

Comparing Content and Input Redirection in MDEs

James R. Wallace
University of Waterloo
Waterloo, ON, Canada
jrwallac@uwaterloo.ca

Regan L. Mandryk
University of Saskatchewan
Saskatoon, SK, Canada
regan@cs.usask.ca

Kori M. Inkpen
Microsoft Research
Redmond, WA, USA
kori@microsoft.com

ABSTRACT

Designers of Multi-Display Environments (MDEs) often use input redirection to allow users to manipulate content on multiple displays with a single interaction device, but users seated at sub-optimal positions (i.e., not facing the display) may find interaction difficult or frustrating. In collaborative MDEs, users should be able to choose their preferred collaborative arrangement, rather than adjusting to the limitations of the technology. We compare content and input redirection from a variety of seating positions in an MDE. Results from our studies show that content redirection does not suffer from performance loss in sub-optimal seating positions, as opposed to input redirection, which does. Content redirection provides a method for all members of a group to interact with shared content regardless of their position relative to a shared display.

Author Keywords

multi-display environments, interaction, redirection, input

ACM Classification Keywords

H5.2 [User Interfaces] Interaction Styles

INTRODUCTION

The vision of collaborative MDEs involves users sharing information and applications in a group setting, interacting with each other using fixed large shared displays, and incorporating their personal devices into the environment. Given that small-group face-to-face interaction is a rich form of co-located collaboration, [20] collaborative MDEs should be designed to support multiple users. While research of MDE interaction has been generally conducted in laboratories, in real-world MDEs, configurations of people and technology will vary; users may choose to sit in a variety of arrangements, displays may be placed anywhere in the environment and people have the freedom to move about the room. Room arrangements which are less than ideal may significantly impact interaction in MDEs.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW '08, November 8-12, 2008, San Diego, California, USA.

Copyright 2008 ACM 978-1-60558-007-4/08/11...\$5.00.

The most common approach for supporting interaction in MDEs involves *input redirection*— moving a user's control focus from one display to another. While this technique has been used in collaborative MDEs [14, 18], it may cause problems when users interact with occluded displays, distant displays, or cases where users simply prefer to sit facing each other. Previous work by [28] demonstrated that the configuration of users and displays can have a significant impact on the effectiveness of input redirection in MDEs. In particular, when users had an unnatural mapping between the control space and the display space (i.e. when the display was behind the user or to their right) performance on a docking task degraded significantly.

An alternative to input redirection is *content redirection*, where content from one device is mirrored onto another. This approach is frequently used to project content from a personal device (e.g. laptop) onto a large shared display. In contrast to previous MDE research we wanted to explore redirecting content from a shared public display onto a user's personal device in a manner that would allow the user to view and interact with the content easily. We expected that content redirection would be less susceptible than input redirection to performance degradations as a result of sub-optimal seating arrangements in MDEs.

This paper compares *input* and *content redirection* for interaction in MDEs under a variety of seating arrangements. The Swordfish framework [26] was used to implement and evaluate four different interfaces. Two of the interfaces used input redirection and have been established in the literature as appropriate for use in MDEs. The second two interfaces employed content redirection and were designed to support cross-display interactions in MDEs. Additionally, each interface used either the keyboard or mouse to transition between displays. Results from this work show that content redirection is a practical approach to interaction in MDEs regardless of seating arrangement. Content redirection was superior to input redirection in terms of performance and preference when users were seated in sub-optimal positions.

RELATED WORK

Interfaces for MDEs

Prior research has explored techniques for multi-user and MDE interaction including functions such as redirecting mouse and keyboard input [14], application relocation [5], shared clipboards [17], and content redirection [25]. As we

develop systems that combine many of these approaches [7], it will be important to know how interaction techniques impact users' abilities to interact in realistic MDEs.

Input Redirection

Input redirection involves moving a user's control focus from one device to another allowing a user to interact with multiple displays. To date, most MDE implementations have focused on input redirection as the primary interaction mechanism. PointRight [14] utilized an Extended Desktop approach where users could move input between displays by dragging the mouse cursor off the edge of one display onto another. Often, the virtual layout of displays is modeled after their spatial layout in the physical world.

Several projects have enhanced basic Extended Desktop functionality. These include modifying the mouse trajectory to provide more intuitive movement between displays [2], simulating stickiness to aid users in targeting objects on the boundary between displays [16], and incorporating head tracking to match the user's perspective [19]. The Extended Desktop approach is often used to support static display configurations [14]; however, there have been projects that emphasize support for dynamic configuration of displays [11, 12, 26]. These projects take into account mobile devices such as handhelds, and tablet PCs, which are unlikely to remain in a single position or orientation. Dynamic display configurations also provide users with a method of customizing their environment to better fit their own needs or preferences.

A second application of input redirection in MDEs is the use of hotkeys to "jump" the mouse cursor between displays [3]. Under this paradigm each display is treated as a separate workspace and users can change input focus to any of these workspaces as needed. Benko and Feiner [3] compared mouse, keyboard, head orientation and mouse placement as methods to transition between displays in MDEs. They found that their hotkey-based technique allowed users to quickly move between displays and complete target selection tasks significantly faster than with a simple Extended Desktop implementation. Other work by Ashdown et al. [1] implemented a head tracking-based implementation of this technique with similar findings, however task time increased in addition to user preference.

A third input redirection approach is a system where icons used to represent displays in the MDE are contained in a palette or world-in-miniature view [5, 7, 8]. This approach is attractive to designers because it provides an overview of the MDE and can provide contextual navigation aids; for example, including physical artifacts in the world-in-miniature can help identify displays within the MDE [5].

Content Redirection

Content redirection allows the contents of one display to be viewed on another display. Virtual Network Computing (VNC), was established as a protocol for interacting with remote displays over low-bandwidth network connections.

Despite being designed for single-computer network administration, VNC provides an established and well-developed technology for MDE interaction. Previous work has shown collaborative benefits from content redirection for MDEs by providing users with the ability to share portions of their private displays on a shared group display [7, 25]. One benefit of this approach is that it enables a user to control what is shared with the group, avoiding potential privacy issues and reducing unrelated information being shared. Other work extends this notion by providing control over what content is shared, and allowing users to conceal lower-level details if desired [4, 7].

Another benefit of content redirection is that it provides an opportunity for users to interact with a shared resource on their preferred display - without disadvantaging other users. With content redirection a user can create a copy of a shared display on their local display and interact with the content without changing its visibility for others. Content redirection provides a malleable display space, and while it has been explored in a number of projects previously, the primary goal has been to share content rather than as a means of interacting with it. In contrast, our work examines the appropriateness of redirecting shared content to a personal device for interaction in collaborative MDEs.

User and Display Arrangements in MDEs

Factors such as the arrangement and number of people and displays in an environment can impact the effectiveness of collaboration. Previous literature has examined these factors and provides insight into how MDE interfaces can better support co-located collaboration.

User Arrangement

MDEs need to support the most natural ways of collaborating instead of forcing users to adapt their collaboration to fit the technologies; this may mean a theatre-style configuration or a face-to-face arrangement. With the exception of tabletop research, much work to date has concentrated on scenarios in which users sit side-by-side and focus on a shared display despite evidence that users may prefer alternate seating arrangements.

Sommer [20] found that participants preferred to sit across from each other when casually conversing or competing, but chose to sit side-by-side for collaboration. The main reason cited for sitting side-by-side was to make it easier to share physical artifacts. While this choice makes sense for physical artifacts, it may be less relevant for computer-supported collaboration given the ease with which digital artifacts can be shared. MDEs that are flexible in sharing group resources can provide advantages over traditional environments. Other work has examined the issue of proximity between users and found that people preferred to work in close proximity with their collaborators [10].

Arrangement of Displays

Researchers have also investigated how factors such as the size and layout of displays can impact co-located

collaboration [13], and where displays and users should be positioned to best support work in MDEs. Su and Bailey [22] found that display position can directly affect performance and subjective workload and they provide guidelines for the positioning of displays within an MDE.

While it is important to understand the impact of user and display arrangements in MDEs, the interaction between these two variables can impact the effectiveness of the environment. Wigdor et al. [28] explored the consequences of seating position when interacting with a large wall display and found that orientation of the control space in relation to the display directly affected performance in the MDE. These results demonstrate the importance of examining MDE interaction techniques with various display configurations to understand any impact that display or user arrangements may have on the technique.

Evaluating MDE Interfaces: Single- vs. Multi- User

The focus in this work was to evaluate MDE interfaces using baseline metrics of performance such as task completion time, accuracy, workload, and preference. These measures can be assessed with a single-user study, without introducing the complexity of a collaborative environment. For example, in [28], the authors focused on individuals' performance and then formed multi-user guidelines based on these results. In our case, multiple users could introduce confounds, as the collaborative process may impact when, where and how interaction took place.

In cases where researchers are interested in understanding group interactions, studies on individual users are likely not appropriate. Biehl and Bailey [6] recognized this limitation and specifically examined multi-user collaboration in MDEs. By collecting data in a group setting, the authors were able to develop guidelines aimed at improving application relocation and sharing in MDEs that would have been difficult to glean through single-user studies. Our interest is in fundamental interactions within collaborative MDEs, and therefore, single-user studies were more appropriate for addressing our research questions.

INTERFACE IMPLEMENTATION DETAILS

We developed four interfaces that redirect input or content using one of two transition mechanisms: keyboard or mouse. With keyboard transitions, redirection is activated by pressing a keyboard button, whereas with mouse transitions this functionality is triggered by moving the mouse cursor (see Figure 1). Each interface was developed using the Swordfish framework [26], a multi-display groupware framework, developed in C#. By using the same framework, we eliminated performance differences between the interfaces. Detailed descriptions of each interface follow, while Figure 1 shows each of the interfaces.

Extended Desktop (ExD)- Input/Mouse

The Extended Desktop (also called Virtual Paths) interface is a standard technique used by many popular operating systems to move the cursor between multiple displays

connected to a single computer. MDEs have applied this metaphor to allow a single mouse cursor to move across displays connected to several computers [21]. In this approach, users move their mouse cursor "off" one edge of a display and it seamlessly appears on the corresponding edge of another display. The ExD interface represents a mouse-triggered approach for input redirection because mouse movement is used to transition between displays.

Multi-Monitor Mouse (MMM)- Input/Key

The Multi-Monitor Mouse interface represents a key-triggered transition for input redirection, recently explored by Benko and Feiner [3]. MMM is a key-switching interface where the mouse jumps between two displays when a key is pressed. Although key combinations such as "ALT-TAB" are used in popular operating systems and are familiar to users, for simplicity, we used a single key ("~"), marked with a green sticker to trigger the switch. When the mouse cursor jumps to an alternate display, it can be positioned in a variety of locations. Two likely placements, explored in [3], are centering the mouse cursor or keeping it in the same relative position that it occupied on the previous display; we used the latter approach.

Edit Blind (EB)- Content/Mouse

The Edit Blind interface represents a mouse-triggered implementation for content redirection. In order to bring remote content onto the laptop's display, users move their mouse cursor up to the top of their display, then back down; like "pulling down a blind". To return to the local display, users moved their mouse to the top of the laptop display as if they were "pushing the blind back up". Although users could partially extend a blind—viewing portions of two displays at once—our study only allowed users to toggle between displays. This design choice was made to reduce functionality in favour of lightweight interaction, and for comparison to the other 'all-or-nothing' techniques.

Key Swap (KS)- Content/Key

The Key Swap interface represents a key-based transition for content redirection. In our studies, we used the "~" key to trigger the content redirection, which caused the entire contents of the public shared display to be mirrored onto a personal laptop (to support local interaction). Again, the trigger key ("~") was marked with a green sticker.

STUDY: CONTENT AND INPUT REDIRECTION

We wanted to compare content and input redirection in varying seating configurations. In real MDEs, not everyone will have an optimal view of the large-shared displays, and it is important to design interaction techniques that support collaboration rather than force users to adopt inferior collaborative arrangements. We used a simple task intended to mimic basic mouse-based interactions, such as target acquisition, with a WIMP-based computer. Our MDE consisted of a large (remote) wall display and a laptop (local) display. Four seating positions relative to the large display were used and named according to a North, South,

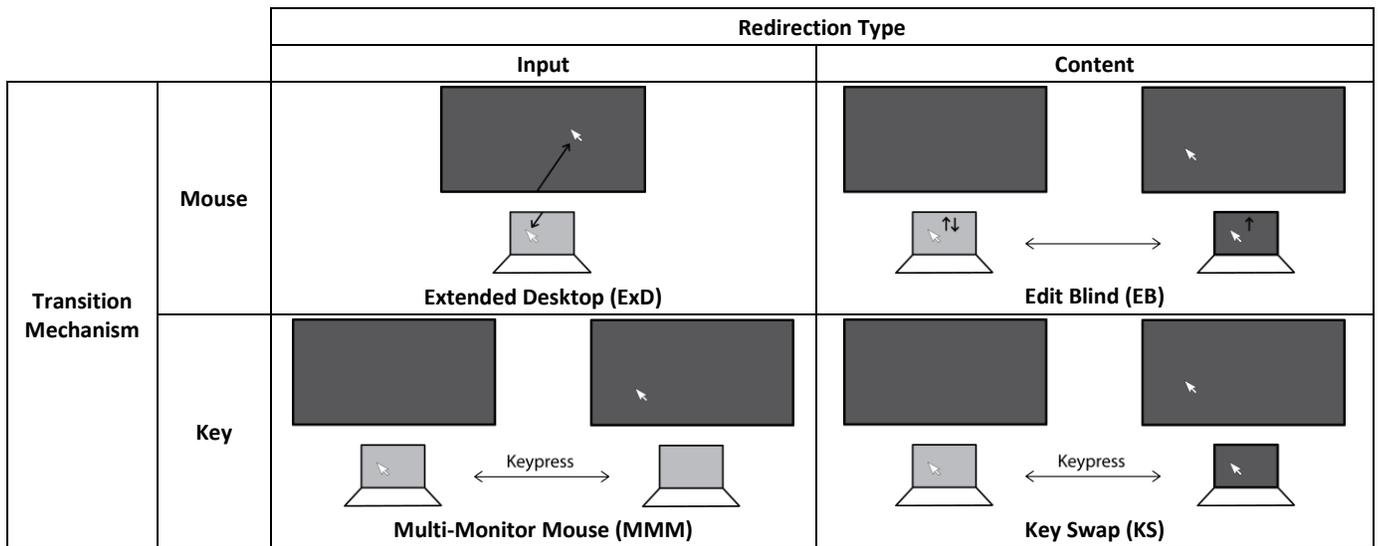
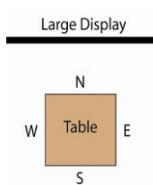


Figure 1: Categorization of the interfaces investigated. Interfaces are based on two main criteria: redirection type (input/content) and transition mechanism (mouse/key).



East and West scheme. The experiment utilized a 2 x 2 x 4 within-subjects design with redirection type (input, content), transition mechanism (mouse, key), and seating position (N, E, S, W) as independent variables, and performance, subjective workload, and preference as dependent measures.

Task

We used a docking task to simulate desktop interaction in a multi-display setting. The task was based on one used by Wigdor et al. [28] since it simulates fundamental desktop interactions such as target acquisition and selection.

Each trial consisted of two stages: a *dock* stage and a *dialog* stage. In the docking stage, participants used a mouse to repeatedly acquire and drag a small blue square and drop it onto a large red square, which changed colour when a user could successfully dock the blue square. There were four successive docks per trial; after each dock, the red square relocated to a new position. At the beginning of each trial the blue square started in the centre of the screen, however it maintained its position when the red square relocated.

Docking locations were pre-calculated based on a randomly assigned angle and three pre-assigned distances; 125, 250 (used twice), and 500 pixels relative to the previous dock location. Each trial used a randomly assigned combination of docking distance and angle to prevent the user from memorizing the order of dock locations, yet each set of four docking movements was comparable in terms of the movement time predicted by Fitts's Law [15]. Target width remained constant throughout the study for similar reasons.

After the four docks were complete the *dialog* stage began. Following the last dock, a dialog appeared on the remote (large) display asking the user to shift to the local (laptop)

display. Then, a dialog box appeared on the laptop display in a random position. The user transitioned to the laptop display, clicked on the local dialog box, then reacquired the blue square on the large display to begin a new trial.

Procedure

Sixteen right-handed participants (9 male), aged 16 to 44, participated in the study. Prior to the study 9 of the participants had never used a multi-display system but had used systems with screens larger than 20" on a monthly or weekly basis. Additionally, 12 of the participants had never used remote administration software such as VNC.

Participants sat with the laptop and mouse on a table positioned approximately 6' from a large projected display. The laptop's display measured 15.5" diagonally, whereas the projected display measured 60" diagonally. Both displays were set to a resolution of 1024 by 768 pixels.

Participants completed four practice trials to become acquainted with the interface. Following this, each participant completed 10 trials (4 docks + 1 dialog per trial) for each of the 16 conditions (4 seating positions x 4 interfaces). Starting position and interface order were counterbalanced across participants. After completing trials in all four seating positions for an interface, the participant completed a brief post-condition questionnaire including a modified NASA-TLX subjective workload assessment. Participants then repeated this process for the remaining three interface conditions. Finally, after all conditions had been completed, the participant filled out a post-experiment questionnaire, which directly compared all interfaces and seating conditions in terms of comfort and usability.

Hypotheses

We formulated 6 hypotheses in terms of performance, subjective workload and preference:

H1: Participants will perform the docking task slower in the North seating position (North) than in the other seating positions (South, East, West).

H2: Participants will perform the docking task slower using interfaces which redirect the input to the remote display (ExD, MMM) rather than interfaces which redirect content to the local display (EB, KS).

H3: Participants will have the slowest docking times with interfaces which redirect the input to the remote display (ExD, MMM) in the North seating position. There will be a performance benefit for redirecting the content (EB, KS) to the local display in the North seating position.

H4: Participants will take longer to transition between displays using interfaces which redirect the input to the remote display (ExD, MMM) rather than interfaces which redirect the content to the local display (EB, KS).

H5: Participants will prefer the content redirection interfaces (EB, KS) over the input redirection interfaces (ExD, MMM) for seating positions other than South.

H6: Participants will rate the content redirection interfaces (EB, KS) as better (lower) on the subjective workload scales than the input redirection interfaces (ExD, MMM).

Data Analyses

Timing information was gathered from computer logs of mouse events. Dock time was calculated from acquisition of the blue square to a successful dock on the red square. The times for each of the four docks in one trial were summed into a single docking time (DT). Since path length and the logarithm (base 2) of path length were comparable for each trial, we treated the series of 4 docks as one trial.

We calculated the time to transition between the local and the remote display (LtoR) before each trial as the time from which the user clicked on the dialog box on the local display until the user acquired the blue square on the remote display. We also calculated the time to transition from interaction on the remote display to interaction on the local display (RtoL) after each trial as the time from which the user clicked on the dialog box on the remote display to the time that the user clicked on the dialog box on the local display. All times were calculated to the nearest ms.

To reduce the effects of transitioning between conditions on our data, we separated the data into two blocks of five trials, and only analyzed data from the second block of five trials for each seating position with each interface.

To investigate the six hypotheses we split the interface factor into its two defining components: whether content or input was redirected and whether the mouse or keyboard was used to perform the redirection. We performed a Repeated Measures Multivariate Analysis of Variance on the timing-based dependent measures with seating position (North, East, South, West), redirection type (Content, Input), and transition mechanism (Mouse, Keyboard) as

factors. All main effects and interaction were tested at $\alpha=.05$, and Bonferroni adjustments were used for all post-hoc analyses. In cases where the sphericity assumption was violated, the degrees of freedom were adjusted using the Huynh-Feldt method. Questionnaire data were analyzed using non-parametric statistical techniques.

Results

We first present the results for docking time (DT), followed by transition time (LtoR, RtoL), and subjective workload and preference measures from the questionnaires.

Docking Time (DT)

Neither the redirection type ($F_{1, 15}=2.2$, $p=.159$, $\eta^2=.13$) nor the transition mechanism ($F_{1, 15}=0.5$, $p=.513$, $\eta^2=.03$) had a significant effect on DT. However, there was a significant effect of seating position on DT ($F_{3, 45}=3.8$, $p=.016$, $\eta^2=.20$). Although participants were slowest to perform the docking task in the North position, followed by the West, the East and the South positions, post hoc analysis revealed that participants were significantly slower ($p=.034$) in the North seating position than the South seating position.

The main effect of seating position needs to be interpreted in the context of a significant interaction effect between redirection type and seating position on DT ($F_{3, 45}=4.66$, $p=.006$, $\eta^2=.24$). Post-hoc analysis revealed that when content was redirected onto the local laptop display there were no performance differences in DT for the different seating positions. However when input was redirected to the remote large display, DT was longer when participants sat in the North position than in the South or East position ($p=.006$, $p=.021$ respectively). Figure 2 shows the performance advantage when redirecting the content to the local display at the North and West positions but not at the South and East positions. There were no interactions between seating position and transition mechanism ($F_{3, 45}=0.94$, $p=.432$, $\eta^2=.06$) or transition mechanism and redirection type ($F_{1, 15}=0.44$, $p=.517$, $\eta^2=.03$) on DT.

Transition Time

Transition time is comprised of two metrics: transitioning from the local display to the remote display to perform the docking task (LtoR), and transitioning back to the local display after finishing the docking task (RtoL). We would expect to see similar results in these two measures.

There was a significant effect of redirection type on both LtoR ($F_{1, 15}=56.98$, $p\approx.000$, $\eta^2=.79$) and RtoL ($F_{1, 15}=36.12$, $p\approx.000$, $\eta^2=.71$). Post-hoc analysis revealed that in both cases, participants took significantly longer to transition between displays when input was redirected than when content was redirected. This is likely due in part to the increased overhead of shifting attention between displays.

There was a significant effect of transition mechanism on RtoL, ($F_{1, 15}=6.96$, $p=.019$, $\eta^2=.32$) but not LtoR ($F_{1, 15}=1.44$, $p=.248$, $\eta^2=.09$). In the case of RtoL participants took significantly longer ($p=.019$) to transition when using

the mouse than when using the keyboard, which is an expected result. The same trend was true for LtoR; however, the difference was not statistically significant. Although there was not a significant main effect of transition mechanism on LtoR, there was a significant interaction between redirection type and transition mechanism ($F_{1, 15}=9.90, p=.007, \eta^2=.40$). Figure 3 shows that for content redirection, using the keyboard was significantly faster than using the mouse ($p=.002$), but for input redirection there was no difference ($p=.201$). This is reasonable considering that for content redirection, a mouse gesture to the "hotspot" takes longer than a key press to execute. However, for input redirection the efficiency of executing a key press was offset by the time needed to shift attention between displays. Additionally, both input redirection techniques had longer transition times than either content redirection approach.

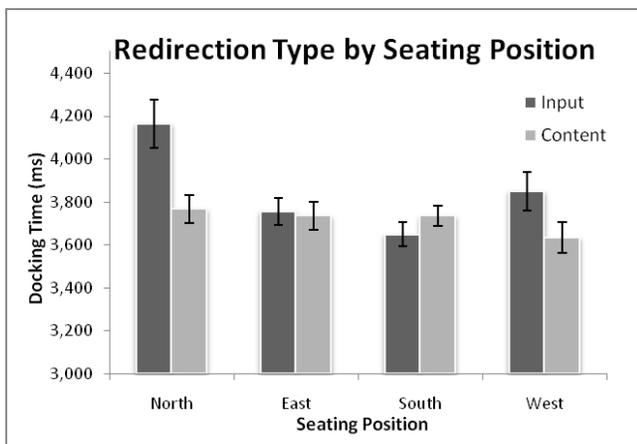


Figure 2: Means and SE for docking time by seating position and redirection type. Content redirection aids at the North and West positions, but not the South and East positions.

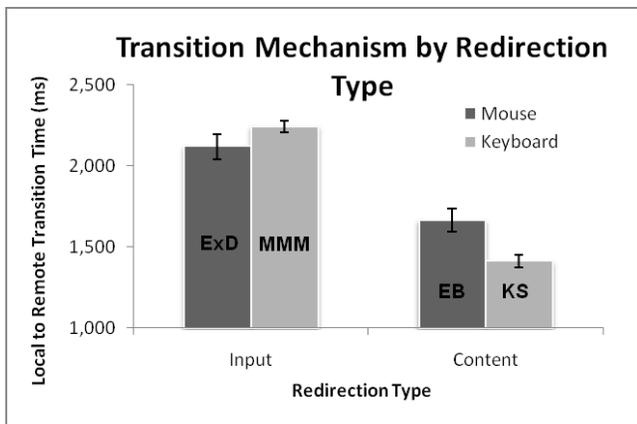


Figure 3: Means and SE for local to remote transition time organized by redirection type and transition mechanism. Using the keyboard is faster than the mouse for content redirection; there is no difference for input redirection.

There were also significant effects of seating position on both LtoR ($F_{3, 45}=5.30, p=.003, \eta^2=.26$) and RtoL ($F_{1.99,$

$29.86=3.70, p=.037, \eta^2=.20$). In both cases participants took longer to transition when seated in the North position than the South position ($p<.03$). However, these effects should be interpreted in the context of the significant interactions between seating position and redirection type on both LtoR ($F_{2.20, 33.06}=9.69, p\approx.000, \eta^2=.39$) and RtoL ($F_{1.83, 27.44}=8.33, p=.002, \eta^2=.36$). For both measures seating position had no impact on time when content was redirected (all $p\approx 1.000$). When input was redirected to the remote display it took participants significantly longer to transition in the North ($p=.008$), East ($p=.028$), or West ($p=.012$) positions than the South position. Similarly, when input was redirected back to the local display, it took participants longer in the North ($p\approx.000$) or East ($p=.008$) positions than the South position. As Figure 4 shows, for input redirection there is a performance loss when seated in suboptimal positions. Redirecting content improves performance overall and negates the disadvantage of seating position.

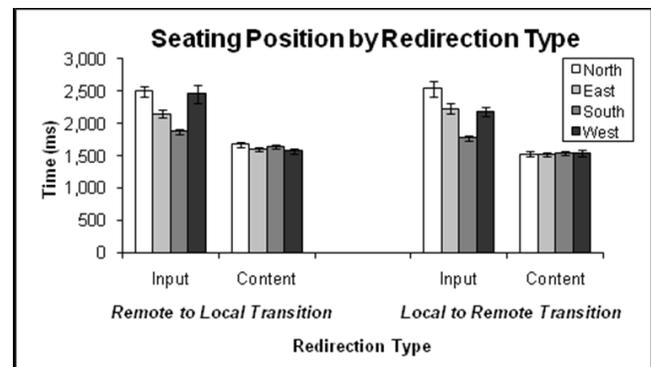


Figure 4: Means and SE for local to remote and remote to local transition times by seating position and redirection type.

There were no interaction effects of seating position and transition mechanism on LtoR or RtoL transitions ($F_{3, 45}=1.86, p=.149, \eta^2=.11$; $F_{1.64, 24.6}=2.39, p=.121, \eta^2=.14$ respectively), although there was a three-way interaction between transition mechanism, redirection type, and seating position on RtoL ($F_{2.19, 32.8}=4.42, p=.017, \eta^2=.227$). Participants transitioned faster using content redirection than by using input redirection at all seating positions whether the mouse or keyboard was used to transition (all $p<.010$) *except* in the case where the participant was seated in the South position, was using input redirection and was using the mouse to transition ($p=.778$). Therefore, the only situation where content redirection did not provide superior performance to input redirection was when participants used the Extended Desktop Interface (ExD) in the South position. This is not surprising given that this is the standard configuration that users are accustomed to. There was no disadvantage to providing content redirection in this case, but for all other cases participants were faster in RtoL transitions when using content redirection.

Subjective Results

Figure 5 shows the mean ratings from the subjective workload measures for each of the four interfaces. No

significant differences were found across interfaces for the amount of physical or temporal demand required from the interfaces. However, significant differences were found for perceived mental load, perceived performance, perceived effort, and level of frustration. Wilcoxon signed ranks tests were used to perform pairwise post-hoc analyses for each of these variables, using a Bonferroni correction ($\alpha=.008$). In terms of the perceived mental demand for the task, the Key Swap interface was found to be significantly less demanding than the Extended Desktop interface and the Multi-Monitor Mouse interface ($p<.008$). In terms of perceived performance, participants felt that they performed significantly better using the Key Swap interface compared to all of the other interfaces ($p<.008$). Participants also expressed much less frustration using the Key Swap interface than either the Extended Desktop interface or the Multi-Monitor Mouse interface. No significant pairwise differences were found for the perceived effort ratings.

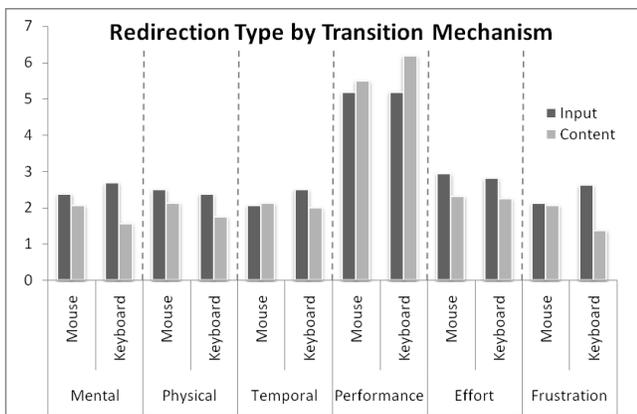


Figure 5: Average subjective workload measures for each transition mechanism. Each measure was rated on a 7 point scale (lower is better).

After completing each interface condition, users were asked to rate each seating position in terms of comfort. Figure 6 shows the mean ratings for each seating position, for each of the four interfaces. Participants reported significantly different comfort levels based on seating position for three of the four interfaces: Extended Desktop, Multi-Monitor Mouse and Key Swap. Post-hoc pairwise analyses revealed that participants found the North position significantly less comfortable than the East position for the Extended Desktop interface ($p=.001$). Additionally, participants found the North position significantly less comfortable than the South position for the Multi-Monitor Mouse interface ($p=.001$). In comparing ratings for each interface, we found that the comfort ratings only differed significantly for the North position with participants reporting that the two content redirection interfaces (EB, KS) were significantly more comfortable than the Multi-Monitor Mouse interface ($p=.001$, $p=.005$). At the end of the session participants were asked what their preferred seating position was; the majority of users (14) preferred the South position, while two users preferred the East position.

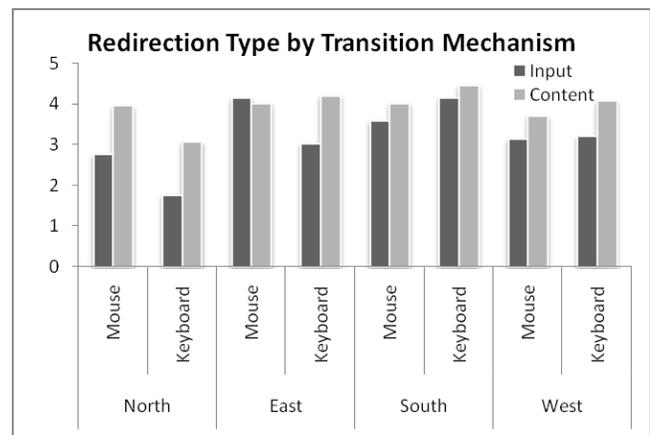


Figure 6: Mean comfort ratings for each transition mechanism by seating position. Participants rated each interface on a 5-pt scale (higher is better).

Participants were asked to rank the four interfaces in terms of how comfortable they were to use, the number of errors they felt they committed, and their overall preference. Overall, no significant differences were found for any of these variables and participants' comments mirrored these findings. Participants tended to like the additional display space found in input redirection-based interfaces; P6 commented that “[with] the large display on my laptop, I cannot open multiple windows to do multiple tasks at the same time”. Participants did however like the Extended Desktop approach; P16 commented that “It was smooth to go from one to the other, and felt like you had just one desktop instead of having two disjoined things you were trying to interact with.”

Participants also mentioned some advantages and disadvantages to the content redirection-based interfaces. Participant 5 recognized that content redirection was advantageous in that “[needing] to only concentrate on one screen the entire time saved time and sanity.” Participant P11 summarized many of the comments, saying “Oddly, I preferred working with the Extended Desktop, but the [Edit Blind] technique was probably easier to use. The Extended Desktop allows one to see a larger work area, but the [Edit Blind] method is easier for mousing”. Given that there was no one interface that ranked consistently as favorite, these comments provide some insight into the reasoning that participants used when rating and ranking the interfaces.

Summary of Results

Our results partially support our six experimental hypotheses, although the interaction among the independent variables was stronger than anticipated.

H1: *Participants will perform the docking task slower in the N than in the other seating positions.* This hypothesis was partially supported. Although users were slowest to perform the docking task in the North seating position overall, the difference was only significant for input

redirection. There were no differences in docking times for content redirection at the different positions.

H2: *Participants will perform the docking task slower with input redirection than content redirection.* Although the mean values of docking time support this hypothesis, we did not find a main effect as the results were impacted by seating position (H3). The hypothesis does hold for the North and West positions.

H3: *There will be a performance benefit for redirecting the content (EB,KS) to the local display in the North seating position.* Redirecting content to the local display rather than redirecting input to the remote display supported faster docking times on average. However, when content was redirected, there was no disadvantage to sitting at a suboptimal location (North or West). When input was redirected to the remote display, there was a marked difference in performance at suboptimal seating positions.

H4: *Participants will take longer to transition between displays using input redirection than content redirection.* We found that it was faster to transition both to and from the remote display with the content redirection-based interfaces than with input redirection.

H5: *Participants will prefer the content redirection interfaces for seating positions other than South.* This hypothesis was partially validated. Participants rated the KS and EB interfaces significantly higher than MMM in the North position, but did not rate them higher than ExD, nor was there a difference in the East and West positions.

H6: *Participants will rate the content redirection interfaces as better on the subjective workload scales.* Participants rated the Key Swap interface as significantly better on a number of the subjective workload measures, however this did not hold true for the EB interface; KS was rated significantly better than EB for some measures. These results are likely a reflection of the reduced transition time with content-redirection based interfaces. Additionally, the mouse-triggered transition in EB seems to have been a factor in its perceived workload; participants seemed to like the simplicity of pressing a key for transitions.

Our main findings include the following:

Content redirection benefits interaction in suboptimal seating positions. Our results show that redirecting content mitigated the negative effects of sitting at suboptimal seating positions for interactions on the remote display and transitions between displays. The corroboration of H3 shows that MDE environments that support content redirection will yield benefits for users interacting with large shared displays while seated at suboptimal positions, without hindering the performance of those seated in the optimal position. In addition, experimental corroboration of H4 shows that it is faster to transition between the displays when redirecting content than when redirecting input. In effect, redirecting content is not only a better approach in terms of *transitioning* between displays, but also supports

comparable *interaction* performance for all collaborators, regardless of where they are seated.

Mouse-based transitions were slower than key-based transitions. Additional results reveal that a mouse-based approach for redirection takes longer than a keyboard-based approach on average, and that this result is stronger for content-based redirection.

Content redirection does not disadvantage optimal interaction situations. Our three-way interaction shows that the only situation where content redirection was not superior to input redirection for transitioning from the remote to the local display is the Extended Desktop Interface (ExD) in the South position. This is the standard configuration to which users are accustomed. Our results also show that there was no disadvantage to providing content redirection in this case. As such, content redirection assists in all non-standard cases without hindering performance in the standard case.

Participants preferred content redirection at the least optimal seating position. In terms of preference, no interface was a clear winner. In the North position, participants preferred the content redirection interfaces for obvious reasons, but this preference did not carry over to other positions. We noticed that users interacted mainly with redirected content on their laptop when it was available, but eye-tracking was not used so it is impossible to quantify or confirm this effect.

Limitations of the Study

The results from this study raised two limiting factors which may have impacted the results. The first limitation is in determining whether the benefits from content redirection are due to the closer proximity of the local display. To investigate this possibility, we conducted a follow up study with 16 participants to investigate the role of display distance in 2D interactions with an MDE. Fitts's Law [15] predicts movement time based on a user's movements in control space, not on the distance traveled on the display. As expected, the results of the follow-up study (details in [27]) showed no systematic differences of proximity to the display, and suggest that the superior performance of participants using content redirection in our main study was not due to their proximity to the display. In addition, these results remove the potential confound that performance differences in our four seating positions could have been due to different distances to the large display.

A second limiting factor was whether the ExD interface would have performed better had it leveraged spatial cues in the environment. In all seating positions, users moved their cursor in the same direction to transition between displays (top of laptop to bottom of large display). The choice to ignore the spatial relationships of the displays was made to reduce the experimental complexity imposed on participants. Because we analyzed only the last half of the trials in any seating position with any interface, we doubt

that our results will change with dynamic bindings, but to investigate this possibility we conducted an additional follow-up study with 8 users that repeated our initial design using static or dynamic bindings. With dynamic bindings, the spatial relationship of the displays was preserved in the mouse movement between displays. Our results (details in [27]) showed no systematic improvement of dynamic bindings on timing, workload, or preference.

IMPLICATIONS AND EXTENSION OF RESULTS

Although our results show that content redirection is a desirable approach for interaction design in MDEs, there are many outstanding issues which require further study.

Scaling up to Multiple Users

Our study was conducted with a single user to establish the interaction costs and benefits of redirecting content to a local display in a collaborative MDE. As the goal is to use content redirection in a collaborative setting, the next step is to see how content redirection scales to multiple users. We feel that the performance benefits seen in our results will transfer, but the social issues surrounding collaboration must be investigated. For example, if a user redirects content to their local display as opposed to redirecting input to a shared display, will that user feel greater ownership over the redirected content as they are interacting with it in their own territory? If so, do we need to provide interfaces that support this territoriality or provide interfaces to neutralize feelings of ownership over shared content?

Another important factor in collaborative MDE design is interface support for group awareness. We need to determine if content vs. input redirection negatively affects group awareness in collaborative MDEs. We have established content redirection as a plausible approach for collaborative MDEs, and will continue to investigate the social issues surrounding its use in collaborative settings.

Scaling up to Multiple Displays

Our study was conducted with one local display and one large display, and was designed to simulate environments where users bring their personal devices into an MDE with an existing large display. Investigating how our results scale to complex MDEs with multiple large, shared displays and multiple personal devices is an important next step.

Although the key press provided a fast transition for content and input redirection, when multiple displays are available will multiple keys need to be dedicated for transition? Will a key combination that scrolls through the options (e.g., "ALT-TAB" in Windows) see drastic performance losses with an increase in displays? Will benefits of physical navigation with mouse transitions and dynamic bindings appear in a more complex environment?

Our results suggest that researchers should consider content redirection when exploring these issues, while designers should provide both redirection options as a responsible interface design without an increase in interaction cost.

Realistic Tasks

Initial investigations into interface design decisions often include a low-level task such as targeting. Although these controlled investigations provide baseline information, studying real-use scenarios is essential for understanding user behavior. Our task involved some transition and some interaction, but in real tasks, the balance of these activities may vary. Some tasks (e.g. repeated copy and paste between the displays) may require numerous transitions with a small amount of interaction on the remote display between transitions, while other tasks (e.g. writing a paragraph on the shared display) will require very few transitions with lengthy interactions on the remote display.

We expect our results to generalize to tasks with similar ratios of interaction and transition, but the benefits of content redirection should be even more apparent for tasks with minimal transition and increased interaction. Given that real tasks generally have a greater cognitive component than our docking task, it will be important to investigate our results in more realistic scenarios.

Dealing with Screen Real Estate

In our study, the large screen and the local display had the same resolution. This will not always be the case, which would impact the choice to redirect content onto a local, lower-resolution display. Although the benefits of content redirection may decrease in this scenario, users may choose to redirect partial content, like an application window [5, 25], rather than an entire display. Or users may choose to redirect content to a local display for interaction in a tightly-constrained virtual space (e.g., editing a paragraph), and redirect input to a large display for tasks requiring a broader overview (e.g., working on the structure of an entire document). Providing both options gives the most freedom and power to users in their MDE. Further developing the interface to incorporate familiar window management tools would provide a functional environment for group work that conforms to current metaphors.

In terms of subjective workload, users preferred the content redirection-based interfaces. This is an interesting result given that these interfaces effectively cut the desk space of users in half and is most likely due to the sparse application layout used in our tasks. In settings where the user is monitoring applications, or has an otherwise cluttered local desktop this tradeoff may be undesirable.

CONCLUSIONS

Our study established significant performance improvements, as predicted, for users in suboptimal seating positions using content redirection. These improvements were also reflected in subjective workload assessments. A follow-up study eliminated display factors such as size and distance as possible causes, while a second follow-up study confirmed that static bindings did not systematically bias the Extended Desktop interface. In addition, we show that keyboard transitions were faster than mouse-based transitions in our particular task.

By concentrating on usability for a single-user in a variety of physical locations we have not only compared four interfaces in an identical context but also gained insight into how design choices affect pointing and selection within an MDE. We have established that content redirection allows users to interact with displays regardless of their seating position as effectively as if they were seated in an optimal position, and consequently enables users to simultaneously interact with shared content and work in a configuration of their choosing. Real-world MDEs will consist of a combination of fixed large displays and portable personal devices. Building interfaces that support users' choice in their physical collaborative arrangement, rather than forcing them to adapt to the limitations of the technology, is an important consideration for HCI researchers. Content redirection is one method that researchers should consider when designing interfaces for collaborative MDEs.

ACKNOWLEDGEMENTS

We thank NSERC and NECTAR for their support, and the members of the EDGE Lab for their insight and assistance.

REFERENCES

1. Ashdown, M., Oka, K., & Sato, Y. (2005). Combining head tracking and mouse input for a GUI on multiple monitors. *Ext. Ab. CHI*, 1188-1191.
2. Baudisch, P., Cutrell, E., et al. (2004). Mouse ether: accelerating the acquisition of targets across multi-monitor displays. *Ext. Ab. CHI*, 1379-1382.
3. Benko, H., & Feiner, S. (2005). Multi-monitor mouse. *Ext. Ab. CHI'05*, 1208-1211.
4. Berry, L., Bartram, L., & Booth, K. (2005). Role-based control of shared application views. *Proc. UIST*, 23-32.
5. Biehl, J.T., & Bailey, B.P. (2004). ARIS: an interface for application relocation in an interactive space. *Proc. GI*, 107-116.
6. Biehl, J.T., & Bailey, B.P. (2006). Improving scalability and awareness in iconic interfaces for multiple-device environments. *Proc. AVI*, 91-94.
7. Biehl, J.T., W.T. Baker, et al. (2008). IMPROMPTU: A New Interaction Framework for Supporting Collaboration in Multiple Display Environments and Its Field Evaluation for Co-located Software Development. *Proc. CHI*, 939-948.
8. Booth, K.S., Fisher, B.D., et al. (2002). The "mighty mouse" multi-screen collaboration tool. *Proc. UIST*, 209-212.
9. Ha, V., Wallace, J., et al. (2006). I want my MDE: configuring virtual workspaces in multi-display environments. *Ext. Ab. CHI*, 1481-1486.
10. Hawkey, K., Kellar, M., et al. (2005). The proximity factor: impact of distance on co-located collaboration. *Proc. GROUP*, 31-40.
11. Hinckley, K. (2003). Synchronous gestures for multiple persons and computers. *Proc. UIST*, 149-158.
12. Hinckley, K., Ramos, G., et al. (2004). Stitching: pen gestures that span multiple displays. *Proc. AVI*, 23-31.
13. Inkpen, K., Hawkey, K., et al. (2005). Exploring display factors that influence co-located collaboration: Angle, size, number, and user arrangement. *Proc. HCII*.
14. Johanson, B., Hutchins, G., et al. (2002). PointRight: experience with flexible input redirection in interactive workspaces. *Proc. UIST*, 227-234.
15. MacKenzie, I. S. (1992). Fitts' Law as a Research and Design Tool in Human-Computer Interaction. *Human-Computer Interaction*, 7, 91-139.
16. Rodgers, M.E., Mandryk, R.L., & Inkpen, K. M. (2005). Smart Sticky Widgets: Pseudo-haptic Enhancements for Multi-Monitor Displays. *Proc. Smart Graphics*, 194-205.
17. Miller, R.C., & Myers, B.A. (1999). Synchronizing clipboards of multiple computers. *Proc. UIST*, 65-66.
18. Myers, B.A., Stiel, H., & Gargiulo, R. (1998). Collaboration using multiple PDAs connected to a PC. *Proc. CSCW*, 285-294.
19. Nacenta, M.A., Aliakseyeu, D., et al. (2005). A comparison of techniques for multi-display reaching. *Proc. CHI*, 371-380.
20. Sommer, R. (1969). *Personal space: The behavior basis of design*. Cliff, N.J: Prentice - Hall.
21. Streitz, N.A., Geibler, J., et al. (1999). i-LAND: an interactive landscape for creativity and innovation. *Proc. CHI*, 120-127.
22. Su, R.E., & Bailey, B.P. (2005). Put Them Where? Towards Guidelines for Positioning Large Displays in Interactive Workspaces. *Proc. INTERACT*, 337-349.
23. Swaminatan, K., & Sato, S. (1997). Interaction Design for Large Displays. *interactions*, 4(1), 15-24.
24. Tan, D.S., Gergle, D. et al. (2004). Physically Large Displays Improve Path Integration in 3D. *Proc. CHI*, 439-446.
25. Tan, D. S., Meyers, B., & Czerwinski, M. (2004). WinCuts: manipulating arbitrary window regions for more effective use of screen space. *Ext. Ab. CHI*, 1525-1528.
26. Wallace, J. R., Ha, V., et al. (2006). Swordfish: user tailored workspaces in multi-display environments. *Ext. Ab. CHI*, 1487-1492.
27. Wallace, J.R., Mandryk, R. and Inkpen, K. (2008). Interaction with 2D Applications on Large and Small Displays. Technical Report 2008-01, Computer Science, The University of Saskatchewan, 1-17.
28. Wigdor, D., Shen, C., et al. (2006). Effects of display position and control space orientation on user preference and performance. *Proc. CHI*, 309-318.