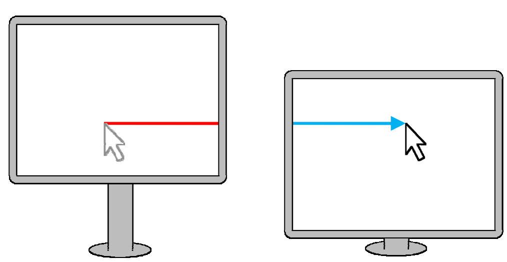

Figure 5. For all techniques, the vertical alignment was taken into account so that horizontal movements of the mouse produced horizontal cursor movements without height jumps.

Mouse Ether

Mouse Ether was implemented according to the description in the original paper by Baudisch et al. [5]. The space gap between monitors was carefully measured and accounted for in the input model. Vertical misalignment of the displays was also included in the model to avoid unnecessary vertical 'jumps'.

As in the original implementation, the cursor could be moved in the ether until it reached the convex hull of the two displays (see Figure 6). Notice that this implementation restricts the cursor position but allows cutting corners. Boundaries belonging to the convex hull behaved as in Stitching (i.e., the cursor would "slide" along the boundary as long as the mouse movement contained that directional component).

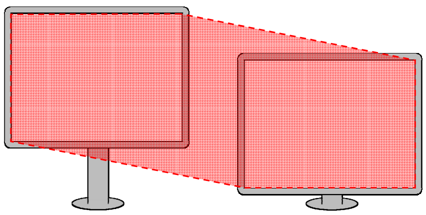

Figure 6. Limits of Mouse Ether as proposed by Baudisch and colleagues [5] (convex hull of display surfaces).

Ether+Halo

Ether+Halo behaved exactly like Mouse Ether except that a red circular halo appeared on screen when the cursor was in displayless space. The halo had a thickness of 3 pixels and was set to appear in at least one display at any time that the cursor was not visible. Halo intrusion borders (the area contiguous to the boundaries that can show a halo) covered a 200px (5cm) thick framework from the borders of each screen (see Figure 7).

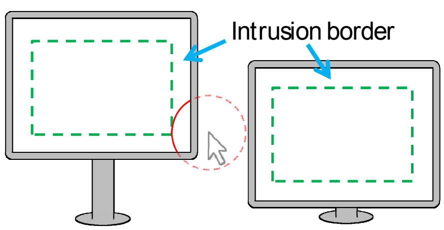

Figure 7. A halo (continuous red line) is always visible in at least one display when the cursor is out of display space. The intrusion border is the area of a display that can show halos.

Apparatus

The experiment was conducted on a Windows™ Pentium IV™ machine connected to two identical 20" monitors (40.6cm by 30.5cm, 1600 by 1200px). The monitors were set at different heights (13.24cm offset) to test for the effect of corners. During the experiment, the monitors were moved between three different positions to test three different gaps between displays (Points A, B, and C in Figure 8). The monitors were arranged along a circumference of radius 90cm centered on the estimated position of the participant's head; this kept the visual distance to the monitors constant for all gap distances and avoided possible perspective confounds (see [24]).

The cursor was controlled by a Logitech G5™ Laser mouse, with a resolution of up to 2000 DPI. The experimental software avoided any quantization – a small movement of the mouse could produce a cursor movement of 1 pixel on the screen. Windows acceleration was disabled by using a flat curve in the registry variable. The CD gain was, therefore, constant at a value of 16 (640px/cm); a relatively high setting that avoided the need for clutching, even at the largest distances.

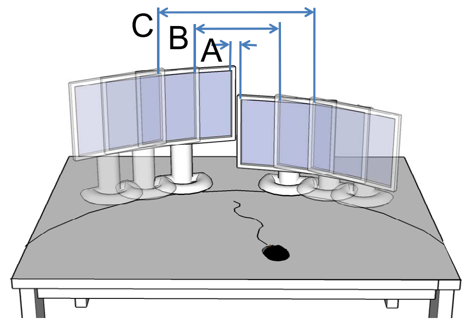

Figure 8. Experimental set-up with three different gaps between monitors.

We chose to avoid clutching because differences in clutching behavior across the experiment and between subjects could be a source of unwanted variability and produce confound effects (see [10]). We disabled acceleration because it implies an extra source of motor-visual space mismatch that could negatively affect the performance of Mouse Ether.

The experimental software was built in C# for .NET 2.0. No delay was noticeable and the frame rate was above 24 frames/second.

Design and Tasks

The experiment followed a 3x3x10 (three factor) within-subjects full-factorial design with cross-display technique, gap, and path as factors. The main factor was cross-display technique (Stitching, Mouse Ether, or Ether+Halo). The gap was either small (the minimum allowed by the monitors' bezels), medium, or large (positions A, B and C in Figure 8), which corresponded to gap angles of 2.4º, 14.4º and 26.4º (3.8cm, 23cm and 42cm along the circumference).

The task consisted of clicking in a blue square and then on a larger red square (4cm side, 158pixels). The initial and target squares were placed according to 10 possible paths designed to test different kinds of situations (see Figure 9).

All paths were left-to-right and diagonal paths were downward; pilot tests showed that completion times differ depending on direction, but direction did not interact with our variables of interest. Therefore it is reasonable to assume that right-to-left and bottom-top motions are affected similarly by cross-display technique and gap.

Paths 0 and 2 were designed to test the effect of corners – rectilinear trajectories are possible with Mouse Ether, but not with Stitching. Path 1 is equivalent to 0 and 2 except that it does not come near any corner.

Paths 8 and 9 represent single-display tasks and serve both as baseline comparisons between single- and multi-display performance and as distracters to avoid rhythmical completion of the tasks.

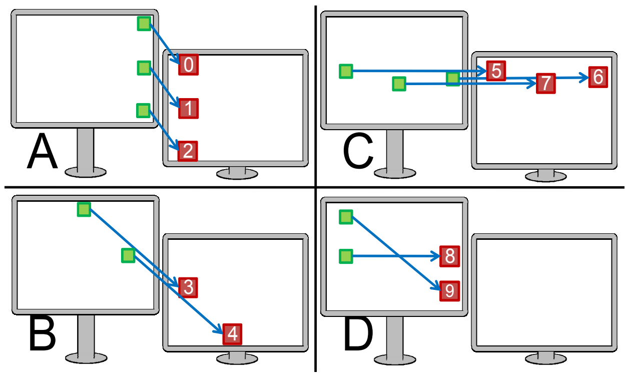

Figure 9. Paths used in the experiment. Note: representation is not to scale and paths 4, 6 and 7 have a vertical offset to improve visibility in the figure.

The main dependent variable of the study was completion time, which was measured from the initial click in the start square until the click in the target square. The software also registered errors, overshoots, times that the cursor entered the target, peak velocity, location of the peak velocity and time-stamped trajectories of the motion; these data, however, are only used to further understand relevant results from the planned tests.

Participants and Procedure

Twelve subjects – 6 females and 6 males, aged between 21 and 33 years – participated in the study for a $10 honorarium. 4 participants had significant experience with multi-display environments and 4 had never used a multi-display system.

The study took approximately 50 minutes to complete. Participants completed consent forms and demographic questionnaires, and were then trained in each of the three cross-display techniques for each of the three possible gaps (30 trials per technique, 90 trials in total). They then performed blocks of 50 trials for each of the cross-display technique/gap combinations for a total of 150 trials per technique and a total of 450 trials in the main blocks.

The order of the inter-display mode condition was fully balanced across subjects (two participants in each possible order). Half of the participants performed the tests with increasing gaps (short, medium, long gaps) and half in the opposite order. In each block, the different paths appeared in a randomized order. The first 10 trials (one per path) of each block were marked as training and excluded from the main analysis to avoid noise and adaptation effects. Trials with errors (e.g., clicking outside the target) were repeated.

After all trials were completed, the participants filled in a post-study questionnaire with specific questions about preference, speed and accuracy of the cross-display techniques for each of the gap distances.

Results

We report our results in three sections: the planned quantitative analyses that correspond to the main questions, the analysis of the subjective data from the post-study questionnaire, and the explanatory analyses.

All analyses were performed on error-free trials that were not marked as training. Four points per user for each cell (same cross-display technique, gap distance and path) were collected (a total of 3*3*10*4 = 360 valid data points per user).

Planned quantitative analysis

An omnibus factorial ANOVA with participant as random factor revealed that, as expected, there were significant main effects of cross-display technique ($F_{2,22}$ = 16.4, $p < 0.001$), gap distance ($F_{2,22}$=283.1, $p<0.001$) and path ($F_{7,77}$=34.7, $p<0.001$). The interaction between cross-display technique and gap distance was also significant ($F_{4,44}$=13.6, $p<0.001$).

To answer question one (Which cross-display technique is better at which gap distance?) and question two (Does Halo help?) we performed planned pair-wise comparisons for each gap distance. All pair-wise comparisons in this section were corrected using the Games-Howell procedure for unequal variance data (analogous to Tukey's HSD).

Short gap. Ether+Halo was the fastest technique (735ms avg completion time), followed by Ether and Stitching (762ms and 768ms respectively). The pair-wise comparisons only show significant completion time differences between Ether+Halo and Stitching ($p < 0.008$) but this represents a small difference (33ms, 4.3%).

Medium gap. Stitching was the fastest technique (865ms avg completion time), followed by Ether+Halo (937ms) and Mouse Ether (995ms). All pair-wise comparisons were significant (all $p < 0.003$). The difference in completion time represents a performance advantage of Stitching of 7% and 13% with respect to Ether+Halo and Mouse Ether.

Large gap. Stitching was again fastest (958ms), followed by Ether+Halo (1100ms) and Mouse Ether (1222ms). The ordering was the same as for the medium gap and the differences were also statistically significant (all $p<0.001$), although proportionally larger; Stitching was 13% faster than Ether+Halo and 21% faster than Mouse Ether.

These analyses and Figure 10 allow us to give explicit answers to the first two questions:

Q1 *(Which technique is better at which distance?):* For the medium and large gaps, Stitching is superior to other techniques. For short gaps, the three techniques are roughly equivalent, although there is a slight advantage of Ether+Halo over Stitching.

Q2 *(Does Halo help?):* Ether+Halo was faster than Ether at all distances. The effect of Halo is not statistically significant for the short gap, but the improvements of Halo on movement time are statistically significant for the other two gap distances.

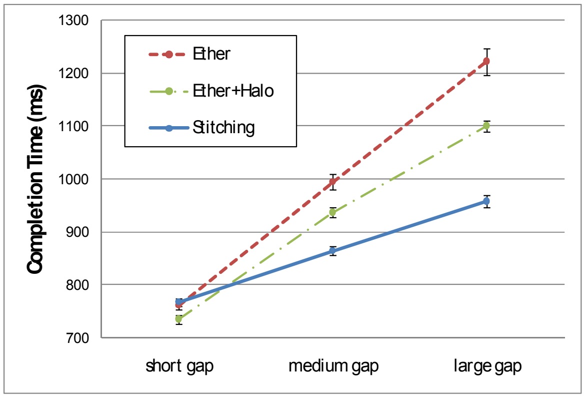

Figure 10. Completion time by gap and technique, paths 8 & 9 excluded (error bars display Standard Error).

To answer question three – relationship between gap distance and performance for Stitching – we planned two comparisons; short to medium gap and medium to large. Both comparisons were significant (p<0.001). This relationship is easiest to see in Figure 11, which compares targeting times in Stitching for paths 3 and 4 with path 8, which is 4 cm longer but does not cross the gap.

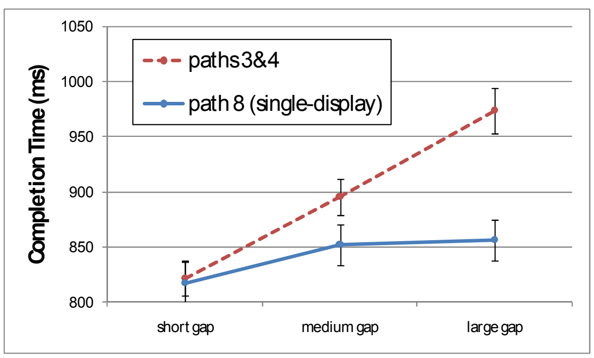

Figure 11. Completion time by gap for paths 8, 3 and 4 with Stitching (error bars display Standard Error).

Q3 *(What is the relationship between gap and completion time with Stitching?):* Although the distance traveled by the mouse in motor space is constant for the Stitching technique at all gaps, the completion time increases with the size of the gap.

Question four is answered by looking at a comparison between paths 1 and 3 (which might benefit from cutting corners) and path 2 (which does not require cutting any corner). In our analysis, all cross-display techniques show completion times similar for paths 1, 3 and 2. A factorial ANOVA with participant as a random factor and using data exclusively from paths 1, 2 and 3 did not show any interaction between path and cross-display technique ($F_{2,22} = 3.423$, p>0.05).

Q4 *(Is cutting corners an advantage of Mouse Ether?):* no significant improvement in target times was observed between techniques in the corner paths, contradicting the hypothesis that cutting corners provides a significant advantage for Mouse Ether. Moreover, these paths show the same pattern than the rest: Stitching is faster or indistinguishable from other techniques.

Across the experiment, errors (i.e., clicking outside targets) were below 6% for any given technique-gap combination; errors were evenly distributed across techniques (5.3% Stitching, 4.7% Ether, 5.5% Ether+Halo).

Subjective Data

The post-study questionnaire asked participants to rank the three cross-display techniques for each of the gaps according to three criteria: personal preference, speed and accuracy (a total of 9 rank sets). Non-parametric Friedman analysis for each of the gap-criteria combinations show a marginally significant difference in the short gap ($\chi^2(2) = 6.16$, $p < 0.046$); only one out of nine tests. No α-level adjustments were performed.

Responses from participants varied widely, as is reflected in their comments; some participants found that Mouse Ether was "more natural" and that "the cursor comes up where I expect it to be". However, other participants stated that "[Stitching] is faster" and that "[in Mouse Ether] without the halo I became lost once or twice". Some participants liked the halo because "[it] helped me get unlost" while others found it "annoying", "distracting" and "confusing", especially for the short gap.

Remarkably, some participants stated that "[with Stitching] I felt lost sometimes" and "the distance is still too far and seems to jump too fast".

Explanatory Analyses

During the study we logged several other measures besides completion time that can help elucidate our findings by revealing user behavior in the context of models for targeting. Using the dependent measures of peak velocity, time to peak velocity, percent time after peak velocity, display where peak velocity occurs, and presence of an overshoot error, we conducted factorial ANOVAs on each dependent measure that offered insights on the results for Q3 and Q2; we report these in this order for clarity.

Q3 (What is the relationship between gap and completion time with Stitching?): Our main findings show that as the physical distance between displays increases, there is a corresponding increase in completion time, even though the physical movement of the mouse remains constant. Fitts's Law [20, 15] predicts movement time based on a user's movements in motor space, not on the distance traveled by the cursor. Although we can assume that switching attention between displays during targeting has a time cost, is this reflected in a slower ballistic movement or a lengthening of the homing-in phase? Previous studies have shown that the timing and magnitude of the peak velocity in a targeting task are a function of movement amplitude regardless of precision constraints, and that the proportion of time spent decelerating to the target is a function of an accuracy constraint, (usually seen through small target widths), regardless of the distance covered [15].

Our analyses on the data from Stitching revealed that participants reached higher top speeds ($F_{2,22}$=5.02, $p<0.05$) with increasing gap and spent longer in the ballistic phase ($F_{2,22}$=33.1, $p<0.001$). These results suggest that movement times should be shorter with increasing gap, not longer as our results indicate. In addition, with larger gaps, participants were not spending a greater proportion of time in the homing-in phase ($F_{2,22}$=1.41, $p>0.05$). Thus, longer movement times are not a result of a slower ballistic phase or a longer deceleration phase; participants were actually moving faster in conditions with larger movement times.

An increase in top speed and time to reach top speed are indicative of longer movement amplitudes, revealing that participants planned and executed the ballistic phase of the movement based on the physical configuration of the displays, even though they were aware that the cursor would immediately warp to the second monitor. In fact, users reached their top speed on the destination display 39% of the time for a small gap, but 58% and 57% of the time for the medium and large gaps respectively. As a result, we can expect that users will consistently overshoot the target when aiming over a gap, which explains the longer movement times. An examination of the percentage of trials with a horizontal overshoot of the target showed that overshoot errors increased dramatically with a physical gap. For the small gap, 36% of trials contained an overshoot as compared to 68% for the medium gap, and 80% for the large gap. These additional overshoots came as a result of users incorporating the physical space between monitors into their motor planning, which subsequently increased movement times when warping the cursor over the gap.

Q2 (Does Halo help?): Our main findings show that Halo reduced movement time when aiming across displayless space. In our study, users were only without visual feedback within the gap, creating a closed-loop to open-loop to closed-loop transition within a single aiming movement. The two-phase model of movement suggests that if users are in the ballistic phase of their movement within displayless space, the halo will not be helpful, and might slow a user down by forcing them into closed-loop aiming. If users are decelerating within the gap, the halo should help because it provides essential visual feedback for movement correction.

Examining only the Mouse Ether and Ether+Halo data, we used peak velocity to determine on which display the ballistic phase ended and the feedback-corrected homing-in phase began. For trials where the ballistic phase ended prior to the gap, or within the gap, there was a movement time advantage with Halo for the medium and large gaps, but not for the small gap. For trials where the ballistic phase ended on the destination display there was no advantage of Halo for the small or medium gap. For the large gap, there was a difference; however there were only 48 trials (less than 2% of analysed trials) that comprised this sample for comparison. In addition, there were no effects of Halo on peak velocity ($F_{1,11}$=2.14, $p>0.05$), time to reach peak velocity ($F_{1,11}$=2.47, $p>0.05$), or percent time spent in the deceleration phase ($F_{1,11}$=1.01, $p>0.05$). Halo did result in decreased time after peak velocity ($F_{1,11}$=7.71, $p<0.05$), likely due to a decreased number of overshoots. Examining the horizontal overshoots showed that there was a smaller percentage of trials with an overshoot when Halo was used for the medium gap (Ether: 36% Halo: 30%), and the large gap (E: 53% H: 35%), but not the small gap (E: 25% H: 23%). These results suggest that the halo is indeed used by participants to know where the cursor is in displayless space, assisting with targeting by decreasing overshoots, thus decreasing movement time.

DISCUSSION

Our discussion focuses on four major issues: the advantage of stitching, the value of off-screen feedback, the costs of warping, and the effects of boundaries.

Advantage of Stitching

Our results indicate that Stitching is better than any of the Mouse Ether conditions at medium and large gaps. This suggests that the costs of absent feedback (or non-direct feedback such as Halo) outweigh the costs of having to reacquire the cursor after a jump. At the short gap distance, techniques were roughly equivalent; this result is hardly surprising, since at that distance (3.8cm), the gap represents a small part of the trajectory and the behavior of the cursor is very similar for the three techniques.

The previous findings may seem to contradict the results of the initial evaluation of Mouse Ether [5]. Although we did not try to replicate Baudisch and colleagues' experiment, we expected that the data would follow the same pattern (superiority of Mouse Ether). In fact, question four (regarding corners) was asked to determine the kinds of tasks that gave Mouse Ether an advantage. The data shows, however, that if there is an advantage to be found for Mouse Ether, it is not in the ability to cut corners, but in some of the factors that we did not test in our experiment but were part of the initial study.

We believe that the discrepancy between the two experiments can be explained through the differences in resolution of displays in the initial study, which would have affected Stitching, but not Mouse Ether. Resolution differences cause serious horizontal misalignment and asymmetry in CD gain (see Figure 2). In our study we did not considered resolution difference as a factor because it is not an intrinsic characteristic of Stitching; it can be corrected, which may reduce the advantage of Mouse Ether.

Although our results do not show performance advantages for Mouse Ether, it is clearly still a valuable technique and may have advantages over Stitching in other situations. For example, Ether is likely to be better when a consistent motor space is needed to determine one display destination among many, and in three-dimensional mappings such as Perspective Cursor [25].

Off-screen feedback

Our study is the first exploration of off-screen feedback for targeting. Results conclusively show Halo's usefulness: we observed that Halo improved cross-display targeting for the medium and large gaps. Moreover, the exploratory analyses confirm that users use Halo to prevent overshoot.

This advantage comes, however, at the cost of possible distraction and increased clutter, likely reasons why several participants felt strongly against it.

The cost of warping

When using Stitching, targeting time increased in proportion to the length of the gap (approximately 7ms per degree of gap). We discuss above how this increase is due to the natural impulse of users to match their motor responses to the visual space instead of to their motor experience. This result is important because it highlights that displayless space is difficult for users to ignore, and therefore should be taken into account in the design of fractured interfaces.

The effect of warp distance in performance might also be relevant for single-display warping techniques. For example, the mismatch between motor and visual space might affect Object Pointing (analogous to Stitching) when the angle between two objects is large.

Boundaries

Display boundaries that keep the cursor in the display (as in most single-display systems) facilitate some targeting tasks because elements on the boundary become infinitely deep in motor space [1]. Current interfaces take advantage of this property and place frequently-used elements on the boundaries (e.g., the task bar in Windows XP and Vista™ or the menu bar in Mac OS X™).

Mouse Ether eliminates some of these hard boundaries in exchange for the ability to cut corners; our analysis could not, however, find any significant advantage of cutting corners. We believe that boundaries should not be sacrificed for extra 'ether' around screens if they can be used to facilitate targeting of frequently used objects.

LESSONS FOR PRACTITIONERS

From our findings we have identified the three main lessons for designers of multi-display environments:

Stitching is still the safest choice. Mouse Ether – with or without Halo – is equivalent or inferior in performance to Stitching, at least in situations with few displays.

With Mouse Ether, use Halo. If Mouse Ether is required, off-screen feedback can help. Halos are disliked by some users, however, and are potentially distracting.

Gaps cost time. Larger gaps cause longer targeting times, even with Stitching. To avoid this cost, displays should be places as close to one another as possible.

OPEN RESEARCH QUESTIONS

The results of our study point at several unexplored issues that we plan to address in future work.

Improving Stitching. Stitching was the clear winner in our performance tests, but its performance was still negatively affected by the size of the gap. It may be possible to design visualization techniques that reduce this effect by showing the user the likely trajectory of the cursor, thus making the jump more predictable.

Improving Mouse Ether's Halo. Our results indicate that off-screen feedback helps improve performance of the Mouse Ether technique. We believe that performance can be further improved with off-screen feedback aids that are specifically designed for fast-moving objects (Halo was designed mostly for relatively static off-screen data).

Hybrid cross-display techniques. Existing implementations of Mouse Ether map motor space exactly to the physical gap between monitors. Other mappings are possible that could achieve the best of both Stitching and Mouse Ether. For example, a reduced CD gain for displayless space might retain the natural feeling of Mouse Ether while performing as fast as Stitching.

A targeting model that accounts for displayless space. Baudisch and colleagues suggested the development of a new model (or extension of a current model) that predicts targeting times across displays. Although our study moves us in this direction, further empirical evaluations have to be designed in order to obtain the equivalent of Fitts's law for multi-display environments.

Effects of resolution. We could only explain the discrepancies in results between our study and the original Mouse Ether study by the use of different resolutions in the two displays. Although resolution differences can be easily overcome, a validation of the effect of unadjusted resolution may help explain other phenomena or pinpoint other important factors that affect multi-display targeting.

Multi-modal feedback. Our study exclusively investigated visual feedback. Other kinds of feedback as auditory or haptic are promising because they can provide information about cursor position when visual feedback is not available.

CONCLUSION

As multi-display environments become more common, the way in which cross-display movement is supported becomes an important issue for designers to consider. Cross-display cursor movement is an important part of most cross-display tasks; therefore, performance can be dramatically improved if we choose optimal movement techniques.

In this research we analyzed two cross-display techniques, Stitching and Mouse Ether, that differ dramatically in how they deal with displayless space. We also tested a version of Mouse Ether that provides off-screen feedback through the Halo technique. We found that Stitching is faster than or equivalent to Mouse Ether for targeting tasks at three inter-display gaps. We also found that Halo helps performance in Mouse Ether, although not enough to be faster than Stitching. The results also reveal that increases in gap distance between displays increase targeting times for all techniques. This was unexpected for Stitching, since the required targeting movement is independent of the distance between displays.

These results translate into direct advice for the design of new multi-display interfaces that rely on indirect input. In addition, our findings uncover important questions for future research, and hint at the development of new and more efficient cross-display techniques.

REFERENCES

1. Accot, J. & Zhai, S. More than dotting the i's - foundations for crossing-based interfaces. *Proc. CHI'02*, 73-80.
2. Asano, T., Sharlin, E., Kitamura, Y., Takashima, K., & Kishino, F. Predictive Interaction Using the Delphian desktop. *Proc UIST '05*, 133-141.
3. Ashdown, M., Oka, K., & Sato, Y. Combining head tracking and mouse input for a GUI on multiple monitors. *Ext. Abs. CHI'05*, 1188-1191.
4. Balakrishnan, R. "Beating" Fitts' Law: virtual enhancements for pointing facilitation. *IJHCS*, 61, 6, (2004), 857-874.
5. Baudisch, P., Cutrell, E., Hinckley, K. & Gruen, R., Mouse Ether: Accelerating the Acquisition of Targets Across Multi-Monitor Displays. *Proc. CHI'04*, 1379-1382.
6. Baudisch, P., Cutrell, E., Robbins, D., Czerwinski, M., Tandler, P. Bederson, B., & Zierlinger, A. Drag-and-Pop and Drag-and-Pick: Techniques for Accessing Remote Screen Content on Touch- and Pen-operated Systems. *Proc. Interact 2003*, 57-64.
7. Baudisch, P. & Rosenholtz, R. Halo: A Technique for Visualizing Off-Screen Locations. *Proc. CHI'03*, 481-488.
8. Benko, H. & Feiner, S. Multi-Monitor Mouse. *Proc. CHI'05*, 1208-1211.
9. Booth, K., Fisher, B., Lin, R. & Argue, R. The "mighty mouse" multi-screen collaboration tool. *Proc. UIST'02*, 209-212.
10. Casiez, G., Vogel, D., Balakrishnan, R. & Cockburn, A. The Impact of Control-Display Gain on User Performance in Pointing Tasks. *JHCI* (2007) (In Press).
11. Crossman, E.. & Goodeve, P. Feedback control of hand movements and Fitts' law. *QJEP*, 35, (1983), 251-278.
12. Dickie, C., Hart, J., Vertegaal, R. & Eiser, A. LookPoint: an evaluation of eye input for hands-free switching of input devices between multiple computers. *Proc. OZCHI'06,* 119-126.
13. Fitts, P.M. & Peterson, J.R. Information capacity of discrete motor responses. *J. of Exp. Psych*., 67, (1964), 103-112.
14. Guiard, Y., Blanch, R. & Beaudouin-Lafon, M. Object pointing: a complement to bitmap pointing in GUIs. *Proc. of GI'04*, 9-16.
15. Graham, E. & MacKenzie, C.L. Pointing on a computer display. *Proc. CHI'95*, 314-315.
16. Ha, V., Inkpen, K., Wallace, J. & Ziola, R. Swordfish: user tailored workspaces in multi-display environments. *Ext. Abs. CHI'06,* 1487-1492.
17. Hascoët, M. Throwing models for large displays. *Proc. British HCI'2003*, Vol. 2, 73-77.
18. Johanson, B., Hutchins, G., Winograd, T. & Stone, M. PointRight: Experience with Flexible Input Redirection in Interactive Workspaces. *Proc. UIST'02,* 227-234.
19. Khan, A., Fitzmaurice, G., Almeida, D., Burtnyk, N. & Kurtenbach, G. A Remote Control Interface for Large Displays. *Proc. UIST'03*, 127-136.
20. MacKenzie, I.S. Fitts' law as a research and design tool in human-computer interaction. *JHCI*, 7,(1992) 91-139.
21. MacKenzie, C.L., Marteniuk, R.G., Dugas, C., & Eickmeier, B. Three-dimensional movement trajectories in Fitts' task: Implications for motor control. *QJEP*, 39A, (1987), 629-647.
22. Mackinlay, J., Good, L., Zellweger, P., Stefik, M. & Baudisch, P. City lights: contextual views in minimal space. *Ext. Abs. CHI '03*, 838–839.
23. 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.
24. Nacenta, M., Sakurai, S., Yamaguchi, T., Miki, Y., Kitamura, Y., Subramanian, S. & Gutwin, C. E-conic: a Perspective-corrected Interface for Multi-display Environments. *Proc. UIST'07,* (in press).
25. Nacenta, M.A., Sallam, S., Champoux, B., Subramanian, S. & Gutwin, C. Perspective Cursor: Perspective-Based Interaction for Multi-Display Environments. *Proc. CHI'06*, 289-298.
26. Schmidt, R., Zelaznik, H., Hawkins, B., Frank, J., & Quinn, J. Motor output variability: A theory for the accuracy of rapid motor acts. *Psych. Review*, 86, (1979), 415-451.
27. Tan, D.S., Meyers, B., Czerwinski, M. WinCuts: Manipulating arbitrary window regions for more effective use of screen space. *Ext. Abs. CHI'04*, 1525-1528
28. Swaminathan, K. & Sato, S. Interaction design for large displays. *Interactions*, 4, 1, 15-24.
29. Woodworth, R.S. The accuracy of voluntary movements, *Psych. Review Monograph*, Supp., 3, (1899), 1-114.