Human-Centered Interfaces for Large, High-Resolution Visualization Systems

David Vanoni
Department of Computer Science and Engineering
University of California, San Diego
dvanoni@cs.ucsd.edu

Abstract

To address the challenge of visualizing and working with the massive amounts of digital data produced today, large, high-resolution visualization systems have been developed that present information to users at scales beyond what is possible with traditional computing systems. While these systems provide an excellent resource for handling substantial visual datasets, they also present an interesting challenge when considering effective methods for user interaction. Because standard desktop interfaces do not inherently scale up for practical use in these large environments, herein we look at a human-centric approach for designing interfaces and present a survey of interaction techniques that have been developed for leveraging the large physical and visual space afforded by these systems. We observe how these various approaches account for the needs and facilities of the human within large-scale interaction spaces and consider possibilities for new interfaces based on emerging technology trends.

1. Introduction

With the ever-growing amount of digital data produced every day, the question arises of how to effectively visualize, consume, and analyze these data. Large, high-resolution visualization systems offer a unique platform for tackling this problem. These systems typically consist of an array of digital projectors or monitors combined into an immense display space aimed at presenting visual information to users at scales not possible with traditional computing systems.

The substantial space provided by large-scale displays offers many advantages when dealing with massive amounts of visual data but also raises interesting new challenges when considering how to effectively interact with these "human scale" systems [3]. Typical desktop computing interfaces employ a "windows, icons, menus, pointer" (WIMP) design accompanied by the time-honored keyboard and mouse, but unfortunately this tried-and-true interaction method fails to scale up and leverage the available space. Effective interfaces for these systems cannot simply evolve from scaling up existing techniques, but rather they must harness the natural tendency of humans, as spatially located creatures, to explore and operate within space [35].

Interaction design for large, high-resolution display systems must not only leverage the affordances of the technology but also consider the needs and facilities of the human. This design strategy, known as human-centered design, is defined by the International Organization for Standardization (ISO) as an "approach to systems design and development that aims to make interactive systems more usable by focusing on the use of the system and applying human factors/ergonomics and usability knowledge and techniques" [28]. Applying a human-centric perspective to the design of interfaces for large-scale visualization systems provides guidance for developing usable, effective interaction techniques.

This paper explores various approaches that have been developed for interacting with large display systems, focusing on human usability. We will begin with an overview of large, high-resolution visualization systems, identifying both the usability benefits and interaction challenges that have been observed when using these systems. We will then discuss two specific interaction themes that arise in discussions of interface design for these systems: embodied interaction and ubiquitous computing. This will be followed by an in-depth investigation of multiple categories of interaction techniques focusing on how they rely on embodied interaction and ubiquitous computing to provide effective interfaces. Finally, we consider opportunities for future research based on emerging technology trends.

2. Large, High-Resolution Visualization Systems

Large-scale visualization systems have been in active development for over two decades. One of the first such systems to gain popularity was the CAVE Automatic Virtual Environment (CAVE) [15]. The original CAVE system consisted of a cube-shaped room made up of 3-meter by 3-meter rear-projection screen walls. While this system provided an impressive immersive experience, it suffered from the limited display resolution of the projectors. Since then, various CAVE systems have emerged, such as the NexCAVE, which are powered by multiple LCD display panels providing vastly increased pixel density [19].

The NexCAVE represents just one of many visualization systems that have adopted the approach of combining multiple commodity displays into a unified visualization platform [20]. Beyond benefiting from improved maintainability and increasingly affordable components such as desktop computer monitors or high-definition television displays, these tiled display systems aim to provide unprecedented pixel real estate for visualization tasks. Over the past decade, various frameworks have emerged for driving the clusters of computers powering these pixels, such as SAGE [54], CGLX [21], and DisplayCluster [33]. These frameworks provide different approaches for delivering visual content to the displays and managing a variety of input devices used to interact with and manage the content displayed. With multiple software frameworks available and established methods for constructing tiled display systems [44], researchers have been able to develop impressively large visualization systems. Currently, the world's highest resolution tiled display system, Texas Advanced Computing Center's Stallion (Figure 1), features a 16×5 array of 30-inch Dell LCD monitors for a combined resolution of 328 megapixels [66].

As large display systems continue to appear in more and more environments, numerous applications have been developed that aim to take advantage of the immense physicality and high visual detail afforded by these systems. Such applications range from digital brainstorming [24] to visualizing and comparing high-resolution multi-spectral and geospatial imagery [50, 73] to interactive museum installations [1], among many others [45]. These applications target a unique platform that provides various usability benefits but also results in a number of interaction challenges compared to that of a typical desktop or laptop computer system.

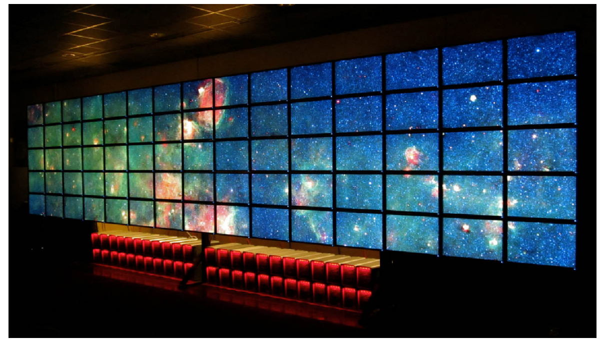

Figure 1: TACC's Stallion, the world's highest resolution display system weighing in at 328 megapixels [66].

2.1. Usability Benefits

Large, high-resolution visualization systems provide unique opportunities for enhanced usability due to their high visual resolution and the substantial physical space available around them. Space plays a vital role when interacting with these systems, both in terms of the floor space in front of the display in which users can act and the screen space used for displaying visual information. Because of our physical presence in the world, the way in which we manage the space around us is "an integral part of the way we think, plan and behave," allowing us to simplify choice, simplify perception, and simplify internal computation [35].

The high resolution afforded through these display systems naturally presents a "focus+context" interface to the user. A user can easily focus on the details of a specific region by moving closer to the display and then simply step back to obtain an overview of the entire visualization. This natural interaction scheme allows a user to maintain an understanding of the context of specific details within the larger visual space because of the continuous transition between the two. Focus+context interfaces have been shown to provide both better performance and higher subjective satisfaction compared to other techniques such as overview+detail and zooming/panning which involve a visual disconnect between the detail and overview viewpoints [10].

Utilizing this focus+context interaction technique involves physically navigating the space in front of the display, and numerous studies have revealed the benefits of such navigation compared to virtual navigation techniques such as zooming and panning. Ball et al. performed user studies comparing the use of a tiled high-resolution display with that of a standard monitor and found that users were able to find and compare targets faster within a visualization of finely detailed data when using the larger display [5]. The participants preferred using physical navigation with the large display as opposed to virtual navigation with the small display which resulted in loss of context and increased frustration.

In a similar experiment, user performance was studied using a range of tiled-display sizes [8]. The study found that as display size increased (i.e. additional display columns were added) virtual navigation decreased and performance times improved. The larger displays provided more overview and details at once which helped to guide physical navigation. Participants again preferred physical over virtual navigation and were observed to first physically navigate as much as possible before resorting to virtual navigation.

Additional studies have shown that users also benefit from the increased visual space of large displays. Tan et al. performed multiple experiments comparing the use of a standard desktop monitor to a large wall-sized display while positioning participants so as to keep the same visual angle between both displays [65]. They found that users performed better on spatial orientation tasks such as 3D navigation and mental map formation and memory which they attributed to the ability of the large display to immerse users within the problem space and bias them into using more efficient cognitive strategies.

Andrews et al. also observed cognitive benefits of increased display space with users of a large, high-resolution workspace [2]. They observed how participants used the large space to support external memory and organization. The visibility and persistence of objects in space provide cues to the user that aid in retrieving information, and the high-resolution of the display provides quick access to highly detailed information, requiring the user to remember less. The user benefits from reduced cognitive load as fewer context switches are needed when large amounts of information are available at a glance. The large display space also promotes arranging on-screen items, providing users a "semantic layer" in which they can encode additional meaning in the spatial relationships between these items.

2.2. Interaction Challenges

While the increased space made available through large, high-resolution visualization systems can enhance how we perceive visual information, it also makes these systems challenging to control and interact with. Kenneth Moreland points out the irony of how large-format display systems aim to increase the data flow from the computer to the human yet tend to simultaneously reduce the data flow from the human to the computer [40].

Various studies [16, 45] have identified multiple challenges faced when interacting with large, high-resolution display systems, which we summarize here:

Losing track of the cursor. In order to quickly traverse a large screen space and avoid having to repeatedly clutch and reposition the input device, higher cursor acceleration is typically used. This increased acceleration makes it difficult to keep track of the cursor during movement. A stationary cursor also poses a problem as it can be hard to find the cursor in such a large visual space.

Reaching distant objects. With increased screen real estate it becomes harder and more time-consuming to reach distant objects and move objects across the entire screen.

Managing space and layout. Increased display size results in increased space available to position on-screen items. It can become difficult to mentally keep track of everything that is happening on the screen at once and can result in complex multi-tasking behavior. With tiled-display systems the display bezels can also result in visual discontinuities and complicate layout management.

Failing to leverage the periphery. Larger displays offer a much larger periphery that can be leveraged to support user activities in ways not possible on standard displays.

Transitioning between interactions. Users may want to work up close to the display when investigating details but then interact from a distance when performing tasks involving an overview of the visualization. Ideally, interaction techniques should be able to support a smooth transition between interactions for up-close and distant tasks.

Moreland claims that "not until we discover appropriate interaction paradigms will large-format displays become truly useful" [40]. With unique challenges to address and the potential for improved visualization workflows, large display systems present an interesting opportunity for studying and developing human-centered interfaces.

3. Interaction Themes

The availability of increased interaction and visualization space provided by large, high-resolution display systems sets these systems apart from standard computing interfaces. Andrews et al. observe that these systems are "human scale" and thus the physicality of the human body is an important element when considering interaction [3]. Understanding how humans, as spatially located creatures, can utilize this space to effectively interface with the system presents an interesting research challenge. Here we explore two interaction themes that arise in discussions of interface design for large-scale visualization systems: *embodied interaction* and *ubiquitous computing*.

3.1. Embodied Interaction

Paul Dourish presents one definition of embodied interaction as "interaction with computer systems that occupy our world, a world of physical and social reality, and that exploit this fact in how they interact with us" [22]. Embodied interaction stresses the connection between the physical body and the mind and focuses on how humans generate and communicate meaning and understanding through physical interactions. Interfaces that utilize embodied interaction principles aim to enhance performance and insight by exploiting embodied resources such as motor memory, proprioception, spatial memory, peripheral vision, and optical flow [6, 7]. Research into the use of large, high-resolution visualization systems has revealed various ways in which humans benefit from the use of these embodied resources.

Ball et al. propose that the performance benefits they observed due to physical navigation give support for embodied interaction theory; users understand the display as a physical real-world object in their interaction space that they can navigate with respect to [8]. "By better utilizing embodied resources such as spatial memory, proprioception, and optical flow, people can more efficiently navigate large information spaces with less disorientation, thus enhancing performance by alleviating the cognitive resources to focus on the analytic task at hand" [6].

Andrews et al. claim that large displays also help to ease cognition by enabling users to arrange objects in space and use the perceptual system to recognize meaning in arrangements of objects rather than having to memorize these characteristics and relationships [2]. They note that the ability to arrange objects also harnesses the brain's ability to encode location information alongside non-spatial information such as text.

Beyond the cognitive benefits seen as large display users take advantage of space, various interaction techniques and devices also aim to leverage users' embodied resources. For example, devices that enable the user to interact with on-screen virtual items by touching or pointing at them present a natural interaction approach that mimics how humans interact with physical objects in our everyday lives. This approach attempts to make these devices and interactions easier to learn how to use by building upon our experience with using our bodies to interface with the physical world [4].

3.2. Ubiquitous Computing

Mark Weiser proposes that "the most profound technologies are those that disappear," enabling us to use them without thinking and focus beyond them on new goals [71]. By providing multiple intuitive ways to interface with computing systems that integrate with our environment, we can be empowered to more easily perform tasks and worry less about the specific technologies involved in these tasks. Ubiquitous computing is the concept that computing devices can be located anywhere and everywhere and can take make different forms, embedded into the natural human environment so that they are both "invisible" yet always accessible.

Large, high-resolution visualization systems have the potential to be an important component involved in ubiquitous computing environments. As these systems present a large space for non-conventional interaction, users can leverage multiple devices of various form factors that provide different modalities for interaction. In the literature this form of interaction is often referred to as "multimodal" interaction. Various research projects have developed multimodal interaction environments, such as the Stanford iRoom [32] and the WILD Room [11], that utilize large displays systems along with a combination of devices used to both control on-screen contents as well as supply new content to be displayed. These environments are meant to be highly collaborative and invite multiple users to interact with the system together using the interface that best suits their current task.

The themes of embodied interaction and ubiquitous computing provide insight into how technologies and interfaces to these technologies can be designed so that humans can use them as effectively as possible. Such interfaces must leverage our understanding of how we interact with the physical world and integrate into this world as seamlessly as possible.

4. Interaction Techniques

The increased physical presence of large display systems poses an interesting challenge when considering effective methods for interaction. These systems are inherently different from standard desktop displays, and thus standard desktop interaction techniques can prove inadequate. The keyboard and mouse have long been staples of desktop computing, optimized for the traditional "windows, icons, menus, pointer" (WIMP) user interfaces found on all modern consumer operating systems. Unfortunately, applying this approach to large-scale display systems restricts the user from moving around in the physical space, which is one of the major benefits of these environments. The mouse is limited to 2D movement and provides only a small number of input buttons, whereas humans can perform many more complex gestures with their bodies (e.g. pointing, pressing, grasping, rotating, snapping, and other movements in full 3D space). Traditional input devices prohibit users from taking full advantage of the 3D interaction space available with large visualization systems.

As such, new forms of interaction are required to fully realize the potential available through these systems. Interfaces for wall-sized displays have been a topic of research since the early 1990s when researchers at Xerox PARC developed *Liveboard* [23], a large interactive display designed to support collaborative group meetings and presentations. The Liveboard system acts as a digital whiteboard which allows interaction through a stylus, enabling users to digitally draw on the board and interact with on-screen items. With the goal of enabling anyone to easily walk up and start using the system, Liveboard represents one of the first attempts at exploring large display interfaces with a specific focus on human usability.

In the remainder of this section we explore various interaction techniques that have been developed for interfacing with large-scale visualization systems. We observe how different techniques attempt to address the interaction challenges inherent in these non-traditional workspaces and how they incorporate the concepts of embodied interaction and ubiquitous computing to enhance the user experience.

4.1. Pen Input

As we saw with Liveboard, pen-based input was one of the earliest forms of interaction developed for large-scale display systems. Pen input is designed to provide direct interaction with the screen and can be seen as a natural evolution of the WIMP paradigm applied to larger displays. Rather than using a mouse on a fixed surface, pen input enables the user to control the pointer of a WIMP-style user interface by operating directly on the display, leveraging the physical space available. It provides a very simple and natural means of interaction that can easily be adopted by users because of the familiarity with analog pens used in everyday tasks.

While the general concept of pen-based input is simple enough, various usability issues arise when put into practice. Elrod et al. designed the Liveboard pen with several input buttons, much like those found on a standard mouse, but observed that users found these buttons to be awkward and avoided using them. This led to the conclusion that input should be based solely on the pen's position and when it comes into contact with the screen. Reducing the physical input channels in this way means that either the graphical interface displayed must provide additional options for triggering the various functions and states previously performed through the buttons, or new approaches, such as gesture interpretation, must be implemented.

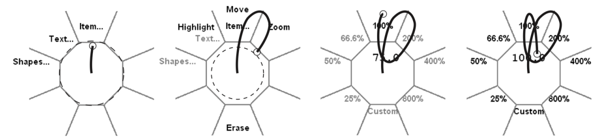

Figure 2: Using the *FlowMenu* to perform a zoom command via a continuous pen stroke. (Figure taken from [25].)

4.1.1. Minimizing Visual Clutter. Attempting to follow the WIMP paradigm for pen input could result in a multitude of on-screen widgets needed to perform various functions. These widgets add undesired visual clutter to the display and quickly become cumbersome to reach for and activate via the pen. In an effort to address these issues Guimbretière et al. developed the *FlowMenu* [25], a visual command-entry system that provides a minimally invasive graphical interface for fluidly performing a number of functions with simple pen strokes. The FlowMenu is presented as a radial menu that provides a hierarchy of options which can be traversed and activated by drawing pen strokes through the desired menu items (Figure 2). Rather than having a collection of graphical widgets continuously cluttering the screen, the FlowMenu condenses the necessary commands into a minimal interface that is optimized for pen gestures and only presented when needed. As users become more familiar with the system they can even rely on motor memory to perform "eyes-free" interactions for commands involving simple sequences of strokes.

FlowMenu presents a clever evolution of menu-based interfaces that is well suited for pen-based input and applications dominated by a focus on visual content where obstructions should be minimized. However, it also suffers from various limitations such as visual obstruction by the user's hand and inaccessibility when activated near screen borders. As visualization systems continue to grow in scale, the localized interaction method provided by FlowMenu fails to scale accordingly as items that are out of reach of the user can-

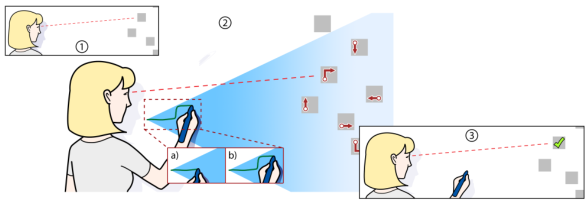

Figure 3: Using *Gesture Select* to select distant objects: User identifies a distant target to select (1), draws an initial stroke in the general direction of the target (2a), continuation marks appear on the targets (2), user continues the initial stroke by drawing the continuation mark of the desired target (2b), the target is selected on pen up (3). (Figure taken from [14].)

not be easily accessed.

4.1.2. Reaching Distant Objects. Several approaches have been taken to address the issue of pen-based access to distant objects on large displays. With the *vacuum* [12], Bezerianos et al. provide a virtual "vacuum cleaner" that brings toward it on-screen items that reside within the arc of influence centered at the point of invocation and spanning the entire display. Users invoke the vacuum by touching the pen to the screen and drawing in the direction in which the arc should extend, after which the arc can be expanded or redirected by further drawing. Once the arc is defined, scaled-down proxies of the items residing within the arc are brought into arm's reach of the vacuum, allowing users to then select an item via its proxy. Bragdon et al. present a similar arc-of-influence technique called *Gesture Select* [14]. Gesture Select allows the user to select a distant object by drawing an initial stroke in the general direction of the target followed by the "continuation mark" corresponding to the desired target, as detailed in Figure 3.

User experiments revealed that both of these techniques are able to reduce the amount of time required to select distant objects as compared to unaided selection where users had to physically move first before being able to select the target via pen. This result complements the observation that users prefer to move as little as possible to select remote items, even if the selection method incurs a slight overhead.

While these techniques can provide improved selection time of distant objects, various issues suggest that on-screen pen input does not provide an ideal interface for working with items spread across large displays. With Gesture Select, users found it difficult to determine which continuation mark was needed when the desired target was either too small, too far away, or too close to other targets, resulting in increased selection time and higher error rates. Additionally, as the vacuum technique gathers scaled-down proxies of on-screen items it changes the scale and spatial arrangement of the visualization; this can be disorienting and possibly negate the benefits of arranging items across a large spatial canvas to take advantage of visual cues and spatial memory.

The fidelity available through direct pen-based input lends itself more useful for localized interactions where a user is focused on up-close details within a specific region of the display. As such, pen input is perhaps best suited as a technique that complements other approaches which enable easier interaction with the visualization at larger overview scales.

4.2. Touch Input

Similar to pen input, touch-based input provides direct interaction with screen content but removes the need for additional devices such as pens. Whereas pen input is limited to a single point of contact and possibly a set of buttons, touch input can leverage the complex movements that can be performed by the human hand. Multi-touch input, in which gestures can involve multiple fingers rather than just a single touch point, can enable a multitude of interactions. By allowing direct manipulation of on-screen objects, touch input leverages our familiarity with real-world interactions, making the experience seem almost instinctual. However, because touch input enables such a wide range of possible interactions, there is still much research to be done regarding the best practices for designing intuitive touch-based interfaces [27].

4.2.1. On-Screen Touch Detection. Numerous approaches have been developed to enable touch-based input on large-scale display systems. Ringel et al. implemented a touch input system for back-projected SMARTBoard displays using behind-screen infrared (IR) illumination and a video camera equipped with an IR filter [55]. When a user touches the front of the screen, her fingertip reflects the IR light which is then visible within the video feed and processed to determine the touch location. Because this solution only works with projection-based displays, a different approach is required to enable touch interaction on tiled-display systems. One method of doing this is to arrange an array of cameras along the plane of the display. The *t-Frame* [61] system uses the known locations of these cameras and background subtraction-based image processing to triangulate the position of multiple fingers on the display surface (Figure 4).

Schick et. al present an interesting extension of this

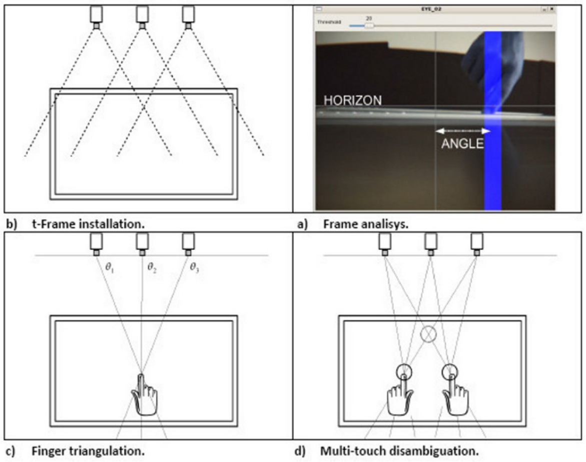

Figure 4: Touch detection via camera-based finger triangulation. (Figure taken from [61].)

Figure 5: Using a transparent touch frame, users see the target on the distant display at their fingertips and perceive touching it directly. (Figure taken from [41].)

camera-based approach by generating a full 3D reconstruction of the space in front of the display [58]. By performing voxel analysis on the scene captured by the cameras they are able to not only receive touch input but also provide users the ability to continue pointing at objects after their finger has left the display surface. Users found that this approach made it much easier to access all parts of the screen, alleviating one of the main drawbacks when using touch input on large displays.

Unfortunately, camera-based touch systems require calibration and can suffer from low precision due to the limited resolution of the cameras. A solution to these issues has emerged recently that involves the use of touch detection “frames” developed by companies such as PQ Labs. The PQ Labs *Multi-Touch Wall* [52] frame can be installed along the edges of displays up to 500 inches diagonal and can detect up to 32 touch points simultaneous with an accuracy of 3 millimeters. The frame contains thousands of embedded IR LEDs with corresponding IR sensors to determine touch locations based on the obstructions of the IR beams. The precision available with these frames nicely complements high-resolution display systems, enabling users to interact with small objects and fine details on the screen [72].

4.2.2. Off-Screen Touch Surfaces. As with pen input, on-screen touch interfaces do not easily allow for working with large-scale visualizations at an overview scale. To address this, various approaches enable touch-based interaction at a distance from the screen. One approach involves the placement of virtual overlays on the visualization to specify the desired interaction region. Touches are then mapped from the input device to this region, allowing interaction with on-screen objects within the region. Input devices for detecting hands have ranged from capacitive touch monitors [57] to overhead camera-tracked desks [38]. Overall users find overlays to be intuitive even though there is no longer direct manipulation of objects via touch. However, it becomes tedious to continually reposition and resize the overlay for the task at hand, whereas with direct manipulation users can select an object immediately, assuming it is within physical reach.

In an effort to provide the feeling of direct manipulation across an entire large-scale display, Müller-Tomfelde et al. [41] developed a transparent touch screen that they placed in front of the display, allowing users to see through the touch area and look at items on the screen (Figure 5). This presents the illusion that users are directly touching on-screen items and allows them to easily access any area of the screen. However, touch accuracy is highly dependent upon the user’s viewing angle as parallax comes into play, forcing the user to stand very still in order to maintain the proper alignment between the touch screen and the visualization.

Fixed, off-screen touch surfaces provide the benefit of interacting at an overview level, but limit the ability of the user to physically navigate within the space and prohibit the user from performing fine-grained interactions with more detailed visualizations.

4.3. Mobile Devices

Within the past several years we have seen a substantial proliferation of mobile devices such as smartphones and tablets. These devices play a substantial role in the widespread evolution of ubiquitous computing in which technology is easily accessible and found almost everywhere. Due to their popularity, it makes sense to leverage mobile devices as an interaction medium that is readily available.

In contrast to pen and touch input methods which can restrict the interaction space to within proximity of the display, smartphones and tablets free the user to move about the entire space and benefit from physical navigation. These devices come equipped with a variety of sensors – accelerometers, gyroscopes, microphones, cameras, multi-touch screens, etc. – as well as outputs such as audio and vibration, enabling numerous possibilities for acquiring user input and providing feedback. They also provide a secondary display that can be used to offload additional controls or information from the main display, alleviating clutter in the visualization and providing the user a personalized miniature interaction space.

4.3.1. Eyes-Free Interaction. One approach to mobile device-based interaction attempts to minimize the need to look at the device, allowing the user to focus on the visualization. Various implementations exist that provide this "eyes-free" interaction style, such as the *levels of precision (LOP) cursor* developed by Debara et al. [18]. The LOP-cursor leverages the multiple input modalities available on smartphones to control an on-screen cursor with two levels of precision (Figure 6). The device's pointing direction, calculated based on input from the on-board motion sensors, is used for coarse-grained placement of a rectangular control canvas on the display. With the canvas positioned in the desired area of control, the user can lock the position of the canvas by tapping and holding on the device's screen and then perform high-resolution control of the arrow cursor by dragging her finger on the screen. Satyanarayan et al. [57] present a similar technique using smartphones, but rather than providing a high-precision drag-able cursor they use a direct mapping from the device screen space to the control overlay space such that touches are visibly mirrored onto the large display within the control area.

Aimed at providing additional proprioceptive feedback with mobile devices, Jansen et al. [30] designed capacitive tangible controls that they affixed to multi-touch tablets. By adding these physical controls, such as sliders, users can modify parameters of the visualization without looking at the device as often, relying more heavily on tactile perception.

4.3.2. Personal Viewports. As opposed to eyes-free interaction, other approaches use the display of the mobile device as a key component of the interface to provide a personal viewing window. One of the most basic

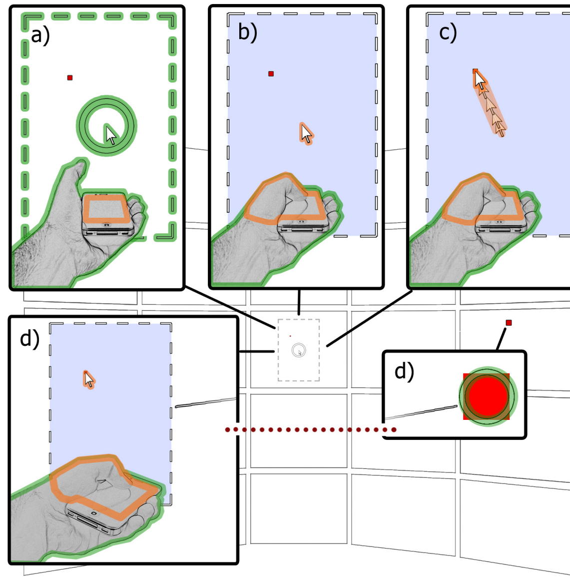

Figure 6: Using the *LOP-cursor* to position a rectangular control canvas and then perform fine-grained cursor movement. (Figure taken from [18].)

solutions involves mirroring the large-scale visualization onto the device's screen and providing touch-based input on the device [51]. Unfortunately this only allows coarse-grained input. A zoom ability could be added, but this would detract from the large-scale visuals and offer little if any advantage over direct on-screen touch input.

Rather than mimicking the visuals already presented on the main display, a secondary display can be especially useful when it provides access to additional information beyond the primary visualization. Applying the concept of "magic lenses" [13], we can think of mobile devices as virtual windows or filters through which we can view an alternate visualization that complements the primary (Figure 7). Advanced augmented reality algorithms optimized for mobile devices [69] utilize the camera to perform natural feature tracking, enabling the device to determine its location relative to the display without the need for external tracking systems. This enables each user to have their own personal window through which they can augment their view of the full-scale visualization [62].

These viewports can also be used as a mechanism for input. Interface elements available on the mobile device can be used to add annotations to items displayed on the large screen so that users can deal with contextual information without cluttering the main visualization [56]. Another technique developed by Jeon et al. [31] allows users to target on-screen items using the camera of their smartphone and then move the item by

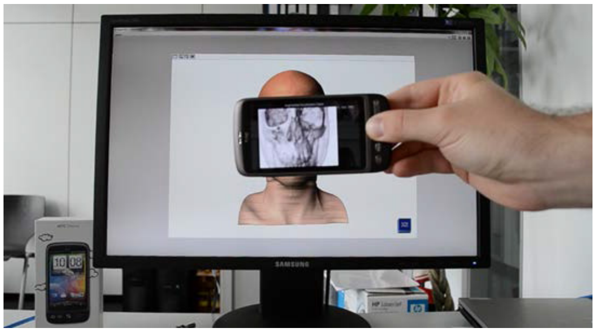

Figure 7: The mobile device acts as a “magic lens” that provides an alternative view of the visualization. (Figure taken from [62].)

simply aiming the device at the desired location. This approach provides immediate control of the visualization to ideally anyone who has a mobile device handy, although further work is still required to make this interaction as seamless as possible.

4.4. Pointing

Perhaps one of the most intuitive means of expressing interest in a remote object is by pointing at it. Because of this, a large body of research has investigated how to support pointing interactions with large-scale display environments in which interacting with objects from a distance occurs on a regular basis.

Ideally, pointing techniques will provide users the freedom to move about the space and not require additional visual attention so that the interaction feels as close as possible to natural pointing and does not distract from the visualization. However, numerous challenges arise when trying to develop natural pointing techniques that can allow both rapid and accurate target acquisition. Kopper et al. [36] identify several issues regarding high-precision pointing with large, high-resolution displays:

Natural hand tremor. The hand naturally trembles at a low amplitude, but when pointing from a distance this can result in an on-screen movement of many pixels. Moving farther from the screen amplifies this issue.

Heisenberg effect. When pressing a button on a device held in free-space, the user often slightly and unintentionally changes the position and orientation of the device. This is known as the Heisenberg effect, and it can cause the cursor to select unintentional locations when a click occurs.

Mapping varies with distance. When standing close to the display, changing the angle of the pointing device results in a smaller on-screen movement than when standing farther away from the display. This makes it difficult to perform small motions when standing far away from the display.

No parkability. With a traditional mouse, the user can stop moving the mouse and the cursor will remain in the same position. Pointing interfaces do not offer this parkability unless some sort of toggle is used, making it unnatural or tiresome to maintain a set cursor position.

No supporting surface. Using a mouse on a supporting surface allows the user to rest their hand and more easily make fine positional adjustments. Pointing interfaces lack a supporting surface, leading to fatigue and difficulty with performing small motions.

4.4.1. Modeless Techniques. Modeless pointing techniques involve a single interaction scheme in which the user is not required to switch between different modes of operation. The most straightforward implementation of such a technique is know as “ray-casting” in which the position of the on-screen cursor is determined by the intersection of a ray extending from the pointing device with the screen. With *Lumipoint* [17], the ray is physically defined by using laser pointers that intersect with a projection-based screen, and the intersection position is determined using cameras located behind the screen. The pointing ray can also be virtual, extending from devices that are either externally motion tracked or provide their own internal tracking via sensors.

Unfortunately, small items cannot be targeted reliably when using absolute ray-casting because of hand tremor and limited input resolution [42]. Various filters and thresholds have been applied to mitigate these issues but have limited benefit [68]. Another technique defines the virtual ray starting from the user’s eye location and then passing through the user’s hand. This perspective-based cursor has been found to improve performance for tasks requiring more accuracy, but it also leads to occlusion of the visualization and greater fatigue [34]. It also requires users to repeatedly switch between two very different focal lengths [42].

An alternative to ray-casting is relative pointing which involves a transfer function for mapping device or body movements into cursor movements. By implementing a dynamically adjusted gain within this transfer function, relative pointing techniques can provide higher precision while still enabling quick movement across the entire screen space [42]. However, indirect mappings have been observed as being less intu-

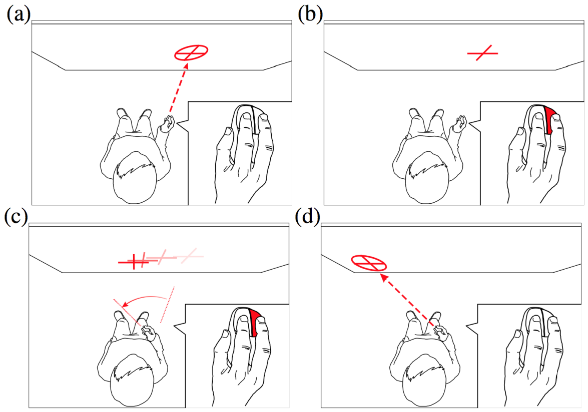

Figure 8: Dual-mode pointing. Ray-casting used for coarse-grained input (a). Activating precise mode via button press (b). Relative rotational movements control cursor movements (c). Switching back to coarse mode be releasing button (d). (Figure taken from [42].)

itive [17], and recalibration is required as the relative offset between the input device and the cursor grows, forcing the user to repeatedly perform a clutch operation [48].

4.4.2. Dual-Mode Techniques. Due to the limitations of modeless pointing techniques, dual-mode techniques have been developed which aim to provide enhanced precision by enabling the user to switch between coarse- and fine-grained input modes, as seen in Figure 8. Although switching between modes requires additional time, dual-mode techniques provide overall performance improvements when dealing with small targets and alleviate recalibration issues so that users do not have to clutch [42].

Kopper et al. [36] present two different implementations of dual-mode pointing. The first, *Absolute and Relative Mapping (ARM)*, works very similarly to the technique depicted in Figure 8 except that an additional device is held in the non-dominant hand to perform the toggle between the two modes so that the dominant hand can more easily perform the selection action. They found that users easily understood how to use ARM for interaction tasks but occasionally had trouble seeing what they were interacting with without moving closer to the display. When dealing with mostly overview tasks, it was observed that users preferred to stay farther away from the screen and only move closer when absolutely necessary. To adapt for this, they implemented another technique called *Zoom for Enhanced Large Display Acuity (ZELDA)* in which activating the precision mode displays a zoom window that not only provides finer-grained cursor movement but also enhances visual acuity, making it easier to see and work with smaller objects while standing far from the screen.

4.4.3. Device-Free Techniques. While pointing is a natural mechanism for targeting items at a distance, ideally interfaces for large-scale displays will enable a smooth transition between distant and up-close interaction. Hand-held pointing devices can make this transition awkward when the device does not have a clear application for up-close interaction with the display. Device-free pointing techniques alleviate this problem and provide a natural transition into close-range input such as touch.

Vogel and Balakrishnan [68] present a device-free pointing implementation that relies on a motion-tracked glove. They propose several techniques for pointing including a dual-mode approach in which the user performs a pointing gesture with her hand for absolute ray-casting and then opens the hand to enable more precise relative movement. To address the absence of input buttons, they also develop the *AirTap* gesture which allows the user to perform a click action by moving the index finger down and back up, similar to how the finger moves when clicking a mouse button.

Banerjee et al. [9] extend this glove-based tracking with *MultiPoint* in which multiple points are tracked on either one or both hands. Several activation gestures were tested, with users preferring the unimanual “breach” gesture in which moving their hands forward past a threshold distance activates interaction with the pointed-at object. This technique essentially provides an invisible multi-touch screen floating in front of the user, supporting actions such as dragging, scaling, and rotating objects. It also maps naturally to multi-touch input directly on the screen, providing a seamless transition between interaction distances.

Unfortunately, while these techniques aim to be device-free they still require users to wear specially designed gloves tracked by an expensive tracking system. The RGB camera-based voxel analysis approach discussed earlier [58] provides a truly device-less experience but suffers from limited resolution and thus cannot achieve the precision required for fine-grained pointing tasks. Advances in marker-less 3D tracking [60] can potentially enable the high precision necessary for a very natural device-free pointing technique for large, high-resolution displays.

4.5. Free-Space

Free-space interaction encompasses a variety of different techniques that make use of the human body operating in space as the main mechanism for input.

This class of techniques, which can be thought of as a superset that includes the various body-based pointing techniques previously discussed, is especially well-suited for large display interfaces as it enables users to freely navigate the space to their advantage while also utilizing embodied resources. It presents many of the same challenges that arise with pointing such as fatigue and lack of precision but offers a wide design space of possible gestures that can leverage users' experiences with everyday interactions.

4.5.1. Extending Reach. When interacting at a distance, one of the most salient challenges is communicating the desired target to operate on. We have seen that pointing provides one mechanism for accomplishing this, but the question arises as to what other body-based interfaces can enable users to extend their reach to items on the display.

One approach involves mapping the motions of the user's hands in 3D space into the 2D motions of two virtual cursors on the display. Lehmann and Staadt [37] present an "interaction scaling" technique in which this mapping scales dynamically based on the user's distance from the display such that moving farther from the display provides a larger range of on-screen interaction space. This scaled interaction corresponds nicely with the tendency of users to perform coarse-grained overview operations from farther away and more precise operations closer up. Unfortunately, it also causes cursor drift to accumulate with prolonged use, forcing users to either trigger a cursor reset or physically move to the appropriate distance so as to realign the cursor with their position.

Extending the idea of mapping hands onto the display, Shoemaker et al. [59] develop *Shadow Reaching* in which a shadow of the entire body is projected onto the display (Figure 9), and different virtual lighting techniques enable multiple methods of distant reaching. For example, the orthographic lighting behavior presents a shadow with minimal distortion regardless of the user's orientation, enabling precise control, whereas with the "user following" behavior the light source stays positioned directly behind the user such that turning the body enables reaching to distant areas of the screen. This technique also makes use of users' spatial memory and proprioception, allowing users to access virtual tools and control surfaces located on their body by simply touching the corresponding body part. The body can also serve as a container for digital information, providing users a natural sense of ownership for personal files that they carry "within their body."

Figure 9: *Shadow Reaching* used to access personal files stored "within the body." (Figure taken from [59].)

4.5.2. Freehand. One of the great benefits of touch screens is their affordance of instant access; anyone can simply walk up and immediately interact with virtual items without the need for anything but their hands. The concept of *freehand* free-space interaction extends this immediacy beyond the screen into the space, preferably eliminating the need to acquire a device or perform any calibration. The idealized conceptualization of such interfaces essentially transforms the user into a sort of magical conductor with the power to command anything on the display with a wave of the arm.

Working toward the goal of "magical" freehand interaction involves various practical considerations such as designing gestures that limit body movements to within the range of users' comfort zones, avoiding unnatural poses and reducing fatigue [53]. Because of the inherent lack of haptic feedback in freehand interactions, alternative forms of feedback – such as visual or auditory cues – can be implemented such that they help guide the user through the task at hand while also minimizing distraction [26].

Without an input device, physical body movements must be interpreted to communicate commands to the display system, requiring the use of body tracking systems. While marker-based tracking systems have been used for many years, recent advances in marker-less tracking systems such as the Microsoft Kinect [60] open up new potential for truly freehand input. Hespanhol et al. [26] used a Kinect to perform a study comparing different freehand gestures for selecting and rearranging items on a large display based on analogous actions performed in the real world (Figure 10). They found that users preferred the *grabbing* gesture and were able to perform the best with it once they overcame an initial learning curve. While the *dwelling* gesture provided the lowest barrier to entry, it also tended to slow down

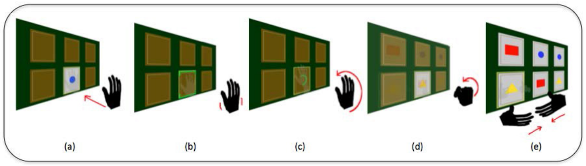

Figure 10: Freehand gestures for selection and rearrangement: (a) pushing; (b) dwelling; (c) lassoing; (d) grabbing; (e) enclosing. (Figure taken from [26].)

the user and was consequently thought to be an obstacle rather than a tool. The *pushing* gesture, while successful in other contexts such as touch screen input, fails to translate effectively to an environment in which no tactile feedback is available.

For pan-and-zoom interactions, Nancel et al. [43] found that two-handed techniques outperformed unimanual approaches as users had better control over the individual functions. They also compared linear and circular gestures for zooming and hypothesized that circular gestures would enable finer control and eliminate the need for clutching. Interestingly however, they found that users performed poorly with circular gestures due to a lack of physical guidance, and linear gestures were better understood as they mapped more naturally to the concept of pulling and pushing an object toward and away from the user.

Results from these studies indicate that free-space techniques perform best when they enable virtual interactions that build on users' familiarity with real-world interactions. As humans perform countless different movements to perform various tasks, an interesting area of research still remains on how to best translate these actions into effective human-computer communication.

5. Moving Forward

Given the number of different interaction techniques that have been researched and developed for large, high-resolution visualization systems, it is clear that no one solution will suffice for every scenario. These systems offer an enormous range of interaction possibilities with much to still explore. Ultimately, the effectiveness of an interface depends on the task at hand; thus a truly usable interface for large-scale display environments will ideally enable *multimodal* interaction, offering a multitude of ways for users to interact while providing a fluid transition between interaction techniques as tasks change.

The unique availability of space afforded by large display systems provides multiple interaction scales that users naturally navigate between when performing different tasks [47]. Multimodal interfaces should build on this behavior so as to provide the best possible form of input when operating within these different contexts. By doing so there is also potential to balance the strengths and weaknesses of individual interaction techniques by intelligently transitioning between techniques based on the task. For example, if users are inclined to move closer to the display in order to investigate details, the challenges of high-precision pointing can be avoided by employing a readily available high-precision touch detection frame.

Research on multimodal interfaces has been ongoing for many years now [32], and as technologies have evolved so too have the frameworks and infrastructures that enable this type of distributed interaction [11, 29, 64, 70], even to the point of commercialization [46]. As digital devices continue to pervade our daily lives we edge ever-deeper into a world in which computing is truly ubiquitous, and interfaces should continue to evolve in order to leverage this. These devices provide not only additional platforms for interaction but also sources of data which can feed the visualizations in question. Enabling personal devices to become part of this interactive ecosystem can help increase the accessibility of these systems so that any user can easily and quickly start using them.

Beyond the immense popularity of laptops, smartphones, and tablets, there has also been an increasing interest in "wearable computing" with devices such as Google Glass [63], the Meta glasses [39], and the MYO armband [67] recently emerging. These devices could open up a whole new range of possible interactions to explore and integrate with existing techniques, but the challenge will be to determine if they ultimately enhance the usability of the overall system.

Another exciting development is the next iteration of the Kinect sensor that will be released later this year alongside the Xbox One. The original Kinect has been one of the most popular devices for developers experimenting with freehand free-space interaction, and the new version promises to provide enhanced resolution and tracking precision that enables it to pick up details such as hand postures and shifts in body weight [49]. Traditionally this sort of tracking required expensive, specialized systems involving wearable markers, and the original Kinect simply did not have the precision necessary to enable fine-grained interactions. It will be interesting to see if the new Kinect can enable users to simply walk up and start using a large-scale display from a distance, providing both immediate and precise interaction that has not yet been possible.

6. Conclusion

Large, high-resolution visualization systems offer a unique platform for new scales of visual presentation and new forms human-computer interaction. Various benefits have been observed while using these systems in such a way that capitalizes on the available space, yet numerous challenges also arise when this space must be navigated and controlled effectively. We have surveyed a variety of different approaches for addressing these challenges and considered how these interaction techniques account for the needs and abilities of the human actor, specifically with regard to principles of embodied interaction and the concept of ubiquitous computing. While individual techniques can provide effective solutions for accomplishing certain tasks on large-scale displays, developing a truly usable interface requires a diverse ecosystem of interaction mechanisms that work seamlessly together. As new technologies continue to evolve and emerge, future research into interfaces for large visualization systems can leverage these advancements to provide even more effective methods for interaction.

Acknowledgments

I would like to thank my Research Exam chair, Charles Elkan, and my committee members, Bill Griswold and Jim Hollan, for taking the time to review my work in this field. I would like to acknowledge Scott Klemmer for kindly volunteering to serve on my committee as well. I would also like to thank my advisor, Falko Kuester, for his input and advice, as well as my good friends and colleagues at CISA3, Andrew Huynh, Vid Petrovic, Ashley Richter, Joe DeBlasio, John Mangan, Aliya Hoff, and TW, for their feedback and support. Finally, a very special and heartfelt recognition goes to Shirin Abdi for her never-ending encouragement during the process of completing this exam.

References

[1] J. Alexander, J. Barton, and C. Goeser. Transforming the Art Museum Experience: Gallery One. In *MW2013: Museums and the Web 2013*, 2013.

[2] C. Andrews, A. Endert, and C. North. Space to Think: Large, High-Resolution Displays for Sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 55–64, 2010.

[3] C. Andrews, A. Endert, B. Yost, and C. North. Information visualization on large, high-resolution displays: Issues, challenges, and opportunities. *Information Visualization*, 10(4):341–355, Aug. 2011.

[4] R. Ball, M. Dellanoce, T. Ni, F. Quek, and C. North. Applying Embodied Interaction and Usability Engineering to Visualization on Large Displays. Technical report, Virginia Polytechnic Institute and State University, Blacksburg, Virginia. USA, 2006.

[5] R. Ball and C. North. Effects of Tiled High-Resolution Display on Basic Visualization and Navigation Tasks. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems*, pages 1196–1199, New York, New York, USA, 2005. ACM Press.

[6] R. Ball and C. North. Realizing embodied interaction for visual analytics through large displays. *Computers and Graphics*, 31(3):380–400, June 2007.

[7] R. Ball and C. North. The Effects of Peripheral Vision and Physical Navigation on Large Scale Visualization. In *Proceedings of graphics interface 2008*, pages 9–16, 2008.

[8] R. Ball, C. North, and D. A. Bowman. Move to Improve: Promoting Physical Navigation to Increase User Performance with Large Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 191–200, 2007.

[9] A. Banerjee, J. Burstyn, A. Girouard, and R. Vertegaal. MultiPoint: Comparing laser and manual pointing as remote input in large display interactions. *International Journal of Human-Computer Studies*, 70(10):690–702, Oct. 2012.

[10] P. Baudisch, N. Good, V. Bellotti, and P. Schraedley. Keeping things in context: a comparative evaluation of focus plus context screens, overviews, and zooming. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 259–266, 2002.

[11] M. Beaudouin-Lafon, S. Huot, M. Nancel, W. Mackay, E. Pietriga, R. Primet, J. Wagner, O. Chapuis, C. Pillias, J. R. Eagan, T. Gjerlufsen, and C. Klokmose. Multisurface Interaction in the WILD Room. *Computer*, 45(4):48–56, 2012.

[12] A. Bezerianos and R. Balakrishnan. The Vacuum: Facilitating the Manipulation of Distant Objects. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 361–370, 2005.

[13] E. A. Bier, M. Stone, K. Pier, W. Buxton, and T. DeRose. Toolglass and Magic Lenses: The See-Through Interface. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 73–80, 1993.

[14] A. Bragdon and H.-S. Ko. Gesture Select: Acquiring Remote Targets on Large Displays without Pointing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 187–196, 2011.

[15] C. Cruz-Neira, D. J. Sandin, T. A. DeFanti, R. V. Kenyon, and J. C. Hart. The CAVE: audio visual experience automatic virtual environment. *Communications of the ACM*, 35(6):64–72, 1992.

[16] M. Czerwinski, G. Robertson, B. Meyers, G. Smith, D. Robbins, and D. S. Tan. Large display research overview. In *CHI '06 extended abstracts on Human factors in computing systems - CHI EA '06*, pages 69–74,

New York, New York, USA, 2006. ACM Press.

[17] J. Davis and X. Chen. Lumipoint: Multi-User Laser-Based Interaction on Large Tiled Displays. *Displays*, 23(5):205–211, Nov. 2002.

[18] H. Debarba, L. Nedel, and A. Maciel. LOP-cursor: Fast and precise interaction with tiled displays using one hand and levels of precision. In *2012 IEEE Symposium on 3D User Interfaces (3DUI)*, pages 125–132. Ieee, Mar. 2012.

[19] T. a. DeFanti, D. Acevedo, R. a. Ainsworth, M. D. Brown, S. Cutchin, G. Dawe, K.-U. Doerr, A. Johnson, C. Knox, R. Kooima, F. Kuester, J. Leigh, L. Long, P. Otto, V. Petrovic, K. Ponto, A. Prudhomme, R. Rao, L. Renambot, D. J. Sandin, J. P. Schulze, L. Smarr, M. Srinivasan, P. Weber, and G. Wickham. The future of the CAVE. *Central European Journal of Engineering*, 1(1):16–37, Nov. 2010.

[20] T. a. DeFanti, J. Leigh, L. Renambot, B. Jeong, A. Verlo, L. Long, M. Brown, D. J. Sandin, V. Vishwanath, Q. Liu, M. J. Katz, P. Papadopoulos, J. P. Keefe, G. R. Hidley, G. L. Dawe, I. Kaufman, B. Glogowski, K.-U. Doerr, R. Singh, J. Girado, J. P. Schulze, F. Kuester, and L. Smarr. The OptIPortal, a scalable visualization, storage, and computing interface device for the OptiPuter. *Future Generation Computer Systems*, 25(2):114–123, Feb. 2009.

[21] K.-U. Doerr and F. Kuester. CGLX: A Scalable, High-performance Visualization Framework for Networked Display Environments. *IEEE Transactions on Visualization and Computer Graphics*, 17(3):320–332, Apr. 2010.

[22] P. Dourish. *Where the Action Is: The Foundations of Embodied Interaction*. MIT Press, 2004.

[23] S. Elrod, R. Bruce, R. Gold, D. Goldberg, F. Halasz, W. Janssen, D. Lee, K. McCall, E. Pedersen, K. Pier, J. Tang, and B. Welch. Liveboard: a large interactive display supporting group meetings, presentations, and remote collaboration. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 599–607, 1992.

[24] F. Guimbretière, M. Stone, and T. Winograd. Fluid interaction with high-resolution wall-size displays. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 21–30, 2001.

[25] F. Guimbretière and T. Winograd. FlowMenu: Combining Command, Text, and Data Entry. In *Proceedings of the 13th annual ACM symposium on User interface software and technology*, pages 213–216, 2000.

[26] L. Hespanhol, M. Tomitsch, K. Grace, A. Collins, and J. Kay. Investigating intuitiveness and effectiveness of gestures for free spatial interaction with large displays. In *Proceedings of the 2012 International Symposium on Pervasive Displays - PerDis '12*, pages 1–6, New York, New York, USA, 2012. ACM Press.

[27] A. Ingram, X. Wang, and W. Ribarsky. Towards the establishment of a framework for intuitive multi-touch interaction design. In *Proceedings of the International Working Conference on Advanced Visual Interfaces - AVI '12*, page 66, New York, New York, USA, 2012. ACM Press.

[28] International Organization for Standardization. ISO 9241-210: Ergonomics of human-system interaction – Part 210: Human-centred design for interactive systems, 2010.

[29] R. Jagodic, L. Renambot, A. Johnson, J. Leigh, and S. Deshpande. Enabling multi-user interaction in large high-resolution distributed environments. *Future Generation Computer Systems*, 27(7):914–923, July 2011.

[30] Y. Jansen, P. Dragicevic, and J.-D. Fekete. Tangible remote controllers for wall-size displays. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12*, pages 2865–2874, New York, New York, USA, 2012. ACM Press.

[31] S. Jeon, J. Hwang, G. J. Kim, and M. Billinghurst. Interaction with large ubiquitous displays using camera-equipped mobile phones. *Personal and Ubiquitous Computing*, 14(2):83–94, Aug. 2009.

[32] B. Johanson, A. Fox, and T. Winograd. The Interactive Workspaces project: experiences with ubiquitous computing rooms. *Pervasive Computing, IEEE*, 1(2):67–74, 2002.

[33] G. P. Johnson, G. D. Abram, B. Westing, P. Navr'til, and K. Gaither. DisplayCluster: An Interactive Visualization Environment for Tiled Displays. *2012 IEEE International Conference on Cluster Computing*, pages 239–247, Sept. 2012.

[34] R. Jota, M. A. Nacenta, J. A. Jorge, S. Carpendale, and S. Greenberg. A Comparison of Ray Pointing Techniques for Very Large Displays. In *Proceedings of Graphics Interface 2010*, pages 269–276, 2010.

[35] D. Kirsh. The intelligent use of space. *Artificial Intelligence*, 73(1):31–68, 1995.

[36] R. Kopper, M. G. Silva, R. P. McMahan, and D. A. Bowman. Increasing the Precision of Distant Pointing for Large High-Resolution Displays. Technical report, Virginia Polytechnic Institute and State University, 2008.

[37] A. Lehmann and O. Staadt. Distance-Aware Bimanual Interaction for Large High-Resolution Displays. In *Computer Vision, Imaging and Computer Graphics - Theory and Application*, volume i, pages 97–111. Springer Berlin Heidelberg, 2013.

[38] S. Malik, A. Ranjan, and R. Balakrishnan. Interacting with large displays from a distance with vision-tracked multi-finger gestural input. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*, UIST '05, pages 43–52, New York, NY, USA, 2005. ACM.

[39] Meta. Meta Glasses. `https://www.spaceglasses.com/`, 2013.

[40] K. Moreland. Redirecting Research in Large-Format Displays for Visualization. In *Large Data Analysis and Visualization (LDAV), 2012 IEEE Symposium on*, pages 91–95, 2012.

[41] C. Müller-Tomfelde, K. Cheng, and J. Li. Pseudo-direct touch: interaction for collaboration in large and high-resolution displays environments. In *Proceedings of the*

23rd Australian Computer-Human Interaction Conference, pages 225–228, 2011.

[42] M. Nancel and E. Pietriga. Precision Pointing for Ultra-High-Resolution Wall Displays. Technical report, INRIA, 2011.

[43] M. Nancel, J. Wagner, E. Pietriga, O. Chapuis, and W. Mackay. Mid-air Pan-and-Zoom on Wall-sized Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 177–186, 2011.

[44] P. A. Navrátil, B. Westing, G. P. Johnson, A. Athalye, J. Carreno, and F. Rojas. A Practical Guide to Large Tiled Displays. In *Proceedings of the 5th International Symposium on Advances in Visual Computing: Part II*, pages 970–981, Las Vegas, NV, USA, 2009.

[45] T. Ni, G. S. Schmidt, O. G. Staadt, M. A. Livingston, R. Ball, and R. May. A Survey of Large High-Resolution Display Technologies, Techniques, and Applications. *Virtual Reality*, pages 223–236, 2006.

[46] Oblong Industries Inc. Mezzanine. `http://www.oblong.com/mezzanine/`, 2013.

[47] S. M. Peck, C. North, and D. Bowman. A Multiscale Interaction Technique for Large, High-Resolution Displays. In *2009 IEEE Symposium on 3D User Interfaces*, pages 31–38, 2009.

[48] C. Pelling, T. Sko, and H. J. Gardner. Be Careful How You Point That Thing: Wiimote Aiming For Large Displays. In *Proceedings of the 21st Annual Conference of the Australian Computer-Human Interaction Special Interest Group: Design*, pages 397–400, 2009.

[49] D. Pierce. The all-seeing Kinect: tracking my face, arms, body, and heart on the Xbox One. `http://www.theverge.com/2013/5/21/4353232/kinect-xbox-one-hands-on`, May 2013.

[50] K. Ponto, K. Doerr, and F. Kuester. Giga-stack: A method for visualizing giga-pixel layered imagery on massively tiled displays. *Future Generation Computer Systems*, 26(5):693–700, May 2010.

[51] K. Ponto, K. Doerr, T. Wypych, J. Kooker, and F. Kuester. CGLXTouch: A multi-user multi-touch approach for ultra-high-resolution collaborative workspaces. *Future Generation Computer Systems*, 27(6):649–656, June 2011.

[52] PQ Labs Inc. Multi-Touch Screen for LCD large screen, plasma TV display - Touch Screen Overlay. `http://multitouch.com/product_wall.html`, 2013.

[53] G. Ren and E. O'Neill. 3D selection with freehand gesture. *Computers and Graphics*, 37(3):101–120, May 2013.

[54] L. Renambot, A. Rao, R. Singh, B. Jeong, N. Krishnaprasad, V. Vishwanath, V. Chandrasekhar, N. Schwarz, A. Spale, C. Zhang, G. Goldman, J. Leigh, and A. Johnson. SAGE: the Scalable Adaptive Graphics Environment. In *Proceedings of WACE*, pages 2004–2009, 2004.

[55] M. Ringel, H. Berg, Y. Jin, and T. Winograd. Barehands: Implement-Free Interaction with a Wall-Mounted Display. In *CHI '01 Extended Abstracts on Human Factors in Computing Systems*, pages 367–368, 2001.

[56] J. Sanneblad and L. E. Holmquist. Ubiquitous Graphics: Combining Hand-held and Wall-size Displays to Interact with Large Images. In *Proceedings of the working conference on Advanced visual interfaces*, pages 373–377, 2006.

[57] A. Satyanarayan, N. Weibel, and J. Hollan. Using overlays to support collaborative interaction with display walls. In *Proceedings of the 2012 ACM international conference on Intelligent User Interfaces - IUI '12*, page 105, New York, New York, USA, 2012. ACM Press.

[58] A. Schick, F. van de Camp, J. Ijsselmuiden, and R. Stiefelhagen. Extending Touch: Towards Interaction with Large-Scale Surfaces. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*, pages 117–124, 2009.

[59] G. Shoemaker, T. Tsukitani, Y. Kitamura, and K. S. Booth. Body-centric interaction techniques for very large wall displays. In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction Extending Boundaries - NordiCHI '10*, pages 463–472, New York, New York, USA, 2010. ACM Press.

[60] J. Shotton, R. Girshick, A. Fitzgibbon, T. Sharp, M. Cook, M. Finocchio, R. Moore, P. Kohli, A. Criminisi, A. Kipman, and A. Blake. Efficient Human Pose Estimation from Single Depth Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PP(99), Oct. 2012.

[61] A. Soro, G. Paddeu, and M. L. Lobina. Multitouch Sensing for Collaborative Interactive Walls. In *Human-Computer Interaction Symposium*, volume 272, pages 207–212. Springer, 2008.

[62] G. Sörös, H. Seichter, P. Rautek, and E. Gröller. Augmented Visualization with Natural Feature Tracking. In *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia*, pages 4–12, 2011.

[63] T. Starner. Project Glass: An Extension of the Self. *Pervasive Computing, IEEE*, 12(2):14–16, 2013.

[64] D. Stø dle, J. M. Bjø rndalen, and O. J. Anshus. A System for Hybrid Vision- and Sound-Based Interaction with Distal and Proximal Targets on Wall-Sized, High-Resolution Tiled Displays. *Human-Computer Interaction*, 4796:59–68, 2007.

[65] D. S. Tan, D. Gergle, P. Scupelli, and R. Pausch. Physically large displays improve performance on spatial tasks. *ACM Transactions on Computer-Human Interaction*, 13(1):71–99, Mar. 2006.

[66] Texas Advanced Computing Center. Visualization Resources. `http://www.tacc.utexas.edu/resources/visualization`, 2013.

[67] Thalmic Labs Inc. MYO - Gesture control armband by Thalmic Labs. `https://www.thalmic.com/myo/`, 2013.

[68] D. Vogel and R. Balakrishnan. Distant freehand pointing and clicking on very large, high resolution displays. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*, page 33, New

York, New York, USA, 2005. ACM Press.

[69] D. Wagner, G. Reitmayr, A. Mulloni, T. Drummond, and D. Schmalstieg. Real-Time Detection and Tracking for Augmented Reality on Mobile Phones. *IEEE Transactions on Visualization and Computer Graphics*, 16(3):355–68, 2010.

[70] N. Weibel, A. Satyanarayan, A. Lazar, R. Oda, S. Yamaoka, K.-U. Doerr, F. Kuester, W. G. Griswold, and J. D. Hollan. HIPerFace: a Multi-channel Architecture to Explore Multimodal Interactions with Ultra-Scale Wall Displays. In *33rd International Conference on Software Engineering (ICSE)*, pages 1–10, 2011.

[71] M. Weiser. The computer for the 21st century. *Scientific American*, 265(3):94–104, 1991.

[72] B. Westing, B. Urick, M. Esteva, F. Rojas, and W. Xu. Integrating multi-touch in high-resolution display environments. In *State of the Practice Reports on - SC '11*, pages 8:1–8:9, New York, New York, USA, 2011. ACM Press.

[73] S. Yamaoka, K.-U. Doerr, and F. Kuester. Visualization of high-resolution image collections on large tiled display walls. *Future Generation Computer Systems*, 27(5):498–505, May 2011.