Using Alternative Views for Layout, Comparison and Context Switching Tasks in Wall Displays

Anastasia Bezerianos[1,2]

[1]Department of Computer Science
University of Toronto

[2]NICTA
a.bezerianos@nicta.com.au

ABSTRACT

In this paper we first present a set of tasks that are relevant to wall display interaction. Among these, layout management, context switching and comparison tasks could benefit from the use of interactive shortcut views of remote areas of a wall display, presented close to the user. Such a shortcut view technique, the ScaleView portals, is evaluated against using a simple magnification lens and walking when performing these tasks. We observed that for a layout and comparison task with frequent context switching, users preferred ScaleView portals. But for simpler tasks, such as searching, regular magnification lenses and walking were preferred. General observations on how the display was used as a peripheral reference by different participants highlighted one of the benefits of using wall sized displays: users may visually refer to the large, spread out content on the wall display, even if they prefer to interact with it close to their location.

Author Keywords

Wall displays, interaction, alternative views

ACM Classification Keywords

H.5.2 User Interfaces: Interaction techniques

INTRODUCTION

Recently, the decreasing cost of projectors and LCD displays has enabled the construction of wall displays by tiling multiple projectors (Guimbretière et al., 2001; Baar et al., 2003) to form a single image. These high-resolution wall displays (resolution of more than 6000 x 2000 pixels) are interesting from an interaction perspective as they enable users to view high quality imagery even when they are up-close to the display (Figure 1). While interacting up-close with wall displays, the increased display real-estate is coupled with direct input affordances and benefits (Sears & Shneiderman, 1991). Such displays are becoming a commercial reality by companies (SMART Technologies http://smarttech.com) and will play a major role in many application domains.

Large sized displays have emerged as valuable tools in visualization and interaction with medical data (Hibbs et al., 2005), brainstorming and meeting facilitators (Elrod et al., 1992), awareness monitors (Greenberg and Rounding, 2001; Bardram et al., 2006), and industrial design canvases (Balakrishnan et al., 1999) among others. As the amount of manipulated information increases, these and other domains may benefit from larger displays.

OzCHI 2007 Proceedings, ISBN 978-1-59593-872-5

Research has identified benefits in using large displays in navigation (Ball & North, 2007), information memorization and recall (Tan et al., 2001), multi-tasking (Czerwinski et al., 2003), as well as general usability benefits (Ball & North, 2007). Despite these benefits, wall displays exhibit unique challenges, as they differ in scale from other direct input interactive surfaces. Some areas of the display are unreachable by the user without a substantial amount of motor effort. Interaction with these remote locations can be challenging, for example when users need to rearrange or switch focus between multiple objects outside arm's reach; or impossible, when areas of the display are inaccessible physically (for instance, very high up), or socially (blocked by another user).

In this paper we first identify a list of tasks that are likely to be performed on wall displays. We then discuss how three of them (layout, context switching and comparison tasks) relate, and how they may be affected by the large scale of the display. We propose using interactive shortcut views of remote display areas, called ScaleView portals (Bezerianos & Balakrishnan, 2005), to aid users in these tasks. The use of these alternative views is examined through a usability session that involves tasks similar to the ones identified in the first part of the paper.

Figure 1. A wall display of 5 x 2 m (16' x 6') with an effective resolution of 6144 x 2304 pixels. A user is manipulating a ScaleView portal (right) that provides a scaled-down view of a remote (left) area of the display.

WALL DISPLAY APPLICATIONS AND TASKS

A variety of direct manipulation tasks performed on physical walls and whiteboards have been identified early on as potential applications for vertical large displays, as seen in classrooms (Abowd et al., 1998), brainstorming and meeting settings (Fitzmaurice et al., 2005), planning (Trimble et al., 2003), and artistic or design processes (Buxton et al., 2000; Guimbretière et al., 2001). Organizational tasks traditionally performed on large physical surfaces, such as magazine layout or story boarding, will be likely transferred to the digital medium as the technology becomes more readily available.

Due to their dynamic nature and scale, wall displays have been used for time critical monitoring tasks, in military command and control centres (Dudfield et al., 2001), emergency response and traffic control centres, or large space security monitoring (often done through a multitude of tiled displays). Command and control units have been experimenting with touch sensitive digital walls (Jedrysik et al., 2000), to enhance such monitoring centres with direct interaction technology.

Large dynamic data visualization, for example in meteorology (Wilhelmson et al., 2002), or network traffic (Wei et al., 2000), is now performed on wall displays, because of the quality and quantity of data that may be simultaneously visible. Interactions in the form of exploration, navigation, visual search, data comparison and switching between different data views are essential in understanding and manipulation of the visualized data.

Finally, large displays emerged as awareness monitors and communication facilitators in hospitals (Bardram et al., 2006), offices (Fass et al, 2002) programming environments (Biehl et al., 2007) and research settings (Greenberg & Rounding, 2001; Huang & Mynatt, 2003).

These applications allowed us to identify a possible subset of frequently occurring interaction tasks likely to be performed on wall displays:

Content creation A large number of wall display applications (as in educational or brainstorming environments) will require content creation.

Content selection and movement As with all interactive environments, wall display interaction will also rely on selecting and moving content. This will be especially the case in educational or brainstorming settings, where new content is created often and periodically moved in different locations of the board as its nature and importance evolves over time.

Layout management Several of the mentioned applications involve organization of material scattered across the display, identifying layout management as another interaction task to consider. Examples of such tasks include arranging and rearranging story-boarding material and design mock-ups, creating and modifying the layout of magazine articles and images, etc.

Search and navigation Since a number of potential wall display applications involve the manipulation and exploration of high-resolution dynamic imagery, search and navigation will likely occur often. Examples of such actions could be following the course of a river on a terrain map or of a particular weather front on a meteorological map, or identifying proteins involved in a particular experiment in a protein connection graph.

Context switching Given the vast amounts of content presented, switching between different views of the displayed data will likely take place in numerous visualization tasks. For example in a protein connection graph, some aspects of the connections may need to be viewed in detail with respect to one experiment, while also viewing the context of the protein connections from other experiments; in terrain maps areas might be viewed in terms of distance from the sea level or of seismographic activity. Context switching will also be relevant in traditional windowing environments used on wall displays, with windows spread across the screen.

Comparison The dynamic nature and quantity of the displayed data on wall displays will require simultaneous views of different contexts in order to compare discrete visual elements, for example of such comparisons may be the need to contrast different meteorological fronts, or cracks in high-resolution structural photographs of construction material. Comparison tasks are a form of context switching, where both contexts need to be visible at the same time.

Monitoring Finally, some applications include a monitoring component of dynamic data. For example in monitoring security videos or real-time feeds from remote sources in command and control rooms.

Although this list may not be comprehensive, it enables us to reason about some of the issues that might arise due to the unique aspects of up-close wall display interaction and to design and perform appropriate evaluations. Already work has investigated some of these tasks. For example (Guimbretière & Winograd, 2001) discuss in-place content creation on wall displays. (Ball & North, 2007) showed that search and navigation tasks are more easily performed on wall displays. (Bezerianos et al., 2006) buffer invisible information to enhance monitoring tasks on wall displays). Also much work (discussed later), has focused on content selection and movement.

We will focus on layout management, context switching and comparison tasks. These share some characteristics and are similarly affected by the large display real-estate. With wall displays, space and layout management issues are somewhat different from regular desktop settings: much more data can be simultaneously displayed in a non-overlapping manner, but some of it will not be in the user's focal visual field, or easily accessible without substantial physical movement. Apart from frequent access and movement of content, organization and layout often requires detailed viewing of the organised content.

As it is likely that in wall displays application windows will not overlap one another much, but rather be spread out over the larger display surface, in-place context switching techniques need to accommodate frequent viewing or accessing of potentially hard to reach areas. Similarly, in comparison tasks, such viewing and accessing of different display areas need to be persistent in nature, and ideally presented close to each other and to the user to enhance user understanding.

While considering layout management, comparison and context switching challenges for wall displays, we observed that all the described tasks could be performed if different views of parts of the display were presented close to the user. This would minimize user physical effort and increase the visibility of remote locations: rearranging content over a large area could be easily performed if the virtual canvas was presented in a scaled down view, where between-content distances are more manageable; remote locations could be visible in detail if the user had a view of the remote location presented in a proximal area; comparison or focus switching between spread out content could be achieved by presenting close to the user views of that content. In our investigation of alternative views we use ScaleView portals (Bezerianos & Balakrishnan, 2005b) that provide this functionality.

PREVIOUS WORK

Issues such as unreachable content, were identified early on (Elrod et al., 1992) in work on large vertical displays integrating pen input. Later (Swaminathan & Sato, 1997) noted problems arising due of the scale of wall displays, like pointer movement and control challenges over large distances. They proposed using a small scale model of the display and its contents to specify pointer movement in the wall display, an effective mechanism both for comfortable reaching and layout management across large distances. Nevertheless, this approach alone cannot accommodate frequent context switching and comparison tasks, as content is presented scaled down.

Other work has focused on reaching across large distances. Approaches for copying content close to the user's location (Baudisch et al., 2003; Bezerianos & Balakrishnan, 2005a; Collomb et al., 2005) are designed for brief reaching actions, small in number, and cannot accommodate more persistent content organization and context switching actions. Research (Reetz et al., 2006; Forlines et al., 2006) has indicated that on large displays visual feedback close to the user is preferred in the case of detailed and precision based tasks, an issue affecting both throwing techniques (Geißler, 1998; Hascoët, 2003; Reetz et al., 2006), as well as input magnification techniques (Robertson et al., 2005; Forlines et al., 2006).

Apart from (Swaminathan & Sato, 1997) that proposed early on using a small-scale view of a wall display to specify pointer movement, other techniques based on alternative views of different display locations may facilitate organization and context switching between content spanning a large interaction surface. (Guimbretière et al., 2001) introduced the use of Zoomspaces, regions on the display with different zoom levels, an excellent space management technique, that nevertheless does not address content organization and layout over extended real-estate. (Tan et al., 2004) augmented window managers by using WinCuts, interactive views of arbitrary regions of existing application windows, allowing for fluid context switching and comparison tasks. Finally, (Khan et al., 2004) introduced a widget that acts as an interactive telescope to a remote area on the display, in order to access remote content, a technique that could be used for layout management, comparison or context switching tasks.

INTERACTIVE VIEWS OF REMOTE AREAS

ScaleView portals are essentially alternative views of different remote areas of a large display presented close to the user, at different scales. Interaction in the interior of a portal is equivalent to interacting with the area depicted inside the portal, which may reside in a remote location. Thus portals act as graphical and interaction shortcuts to the remote depicted areas (Figure 2).

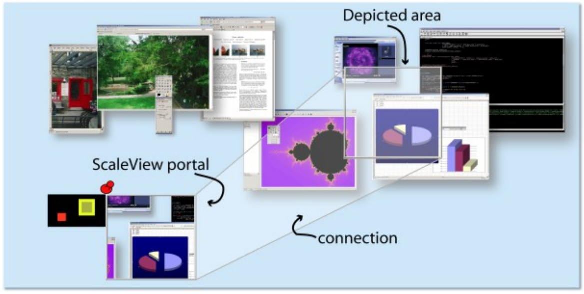

Figure 2. A ScaleView portal (bottom left) depicts a remote area on the display (right). The two are visually connected.

A transition zone (Figure 3) is defined between the portal and the main display. Unlike regular magic lenses (Bier et al., 1993), portals support the passing of objects between the portal, the main display, and other portals. When the centre of movement of an object (typically the cursor) crosses a portal's boundary, the object transitions into the portal and continues its movement inside the portal's coordinate system. The inverse also holds.

The portal provides views of areas of the display at different scales. It has three user-controllable parameters: it's position on the screen, the remote area it depicts, and a zoom factor that determines the scale of the depicted area. Users can reposition the portal by dragging it, or alter its zoom through a menu-invoked command.

The depicted area inside the portal can be altered in three ways. The most direct and precise method is using a menu-invoked command that allows users to define the remote area by walking and touching it. This method however, can be inconvenient if the desired area to depict is on a difficult-to-reach part of display.

Secondly, the user may select the depicted area at a coarser granularity, using a thumbnail representation of the entire virtual canvas we call a minimap. The minimap is attached to the top left corner of the portal and shows an iconic representation of the depicted area as well as the portal's position on the display. The user drags the icon representing the depicted area to reposition it. Finally, we have augmented the portals with a third focusing mechanism, traditional scroll-bars. These ways (Figure 4) of changing the depicted area inside the portal allows users to operate the portal without having to move around the display, and to reach areas at the display's extremities.

One issue that came up in a pilot study of ScaleView portals, was how to visualize the connection between the portal and the remote area that it depicts. Among several design alternatives we chose to draw a semi-transparent connection between the portal and the remote depicted area. To reduce clutter in the presence of multiple portals, this semi-transparent connection is visible while the user manipulates any of the portal attributes (position, scale, depicted area), or hovers over the portal (Figure 2).

In the pilot study we used different ways of presenting multiple portals and found that 2 visualizations work best: a stacking of portals that reduces clutter and side-by-side representation, allowing easy viewing and interaction with more than one depicted areas. We found that up to 4 portals are easily manageable, even in cluttered layouts.

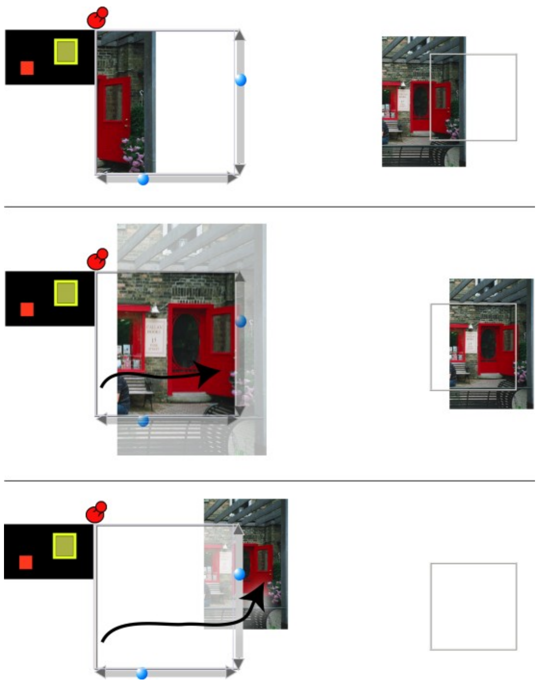

Figure 3. ScaleView portal transitions. (Top) On the left is the portal and on the right the remote depicted area. (Middle) As an object is moved towards the boundary of the portal, it is rendered semi-transparent outside the portal's confines to indicate that it can cross the portal boundary. (Bottom) The object has crossed the boundary, nevertheless it is still semi-transparent inside the portal confines, to indicate it may transition again inside the portal. Resulting effects of the transition in the remote depicted location are seen on the right. Illustrated arrows indicate user dragging.

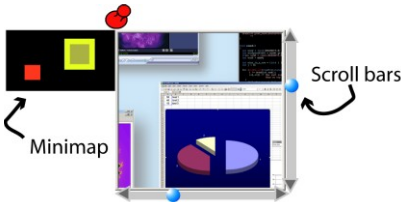

Figure 4. Elements for altering or focusing the area depicted in the portal. On the left an interactive minimap of the display that shows the location of the portal (red) and the depicted area (yellow). On the right and bottom familiar scroll-bars for altering the depicted area.

USABILITY STUDY

The goal of our study was threefold. First, we wanted to investigate the overall usability of ScaleView portals and gather feedback on their design. Second, we wanted to see if interactive views are a viable aid for layout management and context switching on wall displays, or if the complexity of setting up such views is prohibiting. Finally, we wanted to observe how users interacted and viewed the wall display, given that the portal essentially confines interaction close to the user.

Apparatus

For prototyping and testing, a 5x2 m, back projected display was used. Imagery was generated by 18 projectors (1024x768 resolution each) in a 6x3 tiling, for an effective resolution of 6144x2304 pixels. Projectors were driven by a cluster of 18 workstations. Software was written in C++ with Chromium providing graphics rendering over the cluster (chromium.sourceforge.net). A camera-based Vicon (www.vicon.com) motion tracking system tracked a pen's movement over the screen. Although the system could track the pen in 3D space, we used only x-y screen movements, a 10 cm hover zone, touch sensing and a single button for command invocation (a setup similar to most touch sensitive surfaces, like Wacom tablets, tabletPCs, Smartboards).

Users' interactions were logged by the system and an experimenter took notes during the session. Their visual focus was determined by tracking the users' head, a method that predicts eye gaze with 87-89% accuracy (Nickel & Stiefelhagen, 2003).

Tasks

Users were asked to perform a layout management and a comparison task, requiring frequent context switching:

Layout task: A set of images and text from different sources was present on the display, with text being readable without the need for magnification. Participants were asked to semantically group text and images as they saw fit. To ensure they didn't perform the task randomly, they were asked to justify their grouping to the experimenter. This open-ended task simulates real life layout management and organizational tasks, such as story-boarding, magazine layout etc. Furthermore it requires frequent context switching between the overview of the content to be organized and the details of each individual resource.

Comparison task: Participants were asked to locate four differences between two high resolution images of the western hemisphere. Differences were visible with the naked eye, but more prominent when using a magnification aid. This task resembles real life navigation and comparison tasks, such as looking high-resolution images for structural flaws, or examining different meteorological fronts, and then comparing them for similarities.

In our investigation we decided to add a search and navigation task, one that the technique was not designed for, to investigate its usefulness.

Search task: Participants were shown a high resolution image of the western hemisphere and were asked to pinpoint a small red square on the land mass (invisible without a magnification aid)5. This task resembles real life search and navigation tasks in different scales, such as examining high-resolution images, dens graphs etc.

Procedure

The tools available for performing the tasks were ScaleView portals, simple magnification lenses and walking. These techniques were explained to the participants in a brief training session. The training tasks were similar to the ones described above, and were performed using ScaleView portals, magnification lenses, and using no tool if possible (in the layout training task).

Once the training session was over, users could choose whether to use no tools, regular magnification lens(es) and walking, ScaleView portal(s) or combinations of the above to perform the given tasks. This allowed us to observe the use of the tools for the different tasks on wall displays. It also provided us with insights as to how users chose between a simple, easy to use technique (like a magnification lens) and a more sophisticated, but more involved technique (as the ScaleView portal).

As mentioned, one of the goals of this study was to explore the visual use of the wall display real-estate when confining the interaction close to the user (as is the case in the ScaleView portals). To enforce this localization we removed the option of directly selecting through a menu the depicted portal area for the purposes of the study.

Four female and six male participants took part in this session. The following sections summarize user trends while performing the different tasks.

Observations

Layout task: Only one user (code *v*) chose to walk around the display and directly drag content to organize it. She occasionally used a magnification lens to see detail.

The rest of the participants used ScaleView portals for this task. Two users (*iv,vi*), started the task using only a magnification lens, but very quickly switched to using two ScaleView portals, one zoomed-in to focus on particular items, and one zoomed-out to move items around. Nevertheless, these users did not make the connection of passing items "between portals", a functionality not explicitly shown to users in order to examine if the passing between portals felt "natural". Only passing from a portal to the main canvas was shown during the training session. These users moved the focus of the zoomed-in portal around the display using the scroll bars or the minimap instead of directly dragging items from the zoomed-out view. Another user (*x*), used a single ScaleView portal that she panned, zoomed-in to focus on specific content, and zoomed-out to organize content repeatedly. When asked later about passing content back and forth between portals, or between portal and canvas, all three users mentioned they did not recall or infer this functionality of the ScaleView portal.

One user(*i*) (Figure 5 top) was unique in using a combination of both a ScaleView portal and a magnifying lens. She placed the lens next to a zoomed-out ScaleView portal and would drag items from the portal under the lens to view in detail and then return them in the portal to arrange the layout. This user grasped the notion of transitioning between the portal and canvas. When asked later why she did not use a zoomed-in portal, she mentioned that she wanted the two tools to be separate and not attached, as in the current implementation.

Another user (*ii*) did not use a magnification tool. Her strategy was to move items from a single zoomed-out ScaleView portal to the main canvas, look at them and then organize them inside the ScaleView portal.

The remaining four users (*iii,vii,viii,ix*), made use of two portrals, one zoomed-out for organizing content, and one zoomed-in to view content in detail, and passed content between the two portals (an example seen in Figure 5 bottom). This set of users took advantage of the mechanism for passing content between portals.

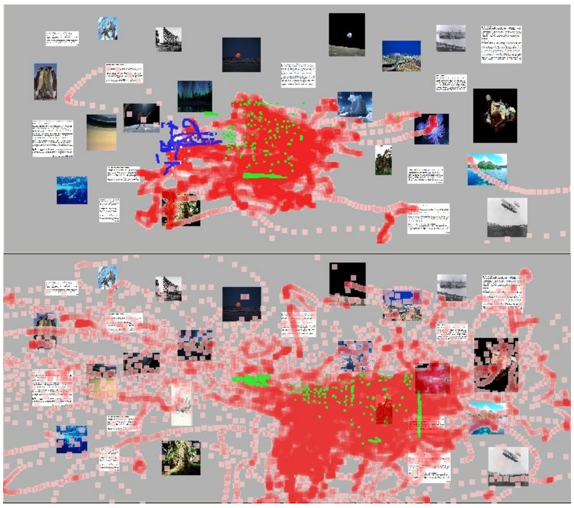

Figure 5. Interaction footprint of 2 users performing the layout task. Interaction with a portal is marked in green, whereas with magnification lens in blue. Red marks represent the user's visual focus. The top user (*i*) used a combination of portals and magnification lens to perform the task. She spent little time looking outside the portal and lens. The second user (*viii*) used 2 portals, making extensive use of the entire display as visual reference.

In this layout task 9 out of the 10 users preferred using the portals to complete the task, indicating that these interactive shortcuts are well suited for layout management tasks. Nevertheless, 3 users did not perceive the passing of content from ScaleView portals to other portals or the main canvas as natural, did not recall or make use of it. Although the remaining 6 users did take advantage of the feature, stronger visual indications of the passing mechanism would serve as a reminder to users.

Comparison task: This was the most challenging task, since it required detailed examination and visual comparisons between remote display locations.

The majority of users completed this task using ScaleView portals (Figure 6 top). One user (*i*) originally invoked two magnification lenses, one for each globe (Figure 6 bottom). Another (*viii*) invoked a single magnification lens that she attempted to move back and forth between the two globes. Both users switched to using two side-by-side ScaleView portals, each one focused on one of the two images. Another two users (*ii,iii*) originally used a magnification lens to focus on one globe and a ScaleView portal to focus on the other. The first completed the task as she started (*ii*), while the other (*iii*) reverted to using two side-by-side ScaleView portals. The remaining six users started and completed the task using two side-by-side ScaleView portals. 9 of the 10 participants completed the task using the portals.

Search task: In this task, the majority of users (8 out of 10) initiated the search using a regular magnification lens, (*i,ii,iii,iv,v,vii,viii,x*). They commented that magnification lenses were easier to use than ScaleView portals. Nevertheless, 3 of these 8 users (*i,ii,iv*) switched to using a single ScaleView portal when they could not quickly perform the task using the simple magnification lens. One of these three, in Figure 7, stated that her search was "not as systematic" with a regular magnification lens, as when using the scrolling functionality of the ScaleView. The remaining 2 users (*vi,ix*) started and finished the task using one ScaleView portal because, as one mentioned ”scroll-bars make it easy to navigate the enlarged screen”.

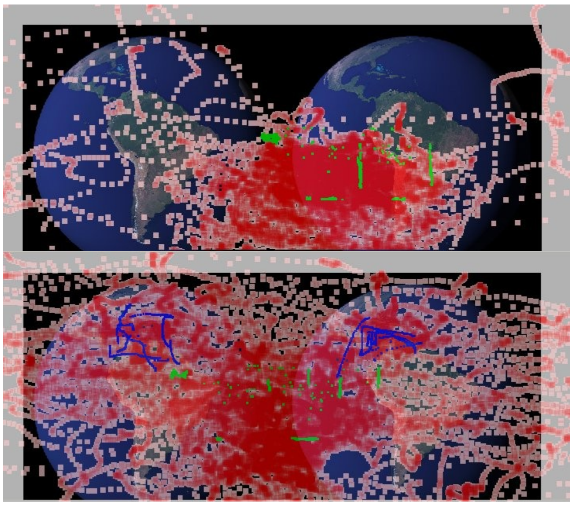

Figure 6. Interaction footprint of users performing the comparison task. Green denotes portal interaction, blue magnification lens movements and red is the user’s visual focus. The top user (*vii*) used 2 portals to perform the task and made small use of the display as reference. The second user (*i*) used a magnification lens (blue) and then switched to 2 portals (green), using extensively the display as reference.

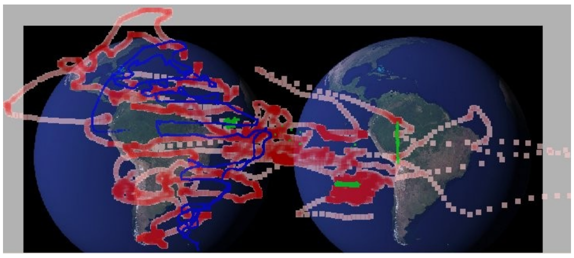

Figure 7. Interaction footprint of a user (*i*) performing the search task. She used first a magnification lens (blue) and then switched to using a portal (green). After invoking the portal the user focused little outside the portal (red marks).

Although ScaleView was preferred by 2 users, most users (8/10) started the task using a simple magnification lens. Interestingly, 3 of these users switched to ScaleView portals. So 5 users completed the task using the portals.

Display use

One question we had regarding using interactive shortcuts (like the portals), was whether users would focus on the portals where interaction took place, or use the overall display as a reference. We note that glances at the display are usually brief, whereas glances on the portals are usually accompanied by interaction that takes more time. Thus, if users spend more than 70% of their time on the portals we claim that they mainly focus on them and did not make heavy use of the large visual display. This decision is backed up by written comments taken by the researcher during the different sessions, where users observed to frequently glance at the main display canvas spend less than 70% of their time focusing on the portals.

From the 9 users that performed the layout task with ScaleView portals, 4 (codes *ii,iv,vi,viii*) used the display as a reference often (percent of time looking at portals less than 70%), whereas five (*i,iii,vii,ix,x*) performed the task by glancing little, if ever, outside the ScaleView portal area (time looking at portals more than 70%). The percentage of time spent looking inside ScaleView portals per user is seen in the first column of Table 1.

Table 1. Percent of time spent looking inside ScaleView portals per user for each task while portals were active.

	% of time looking at ScaleView		
	Layout	Comparison	Search
i	77%	47%	33%
ii	47%	50%	26%
Iii	75%	31%	0%
Iv	59%	85%	20%
V	27%	89%	0%
Vi	64%	29%	54%
Vii	96%	83%	0%
Viii	58%	76%	0%
Ix	88%	80%	70%
X	77%	68%	0%

Similarly, in the comparison task, users utilized the display real-estate differently. Some (*iv,vii,viii,ix*) mostly concentrated on the views inside the ScaleView portals (time looking at portals more than 70%), whereas others (*ii,vi*) made heavy use of the display as a reference for the location of the depicted portal area (column 2 of Table 1).

In the search task, 2 users started the task using portals. Of these two users, one (*vi*) spent time glancing at the display as a reference for the position of the focal area of the ScaleView, whereas the other (*ix*) mostly focused on the displayed content of the ScaleView. The percentage of time spent on looking inside the ScaleView for each user can be seen in the third column of Table 1.

User comments and design improvements

The majority of users (8 out of 10) stated that overall they preferred ScaleView portals for the given tasks. This is not surprising, given that users performed the greater part of the tasks using ScaleView portals, since, as stated, they allowed ”more systematic” interaction and users ”didn’t have to walk around much”. Nevertheless, user comments and observations led to design enhancements.

The major concern expressed was that of focusing on the depicted area inside the portals. In terms of creating a new portal, one user asked for the option to have new portals created as ”duplicates” of an existing portal. This is the default behaviour of the technique that was deactivated during the study to promote navigation and highlight any potential usability issues with the focusing techniques. Another user asked for the option to use one portal as a means to focus another. So if a zoomed-out portal was active, the user could make an area selection inside it to define the depicted area of a new portal.

Two users asked for the option to directly reposition the focal area of a portal at a remote location by walking. The design of the original technique actually provides a menu option for directly focusing the portal. This option was deactivated to confine interaction close to the user.

Interestingly users proposed different alternatives for focusing on the desired depicted area of portals. Users that tended to concentrate their view inside the ScaleView portal, proposed ways to focus it from inside other portals. While those that made use of the entire display as a reference, suggested ways involving direct positioning.

Apart from our existing focusing mechanisms (minimap, scrollbars, "duplicating" new portals from older ones and a menu option for directly refocusing the tool), we added the following functionality to accommodate both types of users: First users may now create a new ScaleView portal by determining its depicted area from inside an existing one. Thus portal focusing and creation are combined in a single action. Second, while the user is manipulating the remote depicted area in the minimap, the representation of that area on the minimap is rendered semi-transparent. When the representation reaches the edges of the minimap, it is rendered outside the confines of the minimap, similar to how items are drawn when they are about to cross the borders of a portal. This is an indication that users can continue dragging the representation outside the minimap. At the crossing point, the representation becomes the remote depicted area itself, which the user can now directly position on the canvas.

The presentation of multiple portals invoked comments from several participants. One user expressed her desire to be able to separate different portals and move them around independently. To address the concern of rearranging the position of individual portals, we now allow multiple separate portals, not necessarily grouped.

Since a number of our study participants did not recall using the "crossing" mechanism between portals and canvas, our design did not highlight this feature adequately. Originally, as a means to indicate the "crossing" option, items that where at the edge of a portal, were rendered semitransparent outside the portal (Figure 3). This indication might have been too subtle, given that part of the moving items had to actually cross the borders of a portal. To enhance the crossing metaphor we added the following behaviour: when a moving window comes close to a border, the border becomes progressively transparent (the more so as the item comes closer), until it disappears when the item starts crossing the border (Figure 3). We felt this visual indication is subtle enough to not interfere with the user's interaction, while highlighting the crossing property.

In the comparison task 1 user had some trouble telling the portals apart and would confuse which portal was focused on which globe. Not surprisingly, she made use of the display to disambiguate the focal points, but also spent a large amount of time looking at the two portals to compare them. The visual representation of the remote focal area was not enough for her and possibly more persistent visualizations of the connection between portal and depicted area, or some extra labelling of portals, would have been beneficial. To this end we the ability to add a written label to portals, for easier identification.

Finally 2 users asked for the option to synchronise the focusing of portals that are side-by-side. This is most likely a side effect of our choice of the comparison task.

CONCLUSIONS AND DESIGN GUIDLINES

Wall displays present unique interaction opportunities since they can display vast amounts of data, and can both be viewed and interacted upon at close proximity. As they become a commercial reality, interaction challenges arising from these unique aspects need to be examined.

We feel these challenges can be better identified through an exploration of potential tasks likely to be relevant in wall display interaction. After examining a set of existing and potential applications for wall displays, we identified a set of such tasks, which may be used as a reference point for future interaction research and evaluations.

By considering the identified tasks, we found that layout management, context switching and comparison tasks could benefit from the use of interactive views of remote areas in the display. An existing such technique, the ScaleView portals, was slightly altered and investigated under the proposed tasks, as well as a search task.

Apart from general design improvements, our findings indicated that complex behaviour in a wall display technique (like the ScaleView portals) is often preferred to simple functionality that requires extended motor effort (for example walking). So in a layout and a comparison task that required frequent context switching between different, potentially remote, areas of the display the vast majority of users preferred the ScaleView portals. Nevertheless, designers should note that in less intensive tasks, like searching, most users opted for a simpler technique, the magnification lens and walking. But even in the searching task some users switched to ScaleView portals when they felt the simpler techniques did not adequately support their task. Thus a combination of a lightweight and a more involved technique might be best.

Moreover, we observed general trends in how different users use the display visually, by tracking their visual focus. Our findings indicate that although many users opted for a technique that minimised motor effort (some explicitly stated this fact), many users suggested interaction with the entire display for determining the area to depict inside a portal. Thus designers should be advised that, whenever possible, they should provide users with the option to move to interact with the display, as well as an alternative to do so close to their location. An example of such alternatives is the ways we propose for determining the depicted area of a portal.

Finally, we observed that participants made different use of the display as a peripheral reference while using a portal. Nevertheless, rarely, if ever, did they simply concentrate on the portal alone, highlighting one of the benefits of using wall sized displays: users visually refer to the large, spread out content on the wall display, even if they prefer to interact with it close to their location.

ACKNOWLEDGMENTS

Thanks to Ravin Balakrishnan for his comments. NICTA is funded by the Australian Government's Department of Communications, Information Technology and the Arts and the Australian Research Council through Backing Australia's Ability and the ICT Centre of Excellence program.

REFERENCES

G. Abowd, C. Atkeson, J. Brotherton, T. Enqvist, P. Gulley, & J. LeMon. Investigating the capture, integration and access problem of ubiquitous computing in an educational setting. In CHI'88, p.440–447.

J. Baar, T. Willwacher, S. Rao, & R. Raskar. Seamless multi-projector display on curved screens. In EGVE, p.281–286, 2003.

R. Balakrishnan, G. Fitzmaurice, G. Kurtenbach, & W. Buxton. Digital tape drawing. In UIST'99, p. 161–169.

R. Ball & C. North. An analysis of user behavior on high-resolution tiled displays. In INTERACT'05 p.350–364.

R. Ball & C. North. Move to improve: promoting physical navigation to increase user performance with large displays. In CHI'07, p. 191-200.

J. Bardram, T. Hansen, & M. Soegaard. Awaremedia: a shared interactive display supporting social, temporal, and spatial awareness in surgery. In CSCW'06, p.109–118.

P. Baudisch, E. Cutrell, D. Robbins M., Czerwinski, P. Tandler, B. Bederson, & A. Zierlinger. Drag-and-pop and drag-and-pick: Techniques for accessing remote screen content on touch- and pen-operated systems. In INTERACT'03, p.57-64.

A. Bezerianos & R. Balakrishnan. Vacuum: facilitating the manipulation of distant objects. CHI'05, p.361–370.

A. Bezerianos & R. Balakrishnan. Canvas portals: View and space management on large displays. In IEEE CG&A, 25(4):34–43.

A. Bezerianos, P. Dragicevic, & R. Balakrishnan. Mnemonic rendering: an image-based approach for exposing hidden changes in dynamic displays. In UIST'06, p.159–168.

J. Biehl, M. Czerwinski, G. Smith, & G. Robertson. Fastdash: a visual dashboard for fostering awareness in software teams. In CHI'07, p.1313–1322.

E. Bier, M. Stone, K. Pier, W. Buxton, a& T. DeRose. Toolglass and magic lenses: The see-through interface. Computer Graphics, 27, p.73–80, 1993.

W. Buxton, G. Fitzmaurice, R. Balakrishnan, & G. Kurtenbach. Large displays in automotive design. In IEEE CG&A, 20(4):68–75, 2000.

M. Collomb, M. Hascoët, P. Baudisch, & B. Lee. Improving drag-and-drop on wall-size displays. In GI'05, p.25–32.

M. Czerwinski, G. Smith, T. Regan, B. Meyers, G. Robertson, & G. Starkweather. Toward characterizing the productivity benefits of very large displays. In INTERACT'03, p.9–16.

H. Dudfield, C. Macklin, R. Fearnley, A. Simpson, & P. Hall. Big is better? human factors issues of large screen displays with military command teams. In People in Control, p.304–309, 2001.

S. Elrod, R. Bruce, R. Gold, D. Goldberg, F. Halasz, W. Janssen, D. Lee, K. McCall, D. Pedersen, K. Pier, J. Tang, & B. Welch. Liveboard: A large interactive display supporting group meetings. In CHI'92, p.599–607.

A. Fass, J. Forlizzi, & R. Pausch. Messydesk and messyboard: two designs inspired by the goal of improving human memory. In DIS'02, p.303–311.

G. Fitzmaurice, A. Khan, G. Kurtenbach, & G. Binks. Cinematic meeting facilities using large displays. IEEE CG&A, 25(4):17–21, 2005.

C. Forlines, D. Vogel, & R. Balakrishnan. Hybridpointing: fluid switching between absolute and relative pointing with a direct input device. In UIST'06, p.211–220.

J. Geißler. Shuffle, throw or take it! working efficiently with an interactive wall. Ext. Abstracts CHI '98.

S. Greenberg & M. Rounding. The notification collage: posting information to public and personal displays. In CHI'01, p.514–521, 2001.

F. Guimbretière, M. Stone, & T. Winograd. Fluid interaction with high-resolution wall-size displays. In UIST'01, p.21–30.

M. Hascoët. Throwing models for large displays. In HCI'03, p. 73–77. British HCI Group, 2003.

M. Hibbs, N. Dirksen, K. Li, & O. Troyanskaya. Visualization methods for statistical analysis of microarray clusters. In Bioinformatics, 6(115), 2005.

E. Huang & E. Mynatt. Semi-public displays for small, co-located groups. In CHI'03, p. 49–56.

P.A. Jedrysik, J.A. Moore, T.A. Stedman, & R.H. Interactive displays for command and control. Aerospace Conference Proceedings, 2:341–351, 2000.

A. Khan, G. Fitzmaurice, D. Almeida, N. Burtnyk, & G. Kurtenbach. A remote control interface for large displays. In UIST'04, p.127–136.

K. Nickel & R. Stiefelhagen. Pointing gesture recognition based on 3d-tracking of face, hands and head orientation. In ICMI'03, p.140–146.

A. Reetz, C. Gutwin, T. Stach, M. Nacenta, & S. Subramanian. Superflick: a natural and efficient technique for long-distance object placement on digital tables. In GI'06, p.163–170.

G. Robertson, M. Czerwinski, P. Baudisch, B. Meyers, D. Robbins, G. Smith, & D. Tan. The large display user experience. IEEE CG&A, 25(4):44–51, 2005.

A. Sears & B. Shneiderman. High precision touch screens: Design strategies and comparisons with a mouse. International Journal of Man-Machine Studies, 43(4):593–613, 1991.

K. Swaminathan & S. Sato. Interaction design for large displays. Interactions, 4(1):15–24, 1997.

D. Tan, B. Meyers, & M. Czerwinski. Wincuts: manipulating arbitrary window regions for more effective use of screen space. In CHI'04, p.1525–1528.

D. Tan, J. Stefanucci, D. Proffitt, & R. Pausch. The infocockpit: providing location and place to aid human memory. In PUI'01, p.1–4.

J. Trimble, R. Wales, & R. Gossweiler. Nasa's merboard: An interactive collaborative workplace platform. In Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies, p.18–44, 2003.

B. Wei, C. Silva, E. Koutsofios, S. Krishnan, & S. North. Visualization research with large displays. In IEEE CG&A, 20(4):50–54, 2000.

R. Wilhelmson, P. Baker, R. Stein, & R. Heiland. Large tiled display walls and applications in meteorology, oceanography, and hydrology. In Proc. Int. Conf. on Interactive Information, p.29–30, 2002.