ChairMouse: Leveraging Natural Chair Rotation for Cursor Navigation on Large, High-Resolution Displays

Alex Endert
Virginia Polytechnic Institute
and State University
aendert@vt.edu

Patrick Fiaux
Virginia Polytechnic Institute
and State University
pfiaux@vt.edu

Haeyong Chung
Virginia Polytechnic Institute
and State University
chungh@vt.edu

Michael Stewart
Virginia Polytechnic Institute
and State University
tgm@vt.edu

Christopher Andrews
Virginia Polytechnic Institute
and State University
cpa@vt.edu

Chris North
Virginia Polytechnic Institute
and State University
north@vt.edu

Copyright is held by the author/owner(s).

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.
ACM 978-1-4503-0268-5/11/05.

Abstract

Large, high-resolution displays lead to more spatially based approaches. In such environments, the cursor (and hence the physical mouse) is the primary means of interaction. However, usability issues occur when standard mouse interaction is applied to workstations with large size and high pixel density. Previous studies show users navigate physically when interacting with information on large displays by rotating their chair. ChairMouse captures this natural chair movement and translates it into large-scale cursor movement while still maintaining standard mouse usage for local cursor movement. ChairMouse supports both active and passive use, reducing tedious mouse interactions by leveraging physical chair action.

Keywords

Large display, interaction design, embodied interaction

ACM Classification Keywords

H5.2. Information interfaces and presentation (e.g., HCI): Input devices and strategies.

General Terms

Design, Human Factors, Performance

Introduction

When appropriately designed, large high-resolution displays can be effective as personal, individual-user workspaces. Such single-machine desktop workstations (Figure 1) can be placed in a standard office, run standard operating systems and tools, and used primarily by individual users for their everyday work. Studies have demonstrated their effectiveness in domains such as intelligence analysis [1], cyber security [2], and other workspace applications. In these scenarios, the primary purpose of the display is not to create an enlarged projected view of information intended for collaboration or presentation. Instead, individual users can organize their information spatially, helping them gain deeper insight and perform their tasks more effectively.

With such workspaces, users *physically navigate* to access the information organized in their workspace, as opposed to relying strictly on traditional virtual navigation strategies (e.g. panning, zooming, etc.) [3]. Curving the workspace around the user (Figure 1) enables efficient access to all areas of the workspace via simple chair rotation [4].

However, the importance of the mouse cursor for interaction and the increased size and pixel density of these workspaces can create usability issues [5]. To illustrate this, we highlight two common large display usage scenarios that can cause problems. First, users *distribute tasks to different regions of the display*. Switching tasks is as simple as turning to a new region of the display, but the cursor is likely to be left behind, causing a disconnect between the focus of attention, and the focus of interaction, which frequently leads to the user loosing track of the cursor's location. Second, users *directly interact with information over large distances* (e.g. dragging a window), which can lead to excessive mouse rowing (or clutching) due to the long distances involved.

Techniques do exist that address these low-level concerns. To find the mouse, the user can activate visual cues that make the cursor easier to find (e.g. a large flashing target appears around the cursor). To move windows over large distances, add-on tools have been developed that can explicitly place windows in pre-defined locations. However, these techniques tax users with additional interactions such as pressing hotkeys to activate the visual cues or navigating menus of window management tools. Essentially, such extraneous interaction is slow and can draw users away from their "cognitive zone" (defined as a mental state where users are focused on their *task* rather than the *tool*) [6] because it does not directly relate to users' tasks or mental models.

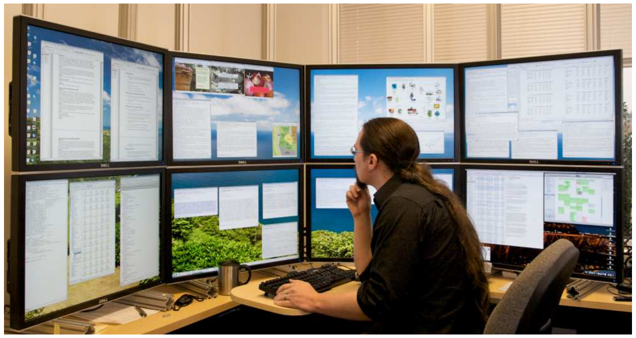

Figure 1 The large, high-resolution workstation (8 30" LCD panels, totaling over 32 megapixels) used in this study.

From observations made during prior studies of users performing everyday work on such displays [2], we view the users' natural chair rotation as an opportunity to transform already occurring physical *action* into meaningful *interaction* in the workspace. ChairMouse provides a second stream of mouse events based on the rotation of the chair that the operating system treats as an additional mouse. As the users rotate their chair clockwise or counterclockwise, ChairMouse generates relative mouse events, which move the cursor right or left, respectively. Additionally, users can continue to interact using their regular mouse for local cursor movement, as shown in Figure 2.

To explain the role we see for ChairMouse, we make the distinction between *active techniques*, which require explicit user intervention (e.g., clicking a button, typing a key sequence, etc.), and *passive techniques*, which leverages existing user behavior to carry out secondary tasks (e.g., keeping the cursor near the user's focus of attention). ChairMouse can be viewed as both an active and passive interaction technique, capable of addressing both scenarios mentioned above. The user can certainly consciously rotate the chair to move the cursor rather than using the mouse. However, it was envisioned to primarily be a passive technique. The user can remain cognitively engaged in a task, physically navigating to various regions of the display to access information, while ChairMouse uses the rotation of the chair to keep the cursor near at hand so that the user does not need to break out of the cognitive zone to relocate it. Since users are moving their chair anyway, we can exploit this passively. The advantage of this type of interaction is that it is essentially free, in contrast to performing previously described extraneous interactions required by the tool. Even in the active case where additional chair movement is used, chair movement closely relates to the task and easily carries over the coupling developed by the passive case.

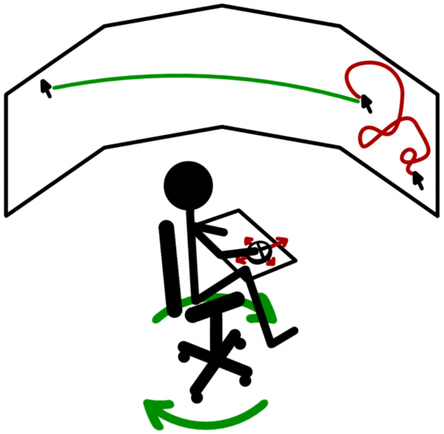

Figure 2 Illustration demonstrating how chair rotation (shown in green) accounts for large distance cursor movements, while regular mouse interaction (shown in red) for local cursor movement is maintained.

From observations during previous studies of the various types of tasks users perform on large displays during daily work, the line between the active and passive technique is fairly blurry. As users change their focus of attention, their task often requires a form of input as well. For example, when switching the primary task, users will require their cursor at corresponding location. In contrast, when referencing visible information, only a quick glance may suffice. As users become increasingly engaged in their task, peripheral cognitive activity, such as deciding on an interaction technique for the planned action, becomes a distraction to users. They should not decide when to use a specific active technique versus another given the situation at hand. Rather, the system should respond to both conditions (i.e. active and passive), without requiring the user to inform the system. ChairMouse offers such capability.

A passive technique also draws from embodied interaction principles, as it strengthens the connection between a physical action and an interaction, or what Klemmer et al. describe as a "relationship between body and task" [7]. When this relationship or link is established and maintained, users can become more embodied in their task. Further, when the physical action upon which the embodied interaction is based occurs naturally (i.e. users currently exhibit this physical action), the interaction becomes more usable. In other words*, physical actions that users perform anyway can be meaningfully leveraged and turned into passive interaction techniques, which then can be effectively exploited for active interaction*.

With ChairMouse, our goal is to unobtrusively leverage the user's chair movement and translate it into cursor movement, preserving the link between the mouse cursor and the user's attention with minimal user interaction.

Large Display Interaction Techniques

Research has been conducted in three main areas to attempt to alleviate mouse problems on large displays. Various ***input devices*** (e.g. joystick, trackpad, touchscreen panels, wand pointing devices [8], etc.) attempt to replace the traditional usage of a mouse with one more focused on large displays [9]. The "Multiple-monitor Mouse" [10] uses multiple mouse pads corresponding to specific monitors. Touchscreens [11] allow users direct access to information, removing the need to maintain cursor position. These techniques are each active techniques, as users have to explicitly activate or use them.

Using the chair as an interaction device has also been previously studied, as ChairIO allows users to interact in a game environment via physically leaning in specific directions on a stool, corresponding to view adjustment and player movement in the game [12]. However, this is clearly designed to encourage new behavior from the user rather than a technique to exploit exiting behavior.

Software solutions exist with the goal of reducing issues of cursor movement and loss by providing functionality accessible through additional commands and options. For example, "Missile Mouse" [9], "Bubble Cursor" [13], and "Drag-and-Pop" [14] are all techniques specifically developed for interacting with the mouse across large distances. Again, these are each active techniques, as users have to explicitly issue commands to activate them. Additionally, most operating systems include methods to adjust the cursor behavior slightly while maintaining standard mouse interaction. Cursor ballistics, enhanced pointer precision, and cursor speed are common options the users can adjust, however their effectiveness diminishes as the display size increases [15]. After initially changing these parameters, these techniques couple to regular mouse movement, much like a passive technique. However, as the tasks require different amount of detail in the cursor movement, the parameters must be explicitly changed again.

Awareness systems attempt to unobtrusively track users' focus of attention and map it to cursor position.

These systems are close to a purely passive technique, as they require close to no user action. Systems capable of gaze tracking exist, using cameras or head mounted tracking equipment to estimate the users' focus of attention and map it to cursor location [16-17]. However, permanently coupling the absolute cursor location to a user's gaze can result in the "Midas touch problem", which has been widely studied [18]. Proposed solutions include the use of heuristic or manually activated triggers for determining when to synchronize the cursor location with user's gaze. While heuristic-based approaches do not require direct user action, they are based on assumptions about general user behavior, which may not hold for all users and all tasks. Alternatively, triggers again require explicit user action. Hence, classifying these techniques as strictly passive would be inaccurate.

Peck et al.'s multi-scale cursor, on the other hand, is an example of a passive, tracking-based system [19]. Their system uses wand-style pointing, so cursor location is not an issue. Instead, the system tracks the user's distance away from the display and uses that the change the area of influence of the wand. In other words, as the user moves back from the display, losing details and perceiving larger scale units, the system makes a similar adjustment in scale, allowing the user to manipulate these larger scale units. Hence, the change in selection modes is done via the user's natural movement rather than explicitly changing selection modes with a menu. ChairMouse also exploits physical navigation, but for physical side-to-side panning rather than forward-backward zooming.

ChairMouse

ChairMouse is not designed to compete with previous active large display interaction techniques – it serves a different purpose and occupies a new place in the interaction design space. In fact it is orthogonal, and could be used in conjunction with those techniques. We believe this opens up an important new design space.

The fundamental design concept of ChairMouse is grounded in *leveraging natural physical navigation*. In doing so, we provide users the ability to interact with the display as they would anyway, preserving the *link between cursor location and user focus.* It is through this coupling that we can consider ChairMouse a *passive technique*.

ChairMouse provides users with a method that seamlessly integrates into their workflow, leveraging an action already being performed (chair rotation) in such a way as to reduce mouse rowing and cursor finding. We believe that chair rotation provides a reasonably good measure of where the user's focus of interaction is, and when the focus shifts, in a controllable fashion and at a stable level of granularity. Additionally, ChairMouse does not replace, but rather complements, standard mouse interaction. ChairMouse focuses on the underlying issue – *the input focus (i.e. cursor) is not where the user's focus of interaction is.*

Figure 3 ChairMouse with keyboard and mouse attached to armrests. A GyroMouse attached under the seat captures chair

Construction and Implementation

ChairMouse is constructed by attaching a Gyration GyroMouse [2] to a conventional office chair (shown in Figure 3). The GyroMouse maps angular movement to linear cursor movement. Thus, when the chair is rotated, mouse events are generated. GyroMouse can be permanently engaged as long as the front of the mouse does not sense a surface. Therefore, we constructed a cradle that holds the GyroMouse so that the front does not touch any surface, leaving it permanently engaged.

When mice generate events, the operating system treats them as relative X and Y coordinate changes, which ultimately update the cursor position. As a result, when multiple mice are attached to the same system, there is no competition, and each can change the position of the cursor. This means that while the rotation of the chair is being translated into cursor movements, the conventional mouse maintains full control while the chair is not moving.

The ability to treat the chair rotation as relative cursor positioning is advantageous for two reasons. First, calibration is not an issue. In absolute positioning systems using gaze tracking, initial and repeated calibration is often necessary to maintain an accurate positioning, even though work is being done to minimize this burden [20]. With ChairMouse, initial calibration is straight forward, as any perceived misalignment can be quickly corrected by the user with the conventional mouse, or by rotating to one of the display bounds. The cursor speed can also be adjusted in the GyroMouse preferences so the chair rotation matches the cursor speed. Secondly, position of the cursor is controlled via the regular mouse, as well as the chair's rotation. Hence, the cursor can, but does not have to, remain directly in front of the chair's orientation. Instead, users have the ability to move the cursor freely with the regular mouse, in addition to moving the cursor via chair rotation. Because the primary action of the chair is swiveling, the GyroMouse currently assists the user only in moving the cursor horizontally.

Evaluating ChairMouse

Evaluating a passive technique such as ChairMouse is challenging. Traditional performance and usibility metrics used for evaluation active interaction techniques, such as Fitts' Law, may not apply. For example, as one of the use cases of ChairMouse is task switching, it is expected that performance time will not improve. Instead, analyzing the amount of extraneous interaction or work (e.g. mouse movement, mouse rowing) may provide a more applicable measure of effectiveness.

We performed a user study, comparing ChairMouse to a traditional mouse setup for a spatial task (consisting of several sub-tasks in different locations) on a large, high-resolution display (shown in Figure 1). The user study consisted of 24 undergraduate students, divided into two groups (control and ChairMouse). The scenario provided to the users was modeled after previously seen usage scenarios [2], and consisted of eight tasks located at pre-defined locations on the workspace, each of which must be completed in a predetermined order. Thus, we could analyze the effectiveness of ChairMouse for each of the task switches. Each of the tasks involved organizing fragments of a story into a timeline. After performing this for each story, users were asked to write a short summary into a movable summary window. The purpose of the study was to shed light on usability issues of embodied passive and active interaction techniques for spatial interaction on large, high-resolution displays. We can show the following outcomes:

Reduced Work for Task Switching

We can show from analyzing task switches between the two groups that ChairMouse users moved the mouse 58% less ($t(20)=12.48536$, $p<0.0001$), and rowed their mouse 61% less ($t(20) = 6.962001$, $p<0.0001$) This is in large part due to the chair rotation occurring naturally, and for ChairMouse users, the rotation translated to cursor movement, reducing the amount of work required by the regular mouse.

Reduction of Extraneous Interaction

As one characteristic of a passive interaction is to not add to the work (or interaction) performed, we analyzed the amount of chair rotation performed by both groups to ensure that ChairMouse users did not simple replace regular mouse movement with additional chair rotation. For the task switches, all but two ChairMouse users rotated their chair to directly face the location of their new task. Of the two who did not, one hardly rotated his chair at all, while the other preferred to not directly face the task, favoring aiming the chair towards the middle of the display. Similarly, we found that of the 12 users in the regular mouse group, 11 of them rotated their chair to directly face their task when task switching. From this, we can show that ChairMouse indeed did not require more work or interaction. Thus, in reducing the amount of regular mouse movement (and rowing), while maintaining similar amounts of chair rotation, we can show the effectiveness of ChairMouse as a passive interaction technique.

Impact on Behavior and Workflow

More striking, however, is the shift in behavior with respect to the actual tasks at hand. This shift was seen during the final phase of the task, which consisted of summarizing each story using the *movable* summary window. To complete this phase, *all 12* ChairMouse users moved the summary window to the monitor corresponding to the story they were currently summarizing in order to type the summary into the window while referring to the story nearby. In contrast, *none* of the control group users elected to do so. By not moving the summary window when summarizing the independent stories, users were forced to choose between facing the summary window or the documents they were asked to summarize. This required the users to repeatedly glance back and forth between the story and the summary window. The common issues faced by the control group were: frequent typographical errors that went unnoticed (all 12 users), inserting summaries into the wrong story text box (7 users), and losing focus from the summary window, resulting in the text not being recorded (9 users). None of these errors were observed in the ChairMouse group. When asked about their choice to not move the summary window, control group users gave justifications such as "I am an expert typer and don't have to see where I type" or "I didn't see the need to". ChairMouse users, however, commented on their choice to move the summary window with comments such as "it makes sense to have the place where I'm typing next to the story" and even "I don't see why I wouldn't [move the window]".

We attribute the difference in behavior to the "soft constraints" that are present in the two conditions that are steering the microstrategies (i.e., low-level interactive behavior) adopted by the users. Unlike hard constraints, which dictate a particular sequence of interactions, soft constraints suggest the most likely patterns of behavior [21]. Both groups had the same goal – to type summaries into the summary window. Both groups could accomplish this by either moving the summary window or leaving it in a fixed position. The difference is that the control group would have to explicitly move the window to each new location, rowing repeatedly over the long spans, while the ChairMouse users could "drag along" the window as they turned to each new story by holding the mouse button while they were turning.

This subtle difference seems to have led to a difference in perception of the effort required to move the window around. Thus, in choosing the microstrategies for achieving the goal, the two groups diverge. As evidenced by the comments we received from the users, it seems that optimization choice was, as Gray and Fu maintain, non-deliberate, in that it the choice seemed clear to both groups of users [21]. The problem, of course, for the control group, was that their assessment was at such a low-level that it didn't take into consideration longer-term consequences, such as the interaction problems that can arise from interacting without visual feedback. While it is difficult to pinpoint whether the active or passive usage of ChairMouse caused this shift in perception, the distinct shift in perception of the amount of work required to move the summary window can most likely be attributed to both.

It is this conflict between low-level and immediate optimization and long-term goals that leads us to label actions like mouse rowing as extraneous (and perceived as "work"). Its presence contributed to the control group adopting a less effective strategy for accomplishing their task.

Discussion

As the results show, coupling relative cursor movement to chair rotation provides a good approximation of user attention, and thus cursor placement. Further, the "weight" of this action is largely what makes it effective. That is, we found that chair rotation only occurs when a shift in user attention occurs (e.g. a task switch). In comparison, eye/head tracking make use of heuristics or triggers to avoid over-tracking (i.e. the Midas' touch problem), as eye and head movement are much more frequent behaviors. The weight of

ChairMouse is light enough to track users' primary task switching, but heavy enough to avoid over-tracking. Task switching is an appropriate goal since local interaction is already handled quite well by the regular mouse, which ChairMouse allows for.

The correctness of the weight of ChairMouse also holds during conditions where a user may habitually "fidget" (e.g. tapping a leg or rotating the chair from side to side slightly as while thinking or typing). After observing one of the users in our study exhibiting this behavior, we asked him how he felt about the cursor moving slightly to reflect his fidgeting. He responded that he "did not even realize that was happening". While this could potentially be a distraction, it was not for this particular user.

Conclusion

Through this work, we present a new research space for large, high-resolution interaction: *passive interaction*. We motivate this space through discussion of previous active interaction techniques, and through this distinction present an instance of a passive technique, ChairMouse – capable of translating chair rotation into relative cursor movement. We evaluate the effectiveness of ChairMouse compared to a traditional mouse setup through a user study.

Our findings show that through leveraging the natural chair rotation of users, ChairMouse reduced regular mouse movement by 58% and reduced mouse rowing by over 61% during task switches in comparison to the traditional mouse only. Meanwhile, the total amount of chair movement remained about the same between ChairMouse and traditional mouse users – indicating that ChairMouse did not add any additional work to function. Additionally, we found that ChairMouse positively impacts users' workflow, resulting in users adopting better task strategies and making fewer mistakes. By giving users a more embodied interaction technique, they work more fluidly in the space, focusing not on minimizing mouse movement, but rather moving based on their task.

ChairMouse operates upon a user's natural embodied interaction with the large display. Therefore, the user is not forced to learn an additional interaction metaphor. All users in this study were able to learn how to use ChairMouse within a few minutes. From these results, understanding the distinction between active and passive interaction techniques is important, as design consideration should be given to support both uses. Finally, ChairMouse can be implemented easily and inexpensively using a GyroMouse and duct tape, and users can quickly benefit.

Acknowledgements

This work was partially supported by the Pacific Northwest National Laboratories, the NSF grant #CCF-0937071, and the National Geospatial-Intelligence Agency (NGA) under Contract Number HM1582-05-1-2001.

References

[1] Andrews, C., Endert, A. and North, C. Space to Think: Large, High-Resolution Displays for Sensemaking. *ACM CHI, 2010.*
[2] Fink, G., North, C., Endert, A. and Rose, S. Visualizing Cyber Security: Usable Workspaces. *VizSec, 2009*.

[3] Ball, R., North, C. and A. Bowman, D. Move to improve: promoting physical navigation to increase user performance with large displays. *ACM CHI, 2007.*
[4] Shupp, L., Ball, R., Yost, B., Booker, J. and North, C. Evaluation of Viewport Size and Curvature of Large, High-Resolution Displays. *Graphics Interface, 2006.*
[5] Ni, T., Schmidt, G. S., Staadt, O. G., Livingston, M. A., Ball, R. and May, R. A Survey of Large High-Resolution Display Technologies, Techniques, and Applications. *IEEE Virtual Reality, 2006.*
[6] Green, T. M., Ribarsky, W. and Fisher, B. Building and applying a human cognition model for visual analytics. *Information Visualization*, 8, 1 2009, 1-13.
[7] Klemmer, S. R., Hartmann, B. and Takayama, L. How bodies matter: five themes for interaction design. *Designing Interactive Systems, 2006.*
[8] Vogel, D. and Balakrishnan, R. Distant freehand pointing and clicking on very large, high resolution displays. *ACM CHI, 2005.*
[9] Robertson, G., Czerwinski, M., Baudisch, P., Meyers, B., Robbins, D., Smith, G. and Tan, D. The large-display user experience. *CG&A, 2005.*
[10] Benko, H. and Feiner, S. Multi-monitor mouse. *ACM CHI, 2005.*
[11] Guimbretiere, F., Stone, M. and Winograd, T. Fluid interaction with high-resolution wall-size displays. *ACM UIST, 2001.*
[12] Beckhaus, S., Blom, K. J. and Haringer, M. *ChairIO - The Chair-Based Interface*.
[13] Grossman, T. and Balakrishnan, R. The bubble cursor: enhancing target acquisition by dynamic resizing of the cursor's activation area. *ACM CHI, 2005.*
[14] Baudisch, P., Cutrell, E., Robbins, D., Czerwinski, M., T., P., Bederson, B. and Zierlinger, A. Drag-and-Pop and Drag-and-Pick: Techniques for Accessing Remote Screen Content on Touch- and Pen-operated Systems. *INTERACT, 2003.*
[15] Dasiyici, M. *Multi-Scale Cursor: Optimizing Mouse Interaction for Large Personal Workspaces*. Masters Thesis, Virginia Tech, 2008.
[16] Ashdown, M., Oka, K. and Sato, Y. Combining head tracking and mouse input for a GUI on multiple monitors. *ACM CHI, 2005.*
[17] Zhai, S., Morimoto, C. and Ihde, S. Manual and gaze input cascaded (MAGIC) pointing. *ACM CHI, 1999.*
[18] Velichkovsky, B., Sprenger, A. and Unema, P. Towards gaze-mediated interaction: Collecting solutions of the "Midas touch problem". *INTERACT, 1997.*
[19] Peck, S. M., North, C. and Bowman, D. A multiscale interaction technique for large, high-resolution displays. *UIST, 2009.*
[20] Ohno, T., Mukawa, N. and Yoshikawa, A. FreeGaze: a gaze tracking system for everyday gaze interaction. *Symposium on Eye Tracking Research & Applications, 2002.*
[21] Gray, W. D. and Fu, W.T. Soft constraints in interactive behavior: the case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science*, 28, 3 2004, 359-382.