Interacting with Large Displays

Patrick Baudisch
Microsoft Research

Wall-size screens and tabletop displays require new user interfaces.

While desktop PC screens used to be the norm, the range of available display form factors has exploded in recent years. In particular, users can now create large personal displays by connecting multiple screens to their PCs, using a projector, or forming interactive digital walls out of multiple projectors.

However, these new capabilities have raised research challenges with respect to interface design, as the straightforward approach of transitioning desktop interfaces to wall-size screens and tabletop displays leads to many serious technical problems. To address these challenges, Microsoft Research is exploring several new interaction technologies.

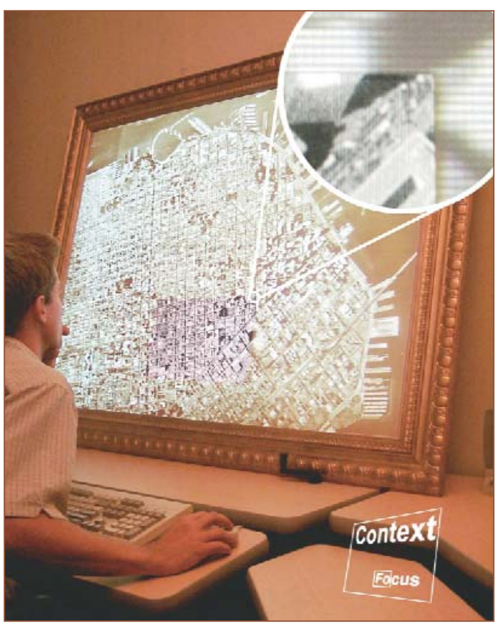

Figure 1. Focus-plus-context screen prototype. Content panned into the focus region is viewed in higher detail, making the focus display behave like a magic high-res lens.

In our lab we have worked with a broad range of large-screen devices including projector-based displays, multiple-smartboard walls, an 18-panel LCD, and interactive tabletops. Figure 1 shows a *focus-plus-context screen*, a particularly inexpensive way to obtain a large display with high resolution. This type of screen basically consists of a small high-resolution display embedded into a large low-res display.

The shown prototype seamlessly integrates an LCD into a front-projection screen. Customized software displays graphical content across both display regions, thereby preserving image scaling. Resolution varies across the two display regions: Content panned into the focus region is viewed in higher detail, making the focus display behave like a magic high-res lens.

USABILITY ISSUES

To explore usability issues related to wall-size and tabletop displays, we developed a series of prototype applications for image viewing and video conferencing, as well as a simple driving simulation, and performed numerous in-house studies.

We found that large displays offer a broad range of benefits, from increasing users' task-management performance to improving their spatial abilities. Alongside these benefits, however, several problems became apparent to us.

Many key desktop interaction paradigms "broke" when we tried to apply them to the focus-plus-context screen. For example, the sheer size of the display caused users to lose track of the mouse pointer; the display's large size also made it difficult to reach distant content using touch or pen input. In addition, those techniques that did seem to transfer often suffered from limited accuracy due to the large screen's inferior tracking capability.

IMPROVING MOUSE TARGETING

On large displays, users employ mouse accelerations to traverse the screen reasonably quickly. The faster the cursor moves, however, the more it seems to jump from one position to the next because it updates only at the monitor's refresh rate.

To deal with this problem, we developed the *high-density cursor,* which helps users track a cursor's movement by filling in additional cursor images between actual cursor positions—a process called *temporal supersampling*. Unlike existing technologies, such as the Windows *mouse trail*, the high-density cursor preserves responsiveness. In one user study we conducted, it significantly improved participants' performance on a Fitts' law target-acquisition task.

Our follow-up project, *mouse ether*, simplifies mouse targeting across multiple monitors by compensating for mouse path distortion caused by bezels, gaps, and resolution differences. We found mouse ether to improve users' targeting performance by up to 28 percent.

REACHING DISTANT TARGETS

Our research indicates that large-screen users have even more trouble with touch and pen input than with mouse input.

Drag-and-pop is a technology we developed that lets users access content on wall-size displays that would otherwise be hard or impossible to reach.

In the example shown in Figure 2, the user is moving an icon located in the right screen into a folder in the center screen. As he drags the icon left toward the folder, all potential target folders temporarily "stretch" toward the current pointer location much like rubber bands. The user can then file the icon in the appropriate folder using a comparably small hand movement.

In one user study we conducted, participants filed icons up to 3.7 times faster using the drag-and-pop interface than using drag-and-drop. A follow-up study revealed that drag-and-pop also outperforms the more traditional approach of extending the user's reach.

We're currently extending this interaction paradigm to work with arbitrary screen content. *Tablecloth* lets users "pull" a distant screen area toward them with their nondominant hand much like a person without table manners might try to get a distant salt shaker by yanking the tablecloth.

PRECISE MANIPULATION

On regular PCs, users can hold down keyboard keys to modify the mouse's functionality. On wall-size screens, however, keyboards generally aren't within easy reach. Consequently, techniques that rely on keyboard input don't work properly on large screens.

Alignment is one example. In our touch-screen version of *snap-and-go*, invisible guides help users drag objects into aligned positions. Snap-and-go is roughly as fast as traditional "magnetic" snapping. The latter technique, however, requires users to hold down a qualifier key whenever placing an object in the immediate proximity of a snap location to prevent it from getting warped there. Snap-and-go isn't subject to this limitation and is thus more practical on wall-size displays.

Figure 2. Drag-and-pop technique. As the user drags an icon from the right panel to the center, potential target folders "stretch" toward him, enabling him to file the icon with a small hand movement.

While large screens often have limited capabilities compared to regular-size screens, some of them are more powerful in other ways. We are currently working with a computer-vision-based tabletop display, for example, that is able to track multiple fingers simultaneously.

Our *dual-finger selection techniques* exploit this extra functionality to help users select very small targets. In the example shown in Figure 3, the user manipulates the pointer with the right hand while adjusting the control-display ratio with the left—in this case, stopping the pointer altogether by selecting "freeze."

Figure 3. Dual-finger selection on multitouch tabletop display. The user manipulates the pointer with the right hand while adjusting the control-display ratio with the left.

The shown *x-menu* technique is one of three we created on our tabletop display. In a user study we conducted, all three techniques achieved significantly faster and more accurate results than a control condition across various target sizes.

Today's computing systems often rely on display technologies originally designed for the desktop, but wall-size screens and large tabletop displays have qualitatively different requirements. Our work thus far has focused on basic input devices and interaction techniques. However, to enable a user experience explicitly based on large screens, future research must address the layers on top of this: new ways of operating applications and, hopefully, entirely new types of applications.

Patrick Baudisch is a research scientist at Microsoft Research in Redmond, Washington. Contact him at www.patrickbaudisch.com.

Editor: Bill Schilit, Intel Research Seattle; bill.schilit@intel.com