

Effects of Display Position and Control Space Orientation on User Preference and Performance

Daniel Wigdor^{1,2} Chia Shen¹ Clifton Forlines¹ Ravin Balakrishnan²

¹Mitsubishi Electric Research Labs
Cambridge, MA, USA
www.merl.com
wigdor | shen | forlines @merl.com

²Department of Computer Science
University of Toronto
www.dgp.toronto.edu
dwigdor | ravin @dgp.toronto.edu

ABSTRACT

In many environments, it is often the case that input is made to displays that are positioned non-traditionally relative to one or more users. This typically requires users to perform interaction tasks under transformed input-display spatial mappings, and the literature is unclear as to how such transformations affect performance. We present two experiments that explore the impact of display space position and input control space orientation on user's subjective preference and objective performance in a docking task. Our results provide guidelines as to optimal display placement and control orientation in collaborative computing environments with one or more shared displays.

Author Keywords

Input-output mappings, spatial transformation, performance

ACM Classification Keywords

H.5.2 [User Interfaces]: Interaction styles

INTRODUCTION

In many contemporary computing environments, especially those where multiple collocated displays are used collaboratively, such as in the war rooms [1, 2, 7, 14, 27, 28], and in operating rooms [9, 17], users and their input devices are often not located directly in front of, or oriented toward, the display of interest (Figure 1). Technical solutions to the problem of allowing multiple participants to make input to multiple displays in such an environment have been examined by Johanson et al. [11]. Unexamined in the literature, however, is a usability problem created in such an environment: how should displays be positioned to optimise their use by multiple participants, and what mapping of pointing-device input to on-screen movement should be employed for any given screen position?

With the exception of direct-touch and tablet interfaces, it is usually the case that the location of user input is offset from

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.

CHI 2006, April 22-27, 2006, Montréal, Québec, Canada.
Copyright 2006 ACM 1-59593-178-3/06/0004...\$5.00.

the visual consequences of those actions. In the case of desktop computers, most users appear to easily handle the transformation of mouse movements on a horizontal surface to cursor movements on a vertical display. However, research [4, 13, 23, 24] has shown that performance of motor tasks can incur significant penalties under more dramatically offset input/output spaces, such as rotated mappings of up to 180°. While these penalties are reduced with practice, they are typically not completely eliminated.

It is important that designers of environments employing shared vertical displays consider these penalties. For example, if a system designer wishes to build a table-centred environment augmented with a single vertical display, where should that display be positioned to allow for optimal use by each participant seated around the table? Furthermore, given that it is impossible for such a display to be located directly in front of all participants, what is the appropriate input/output mapping? Although informative, it is difficult to apply the results of the previous research to these new multi-user scenarios because the experimental setups have typically positioned the display directly in front of the user, resulting in only a simple translational offset (as in the desktop computing scenario) plus the experimentally manipulated rotation of the displayed image.



Figure 1. The Southwest Securities Financial Markets Center at Baylor University is an example of a modern, table-centred environment augmented with multiple vertical displays [10].

Although it is intuitive that the traditional mapping of moving the arm forward in order to move the cursor upward on the vertical display is ideal when facing toward the display, it is not at all clear what happens when the display is moved such that users are no longer facing it directly. Is the ideal mapping body centric, such that forward movement should continue to move the cursor upward; or display centric, such that movement toward the display should move the cursor upward; or, something else entirely? The previous research suggests that selecting the wrong mapping can have a significant effect on performance, resulting in penalties in time and accuracy of well over 100% [3]. In this paper, we present two studies that investigate the effects of display space location and control space orientation on interaction performance. In the first study, we varied the position of the display and gave participants dynamic control over the position and orientation of the input space while performing a docking task. This enabled us to determine how users naturally orient their input spaces when confronted with displays in varying locations. In the second study, we forced participants to use a particular orientation of the input space, allowing us to determine performance in the more fixed configurations typically found in real environments. In combination, these experimental designs cover a broad range of possible display and control space location/orientation scenarios, and the results can help designers of collocated collaborative environments make informed choices as to the placement of shared displays and their input devices.

TERMINOLOGY

To facilitate discussion of these issues, we first define several terms that we will be used throughout this paper:

Display Space

The virtual display where the user sees the results of her input manipulations is defined as the *display space*. We define the display space to be a two-dimensional vertically-oriented plane, located at various positions around the table at which the user is seated. We assume that the display faces the user's body (but not necessarily facing the front or face of the user) such that the centre of the display is the point that is closest to the user. Figure 2 shows the labels we assign to the different display space positions in our experiments. For example, a display facing the user but located behind her while she is seated at the table is at the "S" (south) position; a screen facing the user located to her left is in the "W" (west) position.

Control Space

The area used by the user to provide input to the system is defined as the *control space*. In this paper, control space is a two-dimensional input plane with a particular position and orientation relative to the table on which it is located. We assume that the control-space is on a horizontal plane at right angles to the display space, similar to the situation in a standard desktop computer where movements of a mouse

on horizontal plane are mapped to a vertical screen. The position of the control space on the table can be varied.

Control space orientation refers to the rotational transformation of the control space about the axis perpendicular to the table surface (Figure 3). Note that control space orientation is relative to the top of the table, and not to the display space position. To distinguish between the two, we use compass directions (North, South, etc) for display space position, and angles (0° , 45° , etc) for control space orientation.

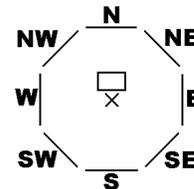


Figure 2. Display space position: location of the screen relative to the position of the user and table where input is made. The "X" marks the centre of the chair where the user is seated, the rectangle above it is the table on which input is made.

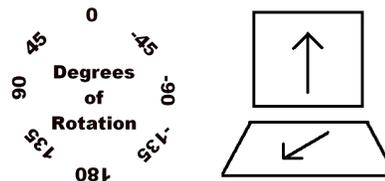


Figure 3. Control space orientation: the rotation of the control space about the axis perpendicular to the table surface. Left: labels used for the various orientations. Right: e.g., with a 135° orientation, to move up on the display (top) the user must move their hand back and to the left on table.

Example

On a standard desktop computer setup, the computer monitor has a display space position of "N", and the control space is the area on which the mouse operates, which is typically a mouse pad with a control space orientation of approximately 0° . The control space orientation can be dynamically changed by rotating the mouse. Figure 4 shows another example, which motivates our exploration.



Figure 4. Example of user orienting his control space to roughly 135° to accommodate a display in the S position.

RELATED WORK

We consider two relevant areas of prior work: collaborative systems where the positions of ancillary vertical display do not allow the familiar *N* display position relative to all users, and results from the psychology literature which have explored how rotationally transformed input-output mappings impact the performance of motor control tasks.

Collaborative Systems

The benefits of a table-centred interaction have recently been explored by several researchers [18-20, 22, 25, 29]. These include social advantages of collaborating face-to-face, the physical convenience of a shared tabletop display, and the enhanced ease of working with a direct-touch interface. Researchers have also found it advantageous to augment table-centred collaboration with additional large-screen displays. It is this pairing of table-centred collaboration with ancillary displays that is most likely to lead to the more extreme placements. When several users are seated around a table, an ancillary display positioned at *N* for one user might be *E*, *W*, or *S* to another. Although it may be possible to limit any adverse effects of display positioning by not allowing users to sit along some edges of the table, it cannot be entirely eliminated if all the advantages of table-centric interaction are to be gained. To date, there has not been any investigation as to the impact of ancillary display positioning in such spaces.

A study of a related issue was undertaken by Su and Bailey [26] who sought to determine the optimal relative distance and angle of two vertical displays when performing a docking task requiring the movement of the docking object from one display to another. In conjunction with this work, our results could be used to inform designers as to the optimal position of multiple displays, in addition to providing the optimal control space orientation.

In some of this previous work [25], users could move to the ancillary vertical displays in order to interact with them, while others advocated a more table-centric approach where all interaction occurs while users remained seated [11]. Nacenta et al. [16] compared techniques for manipulating virtual objects between a tablet and ancillary displays and found that the most effective was a radar view approach, where a portion of the tablet is used to control the input space of the ancillary display. Although they provide a thorough comparison of known techniques, this study only considered the situation where the display is placed at the *N* position, and the control space orientation is always at 0° . The problem is exacerbated when multiple users are seated around a table and it is physically difficult for everyone to be optimally oriented to one or even multiple surrounding shared display(s) (e.g., Figure 1 and [1, 2, 7, 9, 14, 17, 27, 28]). As such, their work does not provide guidance in the more general situation where ancillary displays may be positioned anywhere around the table and/or where the control space may be at a non 0° orientation.

Effects of Control Space Orientation on Performance

Psychologists and cognitive scientists have long studied the effect of distorting the orientation of control space relative to the display space. The earliest work, conducted with optical prisms mounted on eyeglasses, was conducted by Hemholtz [8] and Stratton [23, 24] in the 19th century. Both found that inverting the visual input to the eyes resulted in severe disruptions to motor activity, but that these disruptions were reduced over time.

Cunningham [3] sought to determine the structure of the mapping between the visual and motor systems. Participants sat at a table with a digital tablet, and performed a pointing task on a display oriented vertically and positioned at *N* (directly in front of the user). Performance was measured under control space orientations of 0° , 45° , 90° , 135° , and 180° . Participants were instructed to make straight-line movements between targets in the pointing task, and the effect of the control space orientation was measured as the deviation from the optimal trajectory. A control space orientation of 90° was found to be the most difficult, while 45° and 180° orientations had the lowest rate of spatial error relative to 0° , and the relative difficulty of the 135° orientation varied between participants.

In subsequent work, Cunningham and Welch [4] examined how people learned multiple rotated control spaces. Unlike the previous study [3], participants did not do whole blocks of tasks at a particular orientation, but instead switched back and forth between different orientations. They also measured the effect of different types of cues used to prime participants to the orientation used in the trial. They found that, with practice, these cues could significantly reduce the interference effects of switching between orientations.

GOALS OF THE PRESENT WORK

Although there is substantial subsequent research [6, 12, 21] that extends the work of Cunningham and colleagues, none has examined the issue of how control space orientation impacts performance under different display space positions. User preference for control orientation and display position has also not been investigated. Both these issues are of significant importance not only to the design of collocated collaborative environments but also to our basic understanding of human capabilities when faced with transformed input-output mappings that are more complex than the rotational offsets studied to date. We seek to explore these issues via two experiments, and specifically attempt to answer the following questions:

1. Which display space position do users prefer when given a choice? Which do they prefer the least?
2. Given a particular display space position, what control space orientation do users prefer?
3. Given that in real environments it may not be possible to position displays and orient control spaces to every user's preference, what is the penalty on performance if either or both of these preferences are not met?

EXPERIMENT 1: ADJUSTABLE CONTROL SPACES

Goals and Hypotheses

In this first experiment, we sought to answer our first two research questions: what are users' preferences with respect to display space position and control space orientation? We also partially explored our third question by asking participants to perform a task with the display space positioned at each of the eight possible locations (Figure 2) while allowing them to orient their control space as they wished. The impact of a fixed control space is explored in the second experiment.

Based on the results of previous experiments in the literature and our experiences with collocated table-centric multi-display environments, we formed several hypotheses:

- H1: Participants would most prefer the N display position.
- H2: Participants would least prefer the S display position.
- H3: Display space position would have a significant impact on the selected control space orientation.
- H4: Participants would generally orient their input space such that the traditional mapping of forward/up would be maintained (0° for N, 45° for NW, 90° for W, etc).
- H5: Display position would have a significant impact on performance.
- H6: Performance would be best at display positions most preferred by the participants.

Apparatus

The participant sat in a chair in front of a table, upon which was placed a DiamondTouch [5] tabletop multi-point input surface. Although the DiamondTouch is capable of acting as a touch-input device, we did not make use of this feature; instead, as was done by Cunningham [3], input was made using a stylus, effectively turning the DiamondTouch into a large input tablet. Since our intent in this experiment was to allow participants to orient the control space as they preferred, we built a simple plastic template designed to be manipulated by the non-dominant hand and tracked on the DiamondTouch. The position and orientation of one corner of the control space was mapped to the boundaries of this template, allowing participants to easily reposition and reorient the control space by moving the template appropriately. A ceiling mounted projector displayed a green rectangle on the DiamondTouch to provide the user with a visual indication of the extents of the control space. The control space was 17×13 cm while the DiamondTouch surface was 67×51 cm, thus allowing the participant to manipulate the control space over a reasonable area. The stylus was held in the dominant hand and its input was only effective when used within the control space controlled by the non-dominant hand. Figure 5 illustrates this apparatus.

For the display space, we used a large plasma screen (display area approximately 75×56 cm) positioned atop a wheeled cart. The display was placed at eight different

positions throughout the experiment, each of which was marked with tape on the floor, and positioned 140 cm from the centre of the chair upon which the user sat to perform the study. Figure 6 illustrates.

The software used for the experiment was written in Java, and was executed on a 3.2GHz Pentium PC running Microsoft Windows, which was disconnected from all network traffic. Both the plasma display and overhead projector were driven by a graphics card running at a resolution of 1024×768 pixels with a 60Hz refresh rate.

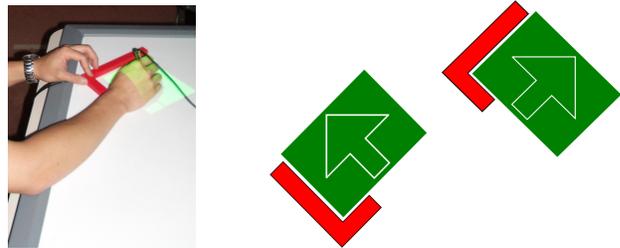


Figure 5. Control space and stylus input. (left): Position and orientation of the control space was achieved by manipulating a physical template held in the non-dominant hand. (right): Diagrammatic illustration of the relationship between the template and control space: the red L-shape represents the template while the green rectangle represents the control area. The arrow indicates the “up” vector for the control space (arrow is shown here for illustration only, and was not displayed during the experiment).

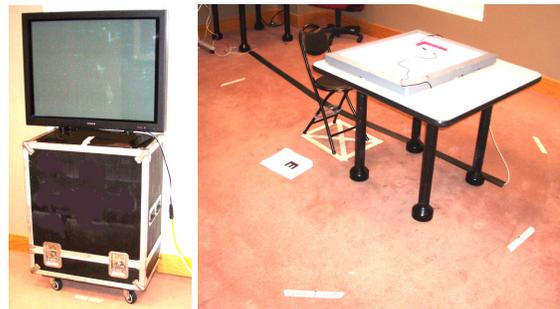


Figure 6. Display and input table. (left): plasma display on cart used to vary the display's position. (right): table and chair used in the study. The 8 possible display space positions, equidistant to the centre of the chair, were marked on the floor with tape and the cart placed accordingly.

Task and Procedure

There are three canonical tasks in a GUI: selection (moving a cursor to an object), docking (selection + drag of an object to a location), and path-following (e.g., steering through levels of a hierarchical menu, or drawing a curve with the pointer). We chose a docking task, since it encompasses the simpler selection task and also evaluates movement trajectories while giving participants freedom to move in the wrong direction and then make corrections to their path. This task also varies from Cunningham's work, where first a selection task [3] and then a path-following task [4] were used, thus our work contributes in terms of task variation.

The stylus, held in the participant's dominant hand, controlled the absolute on-screen cursor position. Selections were made by crossing the stylus over the desired object, and releasing by lifting the pen from the surface of the table. Docking tasks were grouped into several pre-computed "walks", which would begin with a blue square positioned at the centre of the screen. Participants would select this square and then drag it to the position of a larger red square "dock" which would change colour to a light blue to indicate success. Participants would then lift the stylus from the surface of the table to complete a successful dock. The red square would then move to a new location. The blue square remains in the same position so that it does not have to be reselected. The blue square is then dragged again to the red square's new position, and this process continues for four different locations of the red square. Thus, four docking tasks are accomplished in each such "walk". Figure 7 illustrates:

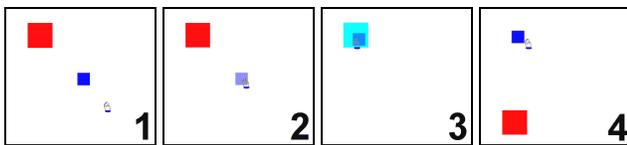


Figure 7. The docking task used in experiment 1: Participant touches the pen to the table, 2: crosses the blue object to select it, 3: drags blue object to the red dock, 4: when released the red dock moves to a new location.

Before beginning the experiment, the procedure and apparatus were explained, and participants were allowed to practice the task until they felt comfortable with it, which usually occurred within 30 practice trials.

We recorded the times taken to make the initial selection and perform the docking task successfully, control space orientation throughout the experiment, and the number of errors. An error occurred when the blue square was released outside the red square. To prevent participants from "racing through the experiment", they had to successfully complete each docking task, even if errors occurred, before the red square would move to a new position.

Participants

8 participants (6 male and 2 female) between the ages of 19 and 28 were recruited from a local university community and our lab, and paid \$20 for their time, irrespective of performance. All were right-handed, had completed at least one year of a Bachelors degree, and had little to no experience working with stylus, tablet, or tabletop input.

Design

Each participant performed 40 docking tasks for each of the 8 display positions. To counterbalance for learning effects, each participant began with a different display space position and then worked through the remaining 7 display positions in counter-clockwise order (i.e., participant #1 began with the display positioned at N, then NW, then W, etc; participant #2 began with the display positioned at NW, then W, etc).. Although we considered randomizing the

order of the display orientation sequence, previous work [3, 4] indicates that radical changes in orientation can result in significant temporary performance degradation before participants adapt to the new orientation. By using a sequential presentation, the orientation changed gradually, thus averting this temporary spike, and reducing the time required to adapt to the new orientation. This was important as our focus was to measure true, adapted performance at each orientation, rather than any transitional effects. Further, by starting each participant at a different orientation, each orientation appears in a different presentation slot per participant (i.e., P1 sees Order 1 first, Order 2 second, etc. P2 sees Order 2 first, Order 3 second, etc) resulting in a between-subjects counterbalanced presentation order. The direction of movement for each docking task was randomized but controlled such that movements in all 8 compass directions were performed an equal number of times by each participant at each display position. In summary, the design was as follows:

8 participants x
8 display positions x
40 docking tasks =
2560 total dockings

Results

Preferred Display Space Position

We administered a post-experiment questionnaire to answer our first question: what are the most and least preferred display space positions? Table 1 summarizes the results.

Participant	Most Preferred	Least Preferred
1	N	S
2	NE	S
3	NE & NW	S
4	NW	S
5	NE	S
6	NE	SW
7	NE & NW	SE & S
8	N	SW

Table 1. Preferred display space position for each participant

That most participants least preferred the S display space position confirms Hypothesis H2. This is unsurprising since S represents the greatest offset between hand and display positions, thus requiring the least comfortable posture. Participants cited body comfort as the primary reason for selecting S as their least preferred display space position. Much more surprising is that participants predominantly (75%) selected a display space location offset 45° from a traditional, N position. Although all participants were asked to provide an explanation for their selection, none was able to articulate their preference to their satisfaction: a typical response, stated by participant 3, was "it just feels better". Based on these results, we reject hypothesis H1, but note that this may well vary by input device, and suggest that designers consider the ergonomics of their input device before applying this particular result.

Preferred Control Space Orientation per Display Position

Our second research question, what is the preferred control-space orientation for a given display space position, was answered by allowing the participants to dynamically reorient their input space throughout the experiment and recording the results. Table 2 summarizes the average control space orientation each participant used across the entire experiment for each display position.

Display Location	Control Orientation (μ)/degrees	σ
S	27.2	45.8
SE	-35.1	21.3
E	-26.0	19.9
NE	-13.3	15.4
N	-4.7	14.0
NW	12.6	18.1
W	22.5	34.1
SW	36.9	32.7

Table 2. Mean control space orientation (μ) and variance (σ) across all participants for each display-space position.

Display position had a significant effect on the control space orientation selected by the user ($F_{7,49} = 9.032$, $p < .001$), thus confirming Hypothesis H3. There was also a significant user x control space orientation interaction ($F_{7,49} = 1181.21$, $p < .001$). Coupled with the high variance, seen especially at the more extreme screen positions, this suggests that significant individual differences between users may play a role in their preferred orientation. Figure 8 shows the mean control space orientation used by each user across the experiment per display space position.

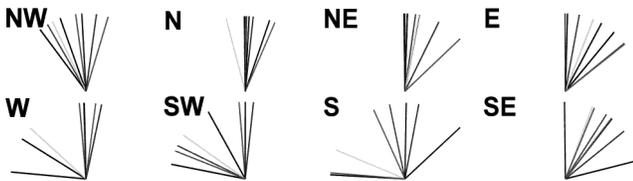


Figure 8. The direction of each line indicates the mean control space orientation for a participant for a display space position.

Predominantly, participants chose their orientation at the beginning of a block of trials at a given display-space position, and rarely changed their control space orientation during a block. We measured both inter-trial reorientation (changes in orientation between trials) and intra-trial reorientation (changes in orientation during a trial). On average, for each participant, when the first trial of each block was excluded, instances of inter and intra-trial reorientation in excess of 1° did not exceed 6 trials.

Figure 8 illustrates the general trend of orienting the control space in the general direction of the display. With the exception of participant 1, who kept the control space at an orientation of 0° for the entirety of the experiment, participants did not strictly maintain the traditional forward/up control space orientation, regardless of whether we consider forward/up to be away from the body or towards the display. We therefore reject Hypothesis H4.

Effect of Display Space Position on Performance

Although we only had eight participants, they completed a total of 2560 docking tasks. This large number allows us to make a number of statistically significant conclusions.

We first measured performance as the time required to perform the entirety of the task from selection of the blue square until it is successfully docked with the red square. Analysis of variance showed a significant main effect for display space position on task completion time ($F_{7,1786} = 10.74$, $p < .0001$), confirming hypothesis H5. Pairwise means comparisons across all participants revealed that performance at screen positions N, NW, E, and W were not significantly different from one another, but were from the rest, as was the case for S and SW; SE and NE were significantly different from one another and from the rest. However, the magnitude of the performance difference between positions was not very large, as shown in Table 3.

	S	SE	E	NE	N	NW	W	SW
Time	1255	1176	1112	1024	1088	1113	1110	1227
%	23%	15%	9%	-	6%	9%	8%	20%

Table 3. Average task completion time in msec for all participants for each display position, as well as how much slower (percentage) participants were to complete the task at that display position relative to the overall fastest one (NE).

We also examined the path traversed by the participants during the docking task. Unlike Cunningham [3], we did not instruct participants to attempt to move in a straight line when performing the task. As such the resulting paths in our experiment are more reflective of how users might perform such tasks in a real application, thus increasing the ecological validity of our results. Motivated by previous research on input coordination [15, 30], we computed the ratio of the total length of the actual path to the ideal minimum-length straight line path. This metric provides an indication of the amount of additional distance participants travelled by deviating from an ideal path. We recognize that this metric only considers path length and not the complexity of a path as might be, for example, measured by the number of turns. However, given that path complexity metrics are not the focus of our research, we chose to rely on the established [15, 30] path length ratio metric. There was a significant main effect for display space position on this path ratio ($F_{7, 1786} = 8.01$, $p < .0001$). There was also a significant correlation ($R^2=.72$) between this ratio and performance time, which is expected as larger rotations imply a longer path which require more time to complete. In combination, this further supports Hypothesis H5.

Errors were measured in two ways: a trial was deemed to have been erroneous if the participant released the blue square before placing it in the red dock, or if the blue square entered the red dock's area and exited again before being released. There was no significant effect for display space position or participant on either error metric.

Relationship of User Preference to Performance

Our third research question: what is the effect on performance of not meeting user preference with respect to display space placement and control space orientation can be partly addressed by the results of this experiment. Although the participants were able to adjust the control space orientation, they had to perform dockings with the display positioned at each of the 8 locations, and as such this experiment provides data as to what happens when display position is not at a user preferred location.

Interestingly, preference did not correlate with optimal performance: only 4 of 8 participants performed fastest with their preferred display placement, while only 2 of 8 participants had the lowest performance at their least preferred display location. We thus reject hypothesis H6

EXPERIMENT 2: FIXED CONTROL SPACES

Goals and Hypotheses

In the first experiment, our third research question: what is the effect on performance of not meeting user preference with respect to display space placement and control space orientation, was only partially explored in that we allowed users to manipulate the control space to their preferred orientation. In this second experiment, we further explore this question, this time using a fixed control space orientation that users could not alter. Thus, this experiment considers the situation that is common in real environments where both display position and control orientation are fixed and users have to work within the given parameters. We formulated the following hypotheses:

H7: Inability to adjust control space orientation will have a significant effect on performance.

H8: Performance at a given control space orientation will vary between display space positions.

Apparatus

The apparatus for this experiment was the same as in experiment 1, except that the physical template's orientation no longer affected the orientation of the control space. To compensate for the gap in feedback created by the removal of this pairing, we added a visualisation to the rendered control space: a gradient from green at the bottom (toward 180°) to blue at the top (toward 0°) of the space (creating a ground/sky effect). To provide the same positioning flexibility as in experiment 1, the template continued to control the position of the control space.

Task and Procedure

The task and procedure were virtually identical to those in experiment 1, except that participants were presented with a particular control space orientation, rather than being allowed to dynamically reorient the space.

Participants

8 participants (4 male and 4 female), different from those in experiment 1, between the ages of 19 and 25 were recruited

from the community, and paid \$20, irrespective of performance. All were right-handed, and had little to no experience working with stylus, tablet, or tabletop input.

Design

The task was performed for the 8 display space positions (Figure 2) and the 8 control space orientations (Figure 3). To reduce the time required to participate, the display control conditions were not fully crossed: each participant performed the task at 4 control orientations for each of 4 different positions. A Latin-square design was used to avoid ordering effects and to ensure that each display space position and control space orientation pairing occurred an equal number of times in the experiment. Because of the learning and interference effects observed by Cunningham [3, 4], we increased the number of docking tasks in each block from the 40 used in the previous experiment to 80. In summary, the design of the experiment was as follows:

8 participants x
4 display positions x
4 control orientations x
80 docking tasks =
10,240 total dockings

Results

Although only eight participants took part in the experiment, their completion of over 10,000 docking tasks at the various pairings of control space orientation and display space position allows us to make a number of statistically significant conclusions. There was a significant interaction between order of presentation of the control orientation and display space position pairs and task performance time ($F_{14,7124} = 31.285$, $p < .001$). This suggests that, as discussed by Cunningham and Welch [4], the transformed spatial mappings of control to display space were interfering with one another. We found that after the first 50 trials per condition, the order effect ceased to be statistically significant, indicating that with sufficient practice the prior spatial mappings ceased to interfere with the one currently being used. Accordingly, in the remaining analyses we consider only the last 30 trials per condition, treating the first 50 trials as practice.

There was a significant main effect for control space orientation on task performance time ($F_{7,1681} = 69.805$, $p < .001$), confirming Hypothesis H7. There was also a significant interaction between control space orientation and display space position on task performance time ($F_{26,1681} = 7.637$, $p < .001$), indicating that the effects of control orientation differ depending on display space position. This confirms Hypothesis H8. Also interesting was that the shortest times were seen to roughly correspond to the preferred range of control space orientations that users chose when given the ability to manipulate the control space in experiment 1. Figure 9 illustrates these effects.

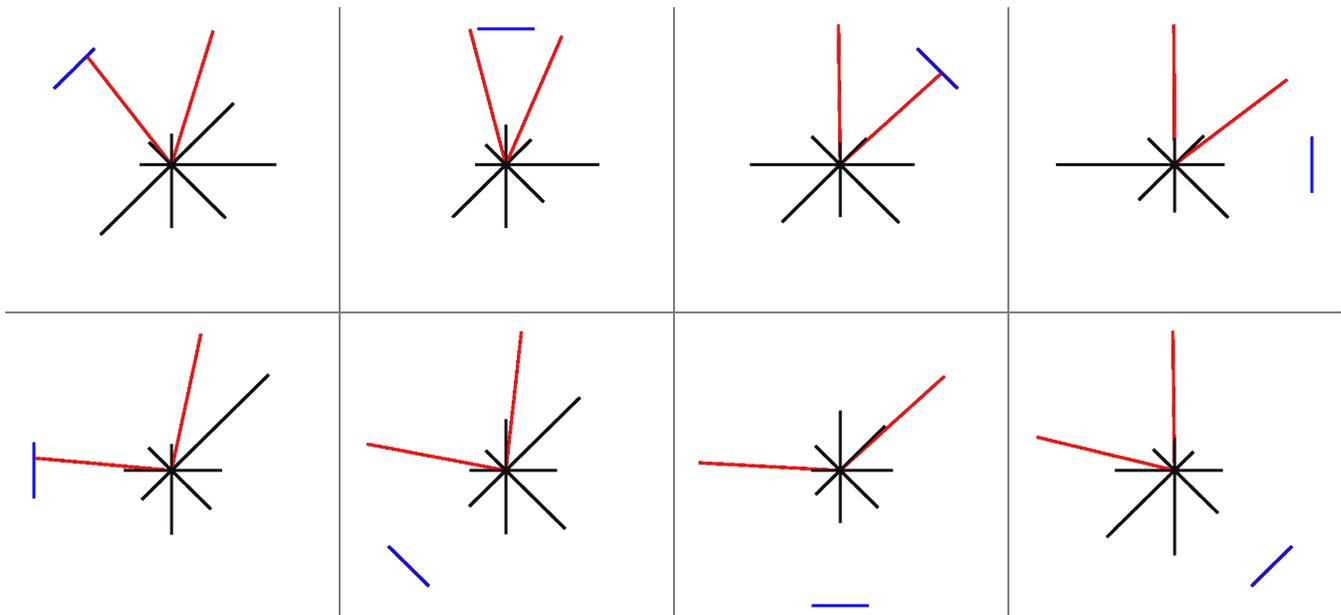


Figure 9. Mean task completion time at a given control space orientation encoded as the length of the line in that direction (longer the line the slower the performance). Display space position indicated by the position of the perpendicular line. Overlaid on each is the range of preferred orientations (longer lines) from experiment 1.

Interestingly, the correlation between actual path to optimal path ratio and task completion time was significantly lower ($R^2 = 0.23$) in comparison to the results of experiment 1. One possible explanation is that several users adopted what we have dubbed the *spiral strategy* to moving under a transformed spatial mapping: rather than attempt a seemingly optimal straight-line movement, they instead chose to move in circular motions. Because the control space was offset rotationally, a circular motion can be more easily anticipated than a straight line – moving in a clockwise circle in the input space produces a clockwise motion in the display space, no matter the control space orientation. Figure 10 illustrates this approach, where we see three distinct anticlockwise spirals as the pointer approaches the red square dock. Note that the blue square was moved very close to the dock near the beginning of spiral S2, but the participant elected to continue the spiral pattern. Although this spiral path clearly deviates from the optimal straight line path, participants who employed it reported that they felt it was faster than trying to learn the more difficult transformed mappings.

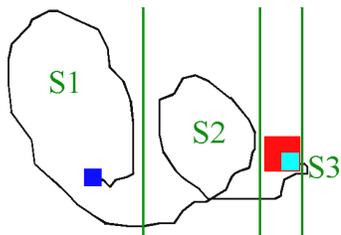


Figure 10. An actual path (in black) followed by a participant using a spiral-strategy to dock the blue square onto the red square under a transformed control – display mapping. Three distinct spirals (S1-S3) are visible.

DISCUSSION AND CONCLUSIONS

The results of our experiments lead to several interesting observations, and suggest design recommendations for designers of systems where a traditional display space position and/or control orientation is not possible,

The lack of correlation between preferred and best performing control space orientation in experiment 1 suggests that either participants are not able to correctly assess their performance, or, more likely, that they consider factors other than performance when determining preference. In particular, the absence of inter-trial reorientation suggests physical comfort may be more important than performance, since it is likely that the initial orientation of the control space was made to optimise comfort. That preference is more closely tied to physical comfort than performance is a likely explanation for the rejection of Hypothesis H6: that performance would be best at those display positions most preferred by the participants. Also interesting was that when asked for their least preferred display position, only 2 of 8 participants chose the position where their performance was worst. The rejection of hypothesis H1, that participants would most prefer the traditional N display space position, provides further evidence that participants were optimising for comfort. Accordingly, our finding that users least preferred the S display position is not surprising, since it requires the most effort to turn the body to allow them to see it. This trade-off between performance and comfort should be considered when designing multi-display environments.

Although significant individual differences were present, some general trends were visible with regards to how display space position influenced the choice of control space orientation (see Figure 8): for all of the east display

space positions (NE, E, SE), participants chose to orient their control space between 0° and -90° , or, generally, to the east. For all of the west display space positions, participants chose to orient their control space between -17.1° and 90° , or, generally, to the west. We suspect that the asymmetry between these two ranges may be due to the fact that our participant population was entirely right handed.

Although we did find a statistically significant effect for display space position on performance, there was on average a maximum 23% penalty when users were able to adjust their control space orientation as in experiment 1. As Figure 9 illustrates, there is a clear performance trend when participants are not able to adjust their control-space orientation. For those display spaces that are in front of them (NW, N, NE), a 45° offset in control-orientation from straight-on produces the best results. For the remaining positions, a 90° offset towards 0° is optimal. These results will be of use to designers of systems where physical constraints limit the users' ability to reorient their control space, such as in operating theatres where it is suggested [17] that if a monitor of a closed-circuit video feed is used by a surgeon to view the movement of her tools, and that monitor is placed directly in front of her, the video image should be rotated in to create a 45° control orientation.

In environments where input devices might be shared by multiple, disparately oriented participants, such as a table-centred environment, care should be taken to allow participants to make input to any ancillary displays at a desirable orientation. For systems with multiple participants collaborating using a single input device to control a vertical display, the data from our second experiment can shed some light on optimal display placement. For example, for a square table with four participants, there are four typical seating positions to be considered, as illustrated in Figure 11. Of the 2^4 possible permutations of user seating positions, 4 are of interest: (1,2), (1,3), (1,2,3), and (1,2,3,4), since all others are repeated cases of these. Table 4 shows, based on our experiment 2 results, the largest performance penalty experienced by any one of the users when the control space is oriented optimally for the given display space position (row) for each of the user position combinations (column).

These results indicate that, if a second surgeon is added to the same theoretical operating theatre described previously, facing the first and performing similar operations on the same patient, the video monitor should then be placed at either the W or E position, and the video rotated to create a control orientation of 45° for the surgeon to whom the screen is to the left, and a -45° control orientation for the other. From this data, it is also evident that for multiple users working in a war room such as the one described in Mark [14], the best arrangement for two participants is to be seated across from one another while using a vertical display located on either side (W or E). Also worth noting is the dramatic increase in penalties paid when moving from three users to four. For such environments, it may

become necessary for: (1) there to exist multiple control spaces, (2) that input to shared control spaces take into account which user is making the input and adjust the control orientation accordingly, or (3) that multiple participants are seated on the same side of the table rather than seating at all 4 canonical positions. How control spaces are shared and positioned is best determined by examining the environment, but it is clear that care should be taken to avoid a high-penalty configuration.

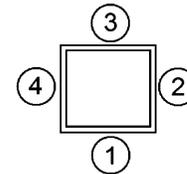


Figure 11. The 4 canonical positions for users seated at a table.

	1 & 2	1 & 3	1,2,3	1,2,3,4
S	53%	32%	53%	183%
SE	32%	38%	38%	273%
E	62%	18%	62%	183%
NE	46%	75%	75%	273%
N	62%	32%	62%	183%
NW	36%	38%	46%	273%
W	41%	18%	62%	183%
SW	38%	75%	75%	273%

Table 4. Best-case performance penalty (as a percentage of optimal, NE/0% pairing) for a display positioned at one of the eight possible positions (rows) for the given user position around a table combination (columns).

In summary, our work has explored the impact of display space position and control space orientation on user preference and performance. The results contribute to the literature on transformed input-output spatial mappings by investigating important transformations not previously tested. These results also allow designers to make more informed choices as to layout of shared displays in multidisplay environments.

ACKNOWLEDGMENTS

We thank our experimental participants, John Barnwell for technology support, Dr. Helen Cunningham for insights into her previous work, Edward Tse for general help, the CHI meta-reviewer for significant assistance in refining the paper, and the CHI reviewers for their insightful comments.

This study was partially supported by the Advanced Research and Development Activity (ARDA) and the National Geospatial-intelligence Agency (NGA) under Contract Number HM1582-05-C-0028. The views, opinions, and findings contained in this report are those of the author(s) and should not be construed as an official Department of Defense position, policy, or decision, unless so designated by other official documentation.

REFERENCES

1. Bentley, R., Hughes, J., Randall, D., Rodden, T., Sawyer, P., Shapiro, D., and Sommerville, I. (1992). Ethnographically informed systems design for air traffic control. *ACM Conference on Computer Supported Cooperative Work*. p. 123-129.
2. Covi, L., Olson, J., and Rocco, E. (1998). A room of your own: What do we learn about support of teamwork from assessing teams in dedicated project rooms? Amsterdam, The Netherlands: Springer-Verlag. 53-65.
3. Cunningham, H. (1989). Aiming error under transformed spatial mappings suggests a structure for visual-motor maps. *J. of Exp. Psychology: Human Perception and Performance*, 15(3). p. 493-506.
4. Cunningham, H.A. and Welch, R.B. (1994). Multiple concurrent visual-motor mappings: implications for models of adaptation. *J. of Exp. Psychology: Human Perception and Performance*, 20(5). p. 987-999.
5. Dietz, P. and Leigh, D. (2001). DiamondTouch: a multi-user touch technology. *ACM UIST Symposium on User Interface Software and Technology*. p. 219-226.
6. Dobbelsteen, J.v.d., Brenner, E., and Smeets, J. (2004). Body-centered visuomotor adaptation. *Journal of Neurophysiology*, 92(1). p. 416-423.
7. Heath, C. and Luff, P. (1992). Collaboration and control: Crisis management and multimedia technology in London underground line control rooms. *Journal of Computer Supported Cooperative Work*, 1(1). p. 24-48.
8. Helmholtz, H. (1866). Treatise on physiological optics.
9. Hindmarsh, J. and Pilnick, A. (2002). The tacit order of teamwork: Collaboration and embodied conduct in anaesthesia. *Sociological Quarterly*, 43(2). p. 139-164.
10. Baylor Photography: <http://www.baylor.edu/pr/>.
11. Johanson, B., Hutchins, G., Winograd, T., and Stone, M. (2002). PointRight: experience with flexible input redirection in interactive workspaces. *ACM UIST Symposium on User Interface Software and Technology*. p. 227-234.
12. Kagerer, F., J.Contreras-Vidal, and Stelmach, G. (1997). Adaptation to gradual as compared with sudden visuo-motor distortions. *Experimental Brain Research*, 115(3). p. 557-561.
13. Krakauer, J., Pine, Z., Ghilardi, M., and Ghez, C. (2000). Learning of visuomotor transformations for vectorial planning of reaching trajectories. *Journal of Neuroscience*, 20(23). p. 8916-8924.
14. Mark, G. (2002). Extreme collaboration. *Communications of the ACM*, 45(6). p. 89-93.
15. Masliah, M. and Milgram, P. (2000). Measuring the allocation of control in a 6 degree-of-freedom docking experiment. *ACM CHI Conference on Human Factors in Computing Systems*. p. 25-32.
16. Nacenta, M., Aliakseyeu, D., Subramanian, S., and Gutwin, C. (2005). A comparison of techniques for multi-display reaching. *ACM CHI Conference on Human Factors in Computing Systems*. p. 371-380.
17. Nardi, B., Schwarz, H., Kuchinsky, A., Leichner, R., Whittaker, S., and Sciabassi, R. (1996). V Turning away from talking heads: the use of video-as-data in neurosurgery. *ACM CHI Conference on Human Factors in Computing Systems*. p. 327-334.
18. Rekimoto, J. (2002). SmartSkin: an infrastructure for freehand manipulation on interactive surfaces. *ACM CHI Conference on Human Factors in Computing Systems*. p. 113-120.
19. Ringel, M., Ryall, K., Shen, C., Forlines, C., and Vernier, F. (2004). Release, relocate, reorient, resize: fluid techniques for document sharing on multi-user interactive table. *Ext. Abs. of the ACM Conference on Human Factors in Computing Systems*. p. 1441-1444.
20. Scott, S., Carpendale, S., and Habelski, S. (2005). Storage bins: Mobile storage for collaborative tabletop displays. *IEEE Computer Graphics and Applications*, July/August 2005. p. 58-65.
21. Seidler, R. (2004). Multiple motor learning experiences enhance motor adaptability. *Journal of Cognitive Neuroscience*, 16(1). p. 65-73.
22. Shen, C., Lesh, N., and Vernier, F. (2003). Personal digital historian: story sharing around the table. *Interactions*, 10(2). p. 15-22.
23. Stratton, G. (1897). Upright vision and the retinal image. *Psych. Review*, 4. p. 182-187.
24. Stratton, G. (1897). Vision without inversion of the retinal image. *Psych. Review*, 4. p. 341-360, 463-481.
25. Streitz, N., Prante, C., Muller-Tomfelde, C., Tandler, P., and Magerkurth, C. (2002). Roomware©: the second generation. *Ext. Abs. of the ACM CHI Conference on Human Factors in Computing Systems*. p. 506-507.
26. Su, R. and Bailey, B. (2005). Towards guidelines for positioning large displays in interactive workspaces. *Interact Conference*. p. 377-390.
27. Teasley, S., Covi, L., Krishnan, M., and Olson, J. (2000). How does radical collocation help a team succeed? *ACM Conference on Computer Supported Cooperative Work*. p. 339-346.
28. Tollinger, I., McCurdy, M., Vera, A., and Tollinger, P. (2004). Collaborative knowledge management supporting mars mission scientists. *ACM Conference on Computer Supported Cooperative Work*. p. 29-38.
29. Wu, M. and Balakrishnan, R. (2003). Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. *ACM UIST Symposium on User Interface Software and Technology*. p. 193-202.
30. Zhai, S. and Milgram, P. (1998). Quantifying coordination in multiple DOF movement and its application to evaluating 6 DOF input devices. *ACM CHI Conference on Human Factors in Computing Systems*. p. 320-327.