

Spotlight: Directing Users' Attention on Large Displays

Azam Khan, Justin Matejka, George Fitzmaurice, Gordon Kurtenbach

Alias

210 King Street East, Toronto, Ontario, Canada

{ akhan | jmatejka | gf | gkurtenbach }@alias.com



Figure 1. The Spotlight technique on a large wall-sized display (enhanced image).

ABSTRACT

We describe a new interaction technique, called a *spotlight*, for directing the visual attention of an audience when viewing data or presentations on large wall-sized displays. A spotlight is simply a region of the display where the contents are displayed normally while the remainder of the display is somewhat darkened. In this paper we define the behavior of spotlights, show unique affordances of the technique, and discuss design characteristics. We also report on experiments that show the benefit of using the spotlight a large display and standard desktop configuration. Our results suggest that the spotlight is preferred over the standard cursor and outperforms it by a factor of 3.4 on a wall-sized display.

Categories and Subject Descriptors: H.5.2 [User Interfaces]: Graphical User Interfaces (GUI), Windowing Systems

Additional Keywords and Phrases: large displays, attention, field of view, user study.

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 2005, April 2–7, 2005, Portland, Oregon, USA.

Copyright 2005 ACM 1-58113-998-5/05/0004...\$5.00.

INTRODUCTION

Large wall-sized display systems have existed for special applications like process monitoring, scientific visualization, and cinema for some time now. Only recently, the price and quality of data projectors has reached the point where it is becoming very affordable to couple many projectors and PCs together to produce large, high quality display landscapes. These economics are enabling meeting rooms to be constructed with multiple projectors and display surfaces rather than the typical single projector and screen. We have constructed and are studying one such room—the Alias Visualization Studio (see Figure 1), where we are experimenting with using multiple projectors to support a variety of types of real meetings and events. In the course of this work we have encountered new interaction challenges peculiar to large wall-sized display not addressed by traditional desktop techniques.

There are two main uses of large displays: (a) multiple people working in parallel on the display, and (b) speaker-audience style presentations. In this paper, we explore (b). Specifically, we examine the issue of directing the attention of multiple people simultaneously viewing data on a very large contiguous display space during a presentation.

During meetings or presentation in our Visualization Studio we have observed that users frequently like to convey an area of interest to other users to focus their attention on a particular area of the displays. There are several ways this can be accomplished:

Verbally - the user can describe either the target or its position verbally to assist others in finding it. This is effective when there are only a few images or blocks of text displayed on the wall. However, when the number of objects displayed is sufficiently high, or they are difficult to uniquely describe, using a very detailed verbal description for specifying the area of interest may become quite cumbersome.

Spatially – the user can spatially point to the target. There are many ways to accomplish this:

The user can simply point with their hand. This is a two step process: the viewers must first focus their attention on the person pointing and then the display. However this is a very coarse style of pointing and may not be accurate enough if the display is dense. Accuracy can be improved but at the cost of walking closer to the target area.

The system cursor can be moved to the target but the cursor is difficult to see across all screens especially from the distance where people are sitting in the space. In addition, to those not controlling the cursor, it is not clear where to expect to see the cursor resulting in the need to scan the full display space. A laser pointer can alleviate some of these problems but it is difficult to keep stationary for prolonged periods, hard to hold steady, not always available, and arguably the laser light can be visually annoying.

Regardless of the strategy used we have observed that when the display space extends beyond the users field of view and users cannot immediately find the target, users do not know whether a new target is already in their field of view and they should examine the display more carefully, or if they should re-scan the full width of the large display area. When this occurs, users may have to re-orient themselves, turn their head to scan the space, and quickly move their eyes to sweep over the display to attempt to find the target. In worst case, a user conducts a visual scan over the entire width of the display.

One way of avoiding rescanning is for the viewer to always keep the target area in sight. For example, a user could follow the laser pointer dot as it travels over the displays. This can be impossible if the pointer jumps or moves too fast or becomes hard to distinguish from the display (for example, moving the cursor over a moving video image). Furthermore, worthwhile distractions, such as looking down at one's notes or looking at other areas of the screen, breaks this tracking process, so rescanning is a necessity.

THE SPOTLIGHT

To address these challenges, we have developed a GUI interaction technique we call the “spotlight” that is based on how a spotlight is used in theatrical productions to direct an audience's attention.

Our GUI Spotlight consists of the interior region, the spot edge, the exterior region, and the cursor (see Figure 2). The

Spotlight technique, once invoked, begins by darkening the contents of the display (about 75%) except for a large circular region around the current cursor position. This serves to draw the viewer's attention to the current region of interest. The spotlight follows the cursor as it is moved. Unlike a real physical spotlight, while the spotlight is moving, the dark exterior region of the spotlight is brightened up (to about 50%) allowing the user to more easily see the intended destination. After the spotlight stops moving, after a short dwell time, the exterior region of the spotlight again slowly darkens to the maximum darkness setting.




Figure 2. Components: The spotlight consists of a darkened exterior region, a fully transparent inner region, and a cursor.

To combine the two functions of pointing with the cursor and repositioning the spotlight, we made the spotlight behave like a *tracking menu* [17]. Tracking menus have the interesting property of allowing the cursor to move about freely within a given region but when the cursor moves beyond the edge of the region, it moves the region in the direction of the cursor movement. In this way, the user can move the cursor to point to a specific location within the spotlight, as well as drag the spotlight to a new position, combining the specification of both a position and a region into one tool (see Figure 3). Moreover, the diameter of the spotlight can be resized by scrolling the thumbwheel on the mouse.




Figure 3. Simplified State Transition Diagram.

A spotlight can have an optional “beam” coming down from the center of the display (see Figure 4). This option is intended to give the user a well-defined starting point to reacquire the spotlight. The user knows that when the spotlight is not in their field of view, they can simply look

toward the center of the display and follow the beam, which quickly guides them to the target spotlight region.




Figure 4. Searchlight Mode: A beam is drawn from the center of the display to the spotlight to further assist the user in quickly acquiring the target.

INITIAL IMPRESSIONS

When we first experienced the Spotlight in action we had several unexpected observations. First, dimming the entire display (72 x 9 feet) dramatically changed the lighting and hence the atmosphere of the room. The lighting in the room changed from an ambient white glow (from the scattering of light from six powerful projectors) to a pleasing, calm, darkened environment reminiscent of a theater. Images in the spotlight appeared sharper and with better color contrast (due to the reduced scattering of light). Also, the scattering of light from the spotlight onto real objects in the room such as tables, chairs, and people also occurred. Thus, even if the spotlight was not in one's field of view, one could still sense the general location of the source of light and hence quickly "know" its location.

IMPLEMENTATION

We implemented a variety of spotlight techniques in a prototype application written in C++ under Windows XP and used the standard Windows API for event handling. In this way, we were able to easily run the prototype on a standard PC using a mouse for input. As we planned on using advanced drawing features such as semi-transparency, stencil masking, and back-buffer drawing, we use the OpenGL API to render the workspace and spotlight.

The prototype application runs on a semi-circular projection wall in 6 panels running on 6 workstations with screen space measuring 72' wide by 9' tall and a combined resolution of 7680 by 1024 pixels (see Figure 5). To achieve this, we used the Chromium networking software (derived from the Stanford Graphics Lab WireGL project [24] code base) for scalable real-time OpenGL rendering on a cluster of workstations. The application is run on a 7th machine, the so-called "mothership," where the input is collected and the OpenGL calls are intercepted by Chromium. The OpenGL calls are then divided, packed, and sent to one of the other 6 machines, the so-called

"servers." These servers each drive a projector and are responsible for rendering and displaying 1/6th of the total display. The mothership is an HP dual Xeon 2.8GHz PC with an nVidia Quadro FX1000 graphics card, while dual Xeon 3.2GHz processors, and one of either an nVidia Quadro FX1000 or an ATI FireGL X1 graphics card powers each of the servers. All of the servers are controlled from one HP workstation using VNC.

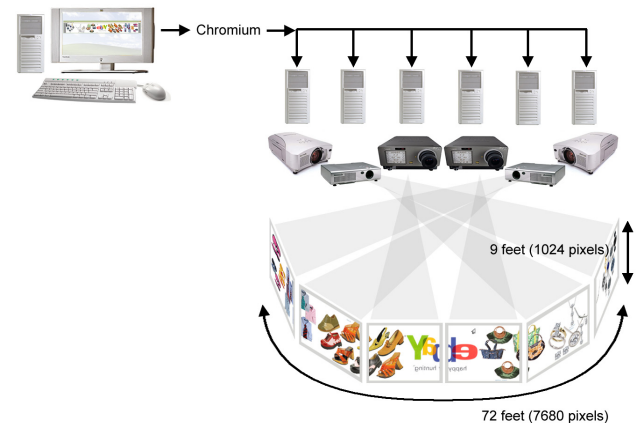


Figure 5. Wall-sized Display Hardware Configuration. Total screen size is 648 sq. ft. (72 ft. by 9 ft.).

The spotlight effect is created by first masking the region of the screen where the spotlight will be shown. The geometry is then rendered as rectangular polygons, quadric disks, and/or irregular polygons (depending on the desired shape of the spotlight) to the stencil buffer. The soft edges are rendered by iteratively moving each vertex of the geometry perpendicular to the local tangent and away from the center and re-rendering to the screen buffer with graduated transparency. One large rectangle is then drawn over the entire screen to cover the areas not lit by the spotlight. The stencil buffer ensures that this final rectangle is not drawn where the gradation and the interior region was drawn.

RELATED WORK

Visual attention has been studied in both HCI and cognitive psychology [31]. Attentional cues are central in decisions about when to initiate or to make an effective contribution to a conversation or project [23]. Most of this work has examined implicit factors of this issue with only a handful of systems looking at explicit mechanisms, like the spotlight, to guide visual attention.

Researchers have studied the effects of managing user's attention between their primary activity and on peripheral information displays [6, 11, 19, 28, 30, 33]. Moreover, there have been a number of studies conducted on the effects of visual separation and physical discontinuities for distributing information across large displays or multiple monitors [37, 39, 40] as well as the effects of field of view for 3D navigation tasks [14, 38].

Furthermore, there has been a great body of research investigating interaction techniques for large displays [10, 15, 20, 21, 29, 35, 36] and interaction techniques that span multiple monitors [4, 8, 16, 18, 19, 22, 25, 40]. However, few have addressed the explicit goal of directing and maintaining users' attention.

There are a variety of additional solutions to consider for directing attention across large distances including (1) animation, (2) in-place motion, (3) graphical trails, (4) sound and (5) attention sensors. Animation can be used to preserve visual coherence while a system state is changing. For example, if a tool palette is being dismissed and put into a system tray, an animation can be designed to de-emphasize the palette (e.g., make it progressively transparent), move it towards the tray and perhaps even shrink the palette to icon size. The user's attention is drawn to a location within the tray for later retrieval. Researchers have also explored how cartoon animation can be applied to user interface dynamics [12].

In-place animation attempts to exploit the user's attention sensitive peripheral capability by generating motion-cues. The use of motion may be useful in visualizing complex information because of its preattentive and interpretative perceptual properties [2]. For example, Microsoft's paperclip will dance in-place to get the user's attention to offer help or advice. Researchers of the multiple display PointRight system [26] noticed that users would "wiggle" the cursor (i.e., move it back and forth rapidly or in a circular path) to facilitate the user finding the cursor. The system detects this action and changes the "pointer" shape cursor to an animated icon (a rotating disk) to draw the user's attention to the cursor. Moreover, animated icons [1] provide additional information and system status for users.

Using graphical trails of visual feedback to hold the user's attention has been investigated by a number of researchers. The Drag+Pop technique [4] uses "rubber-band" lines to visually link the position of file icons that have been temporarily warped closer to the cursor for local manipulation. Several techniques have been developed to graphically show off-screen locations such as Halo [5] and citylights [43]. The high density cursor [7] shows that enhanced cursor tracking can be achieved through additional drawings of the cursor while being moved rapidly. Lastly, graphical streamers indicating the user's past navigational position has been used to predict and highlight future movement through a 3D space [13].

The use of sound can be an effective means of drawing user's attention that does not interfere with the visual system. For example, 2D or 3D audio localization can be used to direct a person's attention. In addition, a verbal description generated by a presenter (such as "look at the blue truck on the second panel") can identify a point of interest.

Actively sensing the attention state of a user has also been studied [42]. For example, EyePliances [34] are small device appliances that detect and respond to human visual attention using eye contact sensors. Applying this approach to our application, for example, could have the system attempt to sense where a particular user is standing or looking and visually guide them to a new focus area.

Combinations of these attention direction approaches can be used. For example, the common laser pointer serves as a derivative narrow spotlight. It offers a visual trail and wiggles in-place to direct the user's attention. It suffers from not being large enough to easily see and, due to the control-to-display ratio, is next to impossible to keep steady.

In the end, the spotlight has a rich history. There is a long tradition in the theater and cinema of utilizing lighting to draw the attention of the audience as well as to evoke emotional responses [9]. The spotlight graphical effect has been used for highlighting in television graphics and also more recently in GUIs, for example, the Mac OS X Tiger file browser.

Beyond cinematic experiences, the spotlight metaphor has been used within computer applications. For example, the Boom Chameleon [41] uses a spotlight technique to point to areas of interest on a 3D model. Computer lighting and shadow design has become more user manipulable [3] and some research has explored elevating shadows to first-class modeling primitives [32]. Lastly, the Light Pen system [27] explores a lighting design system driven by sketching on 3D virtual models. In our spotlight technique, we attempt to direct and maintain users' attention by utilizing similar theatric techniques of simultaneously defining a focus of attention, making surrounding context visible and suppressing distant content.

EXPERIMENT

The Spotlight technique has a number of attributes designed to help users track a target region or to re-acquire the target having, perhaps, looked away or been looking at another part of the large display when the intended target has moved. These attributes include (a) darkening the area outside the desired target area thereby (b) highlighting the area of interest while (c) ensuring the cursor is within the highlighted area. We hypothesize that, on wall-sized displays, users can find the cursor using the Spotlight technique much faster than finding the cursor unassisted. We chose to compare the spotlight technique against the status-quo standard cursor which is used everyday in our Visualization Studio.

We expected that another important factor in the performance of a user finding a target on a large wall-sized display is the field-of-view. When the target is within a user's field-of-view, the user can simply scan the screen by

finding the target in their peripheral vision or by moving their eyes. However, the further a target is from the center of a user's field-of-view, the more they will have to turn their heads to find the target. In extreme cases where the target may be slightly behind the user, they may have to turn their shoulders as well. As we are interested in the use of large wall-sized displays, we also wanted to capture this dimension of a target finding task.

Finally, we wanted to discover how the Spotlight technique might scale down to traditional desktop computer displays. We hypothesize that even on a typical monitor, our technique will still outperform the cursor.

Design

We created an experiment to discover these relative costs and recruited 12 subjects (8 male and 4 female) between the ages of 22 and 45 who were experienced computer users. Our set-up used the hardware configuration as shown in Figure 5. We placed the user in the center of this space to create three field-of-view conditions, so that 180° is the middle condition (see Figure 6).

We wrote an OpenGL application running under Chromium to display a contiguous workspace of 72' by 9' with a combined resolution of 7680 by 1024 pixels. Note that the cursor size matched the standard size used by the Windows operating system (see Figure 7 for relative spotlight and cursor sizing).

We ran the field-of-view conditions in narrow (108°), middle (180°), and wide (240°) order and three target indication techniques: traditional cursor, spotlight, and searchlight. We counterbalanced both field of view and technique.

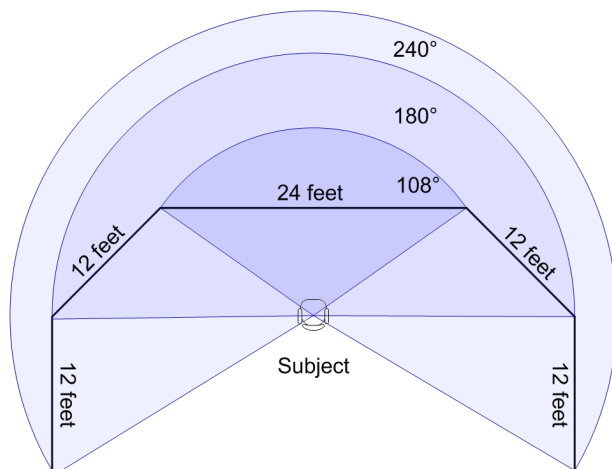


Figure 6. Experimental set-up. Three field-of-view angles were tested: 108°, 180°, and 240°.

On each trial, the user must find the cursor which has been positioned over a target number in a landscape of numbers

(Figure 8). Each trial starts with the cursor in a new location. The subject must find the target and press a single numeric key on a keyboard resting on the lap of the user to complete the trial. The next trial does not begin until the correct key is pressed. Because our experiment focuses on target acquisition, we did not animate the position of the spotlight to the new target position nor did we animate the opacity of the dark exterior region. Instead, once the proper key is hit, we immediately move the spotlight to the next position, to start the next trial.



Figure 7. Close-up of sample trial for Cursor, Spotlight, and Searchlight conditions (left to right). Note that the cursor is present in all techniques.

Subjects were instructed to use only one hand and only the top row of keys on the keyboard and not to use the “numpad” on the side of the keyboard. Also, the system chose target numbers that were a minimum of two keys away from the last key pressed to prevent subjects from simply pressing a neighboring key without looking down. These restrictions were added to force users to look away from the screen and down to the keyboard while selecting the key so they could not use peripheral vision to “see” the target move to the next location. Hence we were able to have each trial test target acquisition instead of target tracking. This design captures the real usage scenario of users often redirecting their attention (e.g., to take notes) and reacquire the focus of interest on the large display.

Before the start of the experiment, we briefly instructed the subjects on how the three target indication techniques worked and allowed them to try each condition to familiarize themselves with the task and the technique. Once comfortable, we asked them to complete the task as quickly and accurately as possible. Figure 8 shows a subject seated before the wall-sized display with the keyboard on his lap, illuminated with a desk lamp.

A within-subjects design was used with each subject using all three techniques in the three field-of-view (FOV) angles. This resulted in 12 subjects × 3 techniques × 3 FOV × 20 trials = 2160 data points for the entire experiment. Trials were grouped by technique and counterbalanced with 1/3 of the subjects using the cursor first, 1/3 using the spotlight first, and 1/3 using the searchlight first. Half of the subjects were presented with the field of view in the narrow-middle-wide order and the other half in the reverse order.

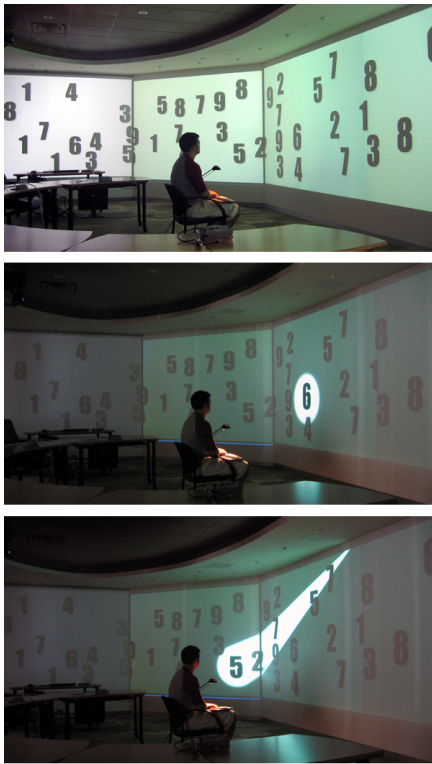


Figure 8. Experiment Set-up. Cursor, Spotlight, and Searchlight Conditions (Top to Bottom).

For every trial we logged the time for the subject to press the key on the keyboard corresponding to the number on the screen. Errors were not possible as the system waited for the correct key before proceeding to the next trial. After completing the experiment, subjects were given a short questionnaire to determine their preference for the three techniques and to provide other comments or observations they may have.

Results

Our data confirms the first hypothesis that users perform much faster at target acquisition with the spotlight technique compared to the cursor. We also found that the cursor performance cost increases as the field-of-view increases. The spotlight outperforms the cursor by a factor of 1.5 for the narrow and by 3.4 for the wide field of view. This confirms our second hypothesis (Figure 9). All subjects had a few cursor trials in which they were unable to find the cursor and had to rescan the full display multiple times resulting in a very long trial time. This is reflected in the high standard deviation for the cursor condition.

We performed an analysis of variance on the performance data and found a significant difference between the standard cursor and spotlight techniques $F(2, 22) = 455.55, p < .0001$. A significant difference was also found between the narrow-middle-wide field of view conditions $F(2, 22) = 119.47, p < .0001$. Moreover, a significant interaction exists between technique and field of view $F(4, 44) = 74.29, p < .0001$.

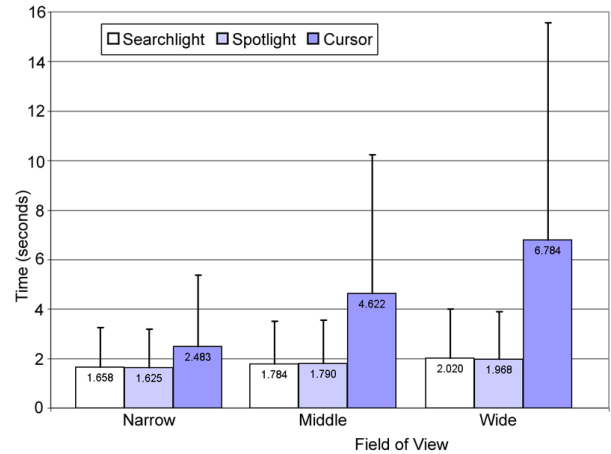


Figure 9. Mean performance on large display.

Experiment 2

After finding a significant benefit of the spotlight technique on a large wall sized display, even in the most narrow field of view, we wanted to determine if the effect would transfer to a desktop monitor configuration. We conducted a second experiment using a 21 inch monitor with the same three techniques. In this case, since the entire monitor surface is in the subject’s field of view, we filled the entire display with potential targets.

We recruited 12 subjects (8 male and 4 female) between the ages of 22 and 45 who were experienced computer users. Beyond removing the field-of-view factor, all experimental procedures and tasks remained the same as the last experiment. This resulted in 12 subjects × 3 techniques × 20 trials = 720 data points.

We performed an analysis of variance on the performance data and found a significant difference between the standard cursor and spotlight techniques $F(2, 22) = 92.7, p < .0001$. We found that the spotlight outperformed the cursor by a factor of 2.3. Figure 10 shows the mean performance differences for the three techniques.

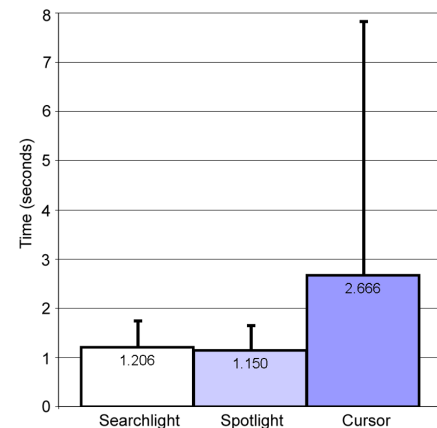


Figure 10. Mean performance on standard 21" monitor.

Subjective Preference

The results of the questionnaire showed that the subjective preference of the spotlight techniques matched subjects' performance. All subjects preferred the spotlight and searchlight techniques to the standard cursor for both the wall display and 21" monitor display (see Figure 11).

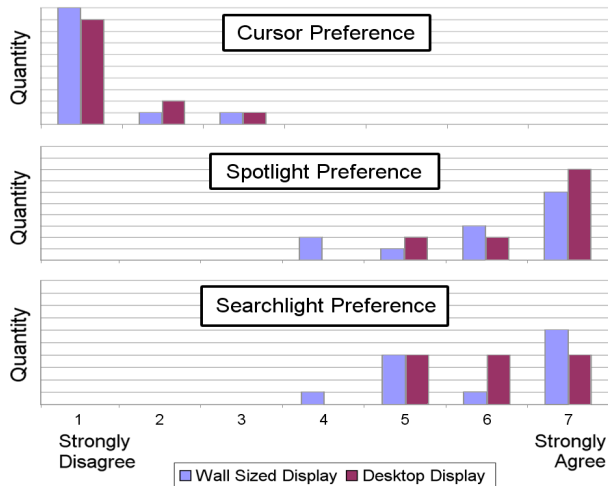


Figure 11. Subjective preference for both the wall-sized display and the desktop display when asked if they liked a given technique.

DESIGN VARIATIONS

We implemented a number of variations of the spotlight technique. Figure 12 (b) shows the simple circular spotlight that we tested in our experiment. We feel that this design was slightly more effective than the searchlight because there are no other distracting elements on screen. The searchlight (see Figure 12 (e)) has an additional beam to direct the user's attention, but we believe that this can slow down the target acquisition in that, once the user looks within the beam, the eye is drawn linearly all the way down to the spotlight region. This effect seems to make it difficult to skip directly down to the spotlight. A problem with the simple circular spotlight is the ongoing need to resize it. In practice, we typically view objects in a variety of sizes. And due to the overall aspect ratio of our display (72" by 9") we often try to layout objects to use the full height of the screen, as we have plenty of width to use. This led us to build the vertical bar spotlight shown in Figure 12 (c). We call this spotlight the *curtain* because when resizing it, it opens and closes in width, much like a curtain. We also made a version of the curtain that increased the cone angle of the light, that is, the curtain width at the bottom of the screen only (see Figure 12 (d)).

To reduce the need to resize the spotlight, we also implemented a shape-based spotlight, shown in Figure 15 (f), which uses the alpha channel of the object to create the spotlight border. While this was effective in always highlighting an entire object, it comes at the cost of not

being able to point the spotlight to an area within an object. We also found that highlighting the shape is slightly distracting as well and thought that a spotlight in the shape of the bounding ellipse of the object would be more appealing (see Figure 12 (i)). Finally, we implemented multiple spotlights (see Figure 12 (a), (g) and (h)) and, in future, would like to explore the use of spotlights to represent the interest of multiple users.

Lastly, we have considered but not yet implemented the idea of using alternatives to the brighten/darken effect in the spotlight/exterior region. For example, one could adjust the content saturation, tint, pixelation or blur in the exterior region, or use other visual effects to differentiate the regions.




Figure 12. Spotlight Designs (a) Multi-light. (b) Circular. (c) Curtain. (d) Cone-light. (e) Searchlight. (f) Shape-light. (g) Multiple Searchlights. (h) Multiple Cone-lights. (i) Elliptical Shape-searchlight.

CONCLUSION

In this paper we introduced a new technique for fast target acquisition, called the Spotlight, designed to draw the users attention to the intended area of a large wall-sized display. We showed the value of augmenting the standard cursor with the spotlight by running an experiment and we reported a significant performance improvement in a target acquisition task. Furthermore, we ran a second experiment showing that the same performance advantages over the cursor exist using the spotlight on a standard desktop monitor.

In the future, we wish to extend the spotlight to support multiple users pointing at the large display, with a variety of input devices such as mouse, laser pointer, or by physically pointing. Also, we plan to continue the development of a content-aware spotlight that will size itself to the relevant content being highlighted. We also believe that more cinematic qualities of lighting and stage design may be applied to large display configurations to assist in managing and maintaining audiences attention as well as to make for more enjoyable experiences.

It remains an open question whether other techniques, different from the Spotlight, such as changing cursor color, size, shape, animation (as in PointRight [26]), etc. can be

used to manage attention. The design, development and evaluation of these techniques are future research.

ACKNOWLEDGMENTS

The authors thank Don Almeida, Ricky Uy at the University of Virginia, Graham Binks, Yit Lam, Xavier Melkonian and Hewlett-Packard, Barco Systems, Steelcase, Viewsonic, ATI, and IBM.

REFERENCES

- Baecker, R., Small, I. and Mander, R. (1991) Bringing icons to life, In *Proceedings of ACM CHI'91*, pp. 1-6.
- Bartram, L. (1997) Perceptual and interpretative properties of motion for information visualization, In *Proceedings 1997 workshop on New paradigms in information visualization and manipulation*, pp.3-7, November 10-14, 1997, Las Vegas, Nevada.
- Barzel R. (1997) Lighting Controls for Computer Cinematography, *Journal of Graphics Tools*, 2(1), pp. 1-20.
- Baudisch, P., Cutrell, E., Robbins, D., Czerwinski, M., Tandler, P., Bederson, B., and Zierlinger, A. (2003) Drag-and-Pop and Drag-and-Pick: Techniques for Accessing Remote Screen Content on Touch- and Pen-operated Systems. *Proceedings of Interact 2003*, pp. 57-64.
- Baudisch, P. and Rosenholtz, R. (2003) Halo: a Technique for Visualizing Off-Screen Locations, In *Proceedings of ACM CHI 2003*, pp. 481-488.
- Baudisch, P., Good, N., and Stewart, P. (2001) Focus Plus Context Screens: Combining Display Technology with Visualization Techniques, In *Proceedings of ACM UIST '01*, pp.31-40.
- Baudisch, P., Cutrell, E. and Robertson, G. (2003) High-density cursor: a visualization technique that helps users keep track of fast-moving mouse cursor, *Proceedings of Interact 2003*, pp. 236-243.
- Baudisch, P. Cutrell, E., Hinckley, K. and Gruen, R. (2004) Mouse Ether: Accelerating the Acquisition of Targets Across Multi-Monitor Displays, In *Proceedings of CHI 2004*, pp. 1379-1382.
- Box, H.C. (1997) *Set Lighting Technician's Handbook. Film Lighting Equipment, Practice, and Electrical Distribution*, Second Ed., Focal Press, Boston.
- Buxton, W., Fitzmaurice, G.W., Balakrishnan, R. and Kurtenbach, G. (2000) Large displays in automotive design, *IEEE Computer Graphics and Applications*, July/August 2000, pp. 68-75.
- Cadiz, J.J., Venolia, G., Jancke, G. and Gupta, A. (2002) Designing and Deploying an Information Awareness Interface, In *Proceedings of CSCW 2002*, pp. 314-323.
- Chang, B-W and Ungar, D. (1993) Animation: from cartoons to the user interface, In *Proceedings of ACM UIST 1993*, pp 45-55.
- Chapman, D. and Ware, C. (1992) Manipulating the Future: Predictor Based Feedback for Velocity Control in Virtual Environment Navigation. Special Issue of Computer Graphics. *Proceedings of ACM 1992 Symposium on Interactive 3D Graphics*, Cambridge Mass, 63-66.
- Czerwinski, M., Tan, D.S. and Robertson, G.G. (2002) Women Take a Wider View, *Proceedings of CHI 2002*, pp. 195-202.
- Elrod, S., Bruce, R., Gold, R., Goldberg, D., Halasz, F., Janssen, W., Lee, D., McCall, K., Pedersen, E., Pier, K., Tang, J., and Welch, B. (1992) Liveboard: a large interactive display supporting group meetings, presentations, and remote collaboration, *Proceedings of CHI 1992*, pp. 599-607.
- Fass, A., Forlizzi, J. and Pausch, R. (2002). MesyDesk and MessyBoard: two designs inspired by the goal of improving human memory. In *Proceedings of ACM DIS 2002*, pp. 303-311.
- Fitzmaurice, G., Khan, A., Pieke, R., Buxton, B. and Kurtenbach, G. (2003) Tracking Menus, In *Proceedings of ACM UIST 2003*, pp. 71-79.
- Geißler, J. (1998) Shuffle, Throw or Take It! Working Efficiently with an Interactive Wall. In *ACM Conference Companion to CHI 1998*, pp. 265-266.
- Grudin, J. (2001) Partitioning Digital Worlds: Focal and Peripheral Awareness in Multiple Monitor Use, *Proc. CHI 2001*, pp. 458-465.
- Guimbretière, F., Stone, M. & Winograd, T. (2001) Fluid Interaction with High-resolution Wall-size Displays, *Proceedings of ACM UIST 2001*, pp. 21-30.
- Hascoët, M. (2003) Throwing models for large displays, *Proceedings of HCI'2003, Designing for Society*, Volume 2, pp. 73-77, British HCI Group.
- Hinckley, K., Ramos, G., Guimbretiere, F., Baudisch, P. and Smith, M. (2004) Stitching: Pen Gestures that Span Multiple Displays, In *Proceedings of AVI'04*, pp. 23-31.
- Horvitz, E., Kadie, C., Paek, T. and Hovel, D. (2003) Model of Attention in Computing and Communication: From principles to applications, *Communications of the ACM*, 46 (3), pp. 52-59.
- Humphreys, G., Eldridge, M., Buck, I., Stoll, G., Everett, M. and Hanrahan, P. (2001) WireGL: a scalable graphics system for clusters, *Proceedings of SIGGRAPH 2001*, pp. 129-140.
- Hutchings, D.R., Smith, G. and Meyers, B. (2004) Display space usage and window management operation comparisons between single monitor and multiple monitor users. In *Proceedings of AVI 2004*, pp. 32-39.
- Johanson, B., Hutchins, G., Winograd, T., and Stone, M. (2002) PointRight: Experience with Flexible Input Redirection in Interactive Workspaces. *Proceedings of UIST 2002*, pp. 227-234.
- Jung, T., Gross, M.D. and Do, E.Y.-L. (2003) Light Pen – Sketching light in 3D, *Proceedings of 10th International Conference on Computer Aided Architectural Design Futures*, Tainan, Taiwan.
- MacIntyre, B., Mynatt, E.D., Volda, S., Hansen, K.M., Tullio, J., Corso and G.M. (2001) Support for Multitasking and Background Awareness Using Interactive Peripheral Displays, *UIST'01*, 41-50.
- Mackinlay, J.D. and Heer, J. (2004) Wideband Displays: Mitigating Multiple Monitor Seams, In *Proceedings of CHI 2004*, pp. 1521-1524.
- Maglio, P.P., Campbell, C. (2000) Tradeoffs in Displaying Peripheral Information, *Proceedings of CHI 2000*, pp. 241-248.
- Pashler, H.E. (1998) *The Psychology of Attention*, "A Bradford book" Publisher.
- Pellacini, F., Tole, P. and Greenberg, D.P. (2002) A User Interface for Interactive Cinematic Shadow Design, In *ACM Transactions on Graphics (TOG)*, 21 (3), July 2002 pp. 563-566.
- Rekimoto, J. and Saitoh, M. (1999) Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments, In *Proceedings of CHI 1999*, pp. 378-385.
- Shell, J.S., Vertegaal, R. and Skaburskis, A.W. (2003) EyePliances: Attention-Seeking Devices that Respond to Visual Attention, In *Proceedings of CHI 2003*, pp. 770-771.
- Streitz, N.A., Geißler, J., Holmer, T., Konomi, S., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., Steinmetz, R. (1999) i-LAND: an interactive landscape for creativity and innovation, *Proceedings of ACM CHI 1999*, pp.120-127.
- Swaminathan, K., and Sato, S. (1997) Interaction Design for Large Displays. *Interactions* 4(1), pp. 15-24.
- Tan, D.S. and Czerwinski, M. (2003) Effects of Visual Separation and Physical Discontinuities when Distributing Information across Multiple Displays. In *Proceedings of Interact 2003*, pp. 252-255.
- Tan, D.S., Czerwinski, M. and Robertson, G. (2003) Women Go With the (Optical) Flow, *Proceedings of CHI 2003*, pp. 209-215.
- Tan, D.S., Gergle, D., Scupelli, P.G., and Pausch, R. (2003) With Similar Visual Angles, Larger Displays Improve Spatial Performance, *Proceedings of CHI 2003*, 217-224.
- Tani, M., Horita, M., Yamaashi, K., and Tanikoshi, K. (1994) Courtyard: Integrating Shared Overview on a Large Screen and Per-user Detail on Individual Screens, *Proceedings of CHI'94*, pp. 44-50.
- Tsang, M., Fitzmaurice, G., Kurtenbach, G., Khan, A., Buxton, B. (2002). Boom Chameleon: Simultaneous capture of 3D viewpoint, voice and gesture annotations on a spatially-aware display. *ACM CHI Letters*, 4(2). p. 111-120. ACM UIST 2002.
- Vogel, D. Balakrishnan, R.. (in press, 2004). Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. To appear in *UIST 2004*.
- Zellweger, P.T., Mackinlay, J.D., Good, L., Stefik, M., Baudisch, P. (2003) City Lights: Contextual views in minimal space, In *Proceedings of CHI 2003*, pp. 838-839.