

Shadow Reaching: A New Perspective on Interaction for Large Wall Displays

Garth Shoemaker, Anthony Tang, and Kellogg S. Booth
University of British Columbia
2366 Main Mall
Vancouver, BC, V6T 1Z4, Canada
garths@cs.ubc.ca, tonyt@ece.ubc.ca, ksbooth@cs.ubc.ca

ABSTRACT

We introduce Shadow Reaching, an interaction technique that makes use of a perspective projection applied to a shadow representation of a user. The technique was designed to facilitate manipulation over large distances and enhance understanding in collaborative settings. We describe three prototype implementations that illustrate the technique, examining the advantages of using shadows as an interaction metaphor to support single users and groups of collaborating users. Using these prototypes as a design probe, we discuss how the three components of the technique (sensing, modeling, and rendering) can be accomplished with real (physical) or computed (virtual) shadows, and the benefits and drawbacks of each approach.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General terms: Human Factors, Design

Keywords: Large displays, interaction techniques

INTRODUCTION

An enduring theme in research on very large wall displays is supporting effective input for users. Two main problems exist: providing fluid access to all areas of the large display for a single user, and conveying awareness of interactions to collaborators. We introduce a novel technique called Shadow Reaching that addresses both of these problems using the perspective information implicit in cast shadows.

Shadow Reaching relies on the underlying interaction metaphor of physical shadows. Unlike previous work employing shadows for interaction [1, 4, 8, 14], central to our technique is a perspective-based transformation of the shadows. Figure 1 illustrates how the size of a user's shadow, and hence the user's effective reach, varies based on the user's movements relative to the light source and the display. The result is fluid, seamless interaction and control over the entire display. Furthermore, because a user's interactions are embodied in the display as familiar shadows, both the user and collaborators can easily understand and interpret interactions.

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.

UIST' 07, October 7-10, 2007, Newport, Rhode Island, USA.

Copyright 2007 ACM 978-1-59593-679-2/07/0010...\$5.00.




Figure 1 A user interacts (left) in close proximity with a small region of the display, and (right) steps back to interact over larger distances.

The primary contribution of this Tech Note is an interaction technique that balances the design tension between the need for efficient reaching (pointing) on very large displays by a single user, and the need for easily interpretable, embodied actions that can be understood by co-present collaborators. We first describe previous solutions to each problem, and then present a high-level description of the Shadow Reaching technique and three prototype implementations that illustrate how the technique is built up from sensing, modeling, and rendering components.

Our first prototype literally uses physical shadows augmented by 6DOF sensing to maintain the model of a user's reachable area. The second prototype applies a vision-based approach to the basic shadow metaphor to facilitate interaction using the entire shadow, and the third extends the second to Magic Shadows, an embodied interaction technique reminiscent of GUI-based Magic Lenses.

We conclude with a discussion of the design space surrounding Shadow Reaching, and also discuss future work, including how the technique could be extended to enforce access control, and how dynamically positioned virtual light sources could enhance the technique.

ISSUES WITH VERY LARGE DISPLAY INTERACTION

There are two design factors that must be addressed by interaction designers for large display: support for interaction over large distances on the display, and support for interactions that are easily interpretable. In this section, we set the stage for Shadow Reaching by illustrating how these two factors are often at odds with each other.

Interaction at a Distance

Large displays can present copious amounts of information spread across the entire work surface. Direct interaction techniques such as touch-sensitive surfaces make it difficult to interact with information in areas not immediately within reach: a user must physically move about the workspace, and may not be able to reach some content at all on large displays (e.g. a 10' tall wall or 6' wide tabletop display).

Indirect interaction techniques fare better in supporting interaction at a distance. Reetz et al. [10] provide a taxonomy of indirect techniques including direct action, cursor extension, long-distance pointing, proxy techniques, radar techniques, and throwing. Each of these techniques effectively supports interaction at a distance for a single user working on a large display. However, as we will see, they all fall short in terms of making interactions interpretable by the user and the user's collaborators.

Interpretable Interactions

A common design heuristic for groupware is to provide consequential communication of activity in the form of continuous feedback of interaction so others can easily interpret and understand those interactions [5]. For large displays, many authors have argued for direct input techniques because they provide physically-grounded embodiments of interaction [12]. For example, Pick-and-Drop provides an interpretable means for transferring information by placing the onus on users to physically move to different areas of the screen to "pick" up and "drop" off information [11]. Similarly, Wu and Balakrishnan [15] describe direct interaction techniques for manipulating data on tabletop displays.

Common to these approaches is the user's physical embodiment within a physical space. When interaction techniques are based on physical properties of a user's embodiment, they become understandable because of our everyday experiences with our physical selves. In the context of large display interaction, we have come to think of *embodiment* as the extent to which a user has a direct connection to the interaction being performed, and ultimately to the data being acted upon. When this connection is strong, interactions are *interpretable*.

Many techniques designed for distance interaction fail to achieve the goal of being interpretable. TractorBeam [9] provides a powerful means of pointing and interacting with remote information, but does not provide a clear link between the user's physical self and the data being acted upon. Proxy techniques such as Frisbee [7] or radar techniques such as Push-and-Pop [3] break down because the original data is isolated at a distance from the user and from the interaction. Throwing techniques [6] do not provide direct links to the initiating user, making it difficult to determine the originator of an action and thus the meaning of the interaction. As a result, these techniques impede collaborators, or even the user, from understanding and predicting the result of actions.



Figure 2 The user can control the reach of her shadow by moving closer to and farther away from the display.

SHADOW REACHING

Shadow Reaching employs a shadow on the display surface through which the user interacts with the scene. A perspective projection applied to the shadow, controlled directly by the user through body positioning, allows the user to increase or decrease the effective range of interaction in a fluid, seamless manner (Figure 2). The tradeoff associated with adjusting range of interaction is a change in interaction resolution (accuracy). The combination of the shadow as interaction proxy, and the ability to control range of interaction, satisfies the dual requirements of interaction at a distance and interaction embodiment, which have thus far been elusive to many interaction designers.

Our shadows-based interaction technique draws on Krueger et al.'s work with VIDEOPLACE, and extends it by making a perspective projection based on the position of the light source a core component of the technique. It is this new perspective component that provides users with the ability to control their range of effective interaction.

Informal observations of six first time users of Shadow Reaching produced encouraging results. Users were first introduced to a puzzle-building application employing shadow input as a cursor replacement. They were shown how to "click" using a handheld device which doubled as a position sensor for determining hand location, and then were left to explore. In general, users immediately understood how to interact, including the significance of the change in shadow size as they moved relative to the display. This is likely a result of experience with shadows in everyday life. In particular, we observed that users naturally stepped back from the display when they wanted to understand the "bigger picture." The related increase in reach maps nicely to a broadened scope of interest.

PROTOTYPE IMPLEMENTATIONS

We implemented three different applications to explore Shadow Reaching. The first is implemented as a general replacement for a mouse cursor for pointing and interacting with the workspace. The second employs a full-body interaction metaphor as did Krueger et al., and the third extends Magic Lens techniques [2] to achieve Magic Shadows. The prototypes use different approaches to sensing, modeling, and rendering. These are discussed in turn.



Figure 3 Using the Shadow Reaching prototype.

Single Point Interaction

The first prototype uses a real-world shadow, generated from a powerful lamp 10 feet from the screen, to support single-point input, as with a mouse (see Figure 3). For easier sensing, the user holds a Polhemus position tracker and a Phidgets button. The modeling stage uses the known geometry of the light source and display and the sensed location of the tracker to determine the location of the shadow of the user's hand on the screen. Button presses trigger click events at that location. The real shadow of the user is a physical embodiment in the workspace, but is not used computationally.

The prototype supports a puzzle-building task, with multiple input devices available for simultaneous bimanual or collaborative interaction. A more practical implementation would use vision-based sensing, either of the 2-D shadow on the screen, or of the 3-D user. The Polhemus was used to guarantee accuracy so we could evaluate the interaction technique without worrying about a vision subsystem.

The real-world light source used for shadow rendering was placed at roughly shoulder height of the user. Critical factors when choosing a light are intensity and beam angle.

Whole Body Interaction

An alternate model of interaction makes use of the entire shadow for whole-body input. Unlike single-point “click” interaction, there is a broad range of interactive possibilities, including use of hand and arm gestures to pick up and manipulate scene objects, use of head positioning for view control, and use of the legs for secondary interaction.

We developed a demonstration application using full body interaction via shadows with dynamic on-screen content. In this second prototype shadow sensing was accomplished using a light source behind the screen, captured with an infrared camera in front of the screen, and extracted by rudimentary computer vision techniques similar to what Tan and Pausch [13] used in a different context. A model of the user's location in space was then computed, and the shadows were rendered onto the screen. In contrast with the first prototype, where everything was “real” except for the computation of hand position, the second prototype's modeling and rendering were accomplished entirely in the vir-



Figure 4 A mockup of how virtual shadows can be used as representation-altering Magic Shadows. In this case the shadows contain satellite photo data, while the surrounding regions hold conventional map data.

tual domain. This illustrates the de-coupling of the sensing, modeling, and rendering components of Shadow Reaching.

In this application, the user's embodied shadow interacts with virtual balls bouncing around the large display. The modeling component constrains balls to bounce off the shadow, and to otherwise follow physical laws. While the application was designed without any intended user goal, we found that users spontaneously developed their own tasks based on the possibilities presented by the system. One user decided to trap balls in outstretched and joined arms, while another attempted to keep balls from hitting the ground. From this we conclude that whole body interactions present a host of affordances that can be exploited.

Shadows as Magic Lenses

Shadow embodiments are very personal. Like Krueger et al., we have found that users generally do not intrude on others' shadows. As a direct result of this, it could be helpful to use shadows to personalize the display of on-screen data in collaborative scenarios.

Our third prototype used shadows to define the boundaries of a Magic Lens (see Figure 4). Magic Lenses are movable see-through widgets that are used to visually filter on-screen data. They can perform arbitrary transformations on the data, including altering representation or presentation of secondary information. Magic Shadows provide a natural means of defining personal views of data, and moving a lens about the workspace. As in the second prototype, we used a vision-based method for generating virtual shadows.

EXPLORING THE DESIGN SPACE

When designing shadow-based interactions, certain choices must be made. The projection used for shadow generation is important. We used a perspective projection with the goal of enabling distance reaching, but Krueger et al.'s orthographic projection may be more appropriate for detailed interaction when a user is standing at a distance from the display. The method for rendering shadows is also an important consideration. Using real shadows is easy and powerful, but provides no support for modeling and customiza-

tion such as color-coding to distinguish users or interaction modes. The effectiveness of real shadows also depends on lighting conditions (bright rooms cause washed-out shadows). Vision-based sensing, on the other hand, opens up many possibilities for processing the captured data and modifying it before it is ultimately rendered.

The use of real or virtual shadows as interface elements raises a number of possibilities for user interactions. We discuss two of these here.

Access Control

A shadow is a very personal embodiment, and as such it may be useful as a means of controlling access to data in a collaborative setting. When users are gathered around a wall display, it is important that work be coordinated so that users avoid interfering with one another. There are natural tendencies that help govern this coordination, such as users avoiding making contact with other users' shadows. This could be made explicit by restricting a user's ability to edit data within a collaborator's shadow (perhaps extended by penumbra), or even inhibiting a user's virtual shadow from intruding on another user's virtual shadow.

Dynamic Light Source Positioning

Our prototypes made use of stationary real and virtual light sources. Assuming that a light source is centered on the display, the perspective distortion increases as a user approaches the edges of the display, potentially rendering interaction difficult. It would be desirable to explore more complex rules governing the movement of virtual light sources. For example, a virtual light source could move relative to a user's position in the physical world, or relative to the orientation of their body. Furthermore, independent light sources could exist for each user. Lastly, distortions other than simple perspective distortions could be employed that would ease distance interaction and reduce the need for physical movement.

CONCLUSIONS AND FUTURE WORK

The promise of ubiquitous large computer displays, including wall mounted and tabletop units, has led to significant research activity attempting to define how users will interact with those displays. A number of different interaction techniques have been proposed, but we believe that none of them adequately satisfy the twin factors of embodied interactions and interaction at a distance. We propose a new interaction technique, dubbed Shadow Reaching, which employs a perspective-projected shadow of the user on the display for interaction. The shadow maintains physical embodiment, and perspective allows interaction at a distance.

Our prototypes have revealed many positive aspects of Shadow Reaching. There is still, however, much interesting work to be done. First, we plan to make a quantitative comparison of pointing efficiency using Shadow Reaching and other large screen interaction techniques. Second, we will investigate the degree to which shadow embodiment aids the user and collaborators in understanding the interactions being performed. Third, we need to further explore the possibilities identified earlier, including shadow-based access control, and dynamic positioning of virtual light

sources. Lastly, we will investigate ways for users to transition from direct touch interactions to Shadow Reaching.

REFERENCES

1. Apperley, M., McLeod, L., Masoodian, M., Paine, L., Phillips, M., Rogers, B., and Thomsen, K. (2003). Use of video shadow for small group interaction awareness on a large interactive display surface. *Proc. AUIC 2003*, 81-90.
2. Bier, E.A., Stone, M., Pier, K., Buxton, W., and DeRose, T. (1993). Toolglass and Magic Lenses: the see-through interface. *Proc. SIGGRAPH 1993*, 73-80.
3. Collomb, M., Hascoet, M., Baudisch, P., and Lee, B. (2005). Improving Drag-and-Drop on wall-size displays. *Proc. GI 2005*, 25-32.
4. Davis, J. and Bobick, A. (1998). Virtual PAT: a virtual personal aerobics trainer. *Workshop on Perceptual User Interfaces*, 13-18.
5. Gutwin, C. and Greenberg, S. (2002). A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work*, 11(3):411-446. Kluwer Academic Press.
6. Hascoet, M. (2003). Throwing models for large displays, *HCI 2003*, Volume 2, 73-77.
7. Khan, A., Fitzmaurice, G., Almeida, D., Burtnyk, N., and Kurtenbach, G. (2004). A remote control interface for large displays. *Proc. UIST 2004*, 127-136.
8. Krueger, M.W., Gionfriddo, T., and Hinrichsen, K. (1985). VIDEOPLACE – an artificial reality. *Proc. CHI 1985*, 35-40.
9. Parker, J.K., Mandryk, R., Nunes, M., and Inkpen, K. (2005). TractorBeam selection aids: improving target acquisition for pointing input on tabletop displays. *Proc. INTERACT 2005*, 80-93.
10. Reetz, A., Gutwin, C., Stach, T., Nacenta, M., and Subramanian, S. (2006). Superflick: a natural and efficient technique for long-distance object placement on digital tables. *Proc. Graphics Interface 2006*, 163-170.
11. Rekimoto, J. (1997). Pick-and-Drop: A direct manipulation technique for multiple computer environments. *Proc. UIST 1997*, 31-39.
12. Scott, S.D., Grant, K.D., and Mandryk, R.L. (2003). System guidelines for co-located, collaborative work on a tabletop display. *Proc. ECSCW 2003*, 159-178.
13. Tan, D. and Pausch, R. (2002). Pre-emptive shadows: eliminating the blinding light from projectors. *Extended Proc. CHI 2002*, 682-683.
14. Utterback, C. and Achituv, R. (2002). Text Rain. *SIGGRAPH Electronic Art and Animation Catalog*, 78.
15. Wu, M., and Balakrishnan, R. (2003). Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. *Proc. UIST 2003*, 193-202.