

A Survey of Large High-Resolution Display Technologies, Techniques, and Applications

Tao Ni¹ Greg S. Schmidt² Oliver G. Staadt³ Mark A. Livingston² Robert Ball¹ Richard May⁴

¹Department of Computer Science, Virginia Tech*

²3D Virtual & Mixed Environments Lab, U.S. Naval Research Laboratory†

³Department of Computer Science, University of California, Davis‡

⁴Pacific Northwest National Laboratory§

ABSTRACT

Continued advances in display hardware, computing power, networking, and rendering algorithms have all converged to dramatically improve large high-resolution display capabilities. We present a survey on prior research with large high-resolution displays. In the hardware configurations section we examine systems including multi-monitor workstations, reconfigurable projector arrays, and others. Rendering and the data pipeline are addressed with an overview of current technologies. We discuss many applications for large high-resolution displays such as automotive design, scientific visualization, control centers, and others. Quantifying the effects of large high-resolution displays on human performance and other aspects is important as we look toward future advances in display technology and how it is applied in different situations. Interacting with these displays brings a different set of challenges for HCI professionals, so an overview of some of this work is provided. Finally, we present our view of the top ten greatest challenges in large high-resolution displays.

CR Categories: A.1 [General Literature]: Introductory and Survey— [H.5]: Information Systems—Information Interfaces and Presentation H.1.2 [Information Systems]: Models and Principles— User / Machine Systems

Keywords: Large high-resolution displays, visualization, virtual environments, multi-monitor, projector array, distributed rendering, collaboration, user interfaces, interaction techniques, evaluation

1 INTRODUCTION

With technological advances, large high-resolution* displays are becoming prevalent in many fields. From multi-monitor configuration to tiled LCD panels to projection-based seamless displays, researchers have been constructing large displays with various hardware configurations. Two common features of such displays are increased physical size and higher resolution. Researchers, however, have encountered numerous technological difficulties with building large-format high-resolution displays, especially with tiling commercially available projectors into a seamless display landscape. Accordingly, they have proposed a number of techniques to address those fundamental problems and matured the construction of large high-resolution displays.

Meanwhile, researchers have been designing and implementing software toolkits that support large high-resolution displays. A single PC workstation is far from sufficient to drive a display made up of more than ten monitors or projectors. It is now common to use a

PC cluster, a set of computers connected via a high-speed network, for rendering on large high-resolution displays. Cluster rendering algorithms and systems have been interesting to computer scientists, and distributed rendering software as well as data streaming architectures have been made available to the public.

As an emerging technology, large high-resolution displays have been widely applied in various domains. The increasing popularity of large high-resolution displays [37] is paralleled by booming research efforts in addressing a fundamental question: how do users benefit from increased size and resolution? Many intuitively believe that large displays automatically outperform small ones. It is desirable, however, to understand why increased size and resolution are advantageous, and how we benefit from large high-resolution displays in accomplishing general or domain-specific tasks. Quantitative and qualitative experiments have been conducted, gathering empirical evidence to demonstrate the relationship between the changing visual effects afforded by emerging technologies and users' productivity and performance in collaborative and individual work.

As integral components in a computing environment, user interfaces and interaction techniques are constantly drawing attention of researchers. Many traditional user interfaces and interaction techniques become awkward or next to impossible to operate on large high-resolution displays [63]. Researchers have been attempting to modify or extend existing interface metaphors for large-format displays. They have also been creating novel interface techniques that scale well on large displays. In addition, solutions for input technologies other than mouse and keyboard have been proposed to interact with large-format displays. Pen-based techniques, laser pointers, and gestures are just a few examples.

Many challenges remain for large high-resolution displays in various aspects, and research on large high-resolution displays is very active and progressing rapidly, partially pushed by industry (e.g. high-definition TV systems) [63]. We survey the literature on large high-resolution displays, covering research aspects of hardware configurations (Section 2), rendering and streaming software (Section 3), applications of large displays (Section 4), visual effects and human performance (Section 5), and user interfaces and interaction techniques (Section 6). We conclude by proposing ten research challenges for utilizing and interacting with large high-resolution displays. By specifying these challenges, we hope to inspire original future research.

2 HARDWARE CONFIGURATIONS

2.1 CAVE and Derivatives

Among the first large-format display systems to receive widespread use was the CAVE™ (CAVE Automatic Virtual Environment), a projection-based *Virtual Reality* (VR) system that surrounds viewers in an immersive environment with four or more large display walls.

*The word resolution historically means the density of pixels on the screen, usually in terms of dots per inch (DPI) [13]. However, it is becoming common practice to refer to resolution as the number of pixels on a display, especially when people use the term “high-resolution displays.”

*{nitao,rgb6}@vt.edu

†{gregory.schmidt,mark.livingston}@nrl.navy.mil

‡staadt@cs.ucdavis.edu

§richard.may@pnl.gov

A CAVE system typically arranges four 10 ft \times 10 ft screens in a cube made up of three rear-projection screens for walls and a front-projection screen for the floor [23]. Five-wall [134] and six-wall configurations [99] exist as well, but they require special screens that can support the weight of users and/or movable screens to enable entry into the facility. One concern with constructing CAVE systems is the space for the optical path between projectors and screens. Most implementations fold each optical path with at least one mirror in order to shrink the footprint of the total system. Typically, all users wear stereo shutter glasses and one user wears a head tracker with six *degrees of freedom* (DOF). This enables all users to see stereo imagery, although only the tracked user will have a correct perspective projection. Tracked 3D input devices afford multiple DOF interaction; data gloves and joysticks are commonly employed in immersive virtual environments. Factors including an increased display area, a wider field of view and field of regard, multiple DOF, and stereopsis combine to generate an increased level of immersion [89].

The current highest resolution visualization room in the world [34] (Colorplate Figure C1) has been developed by Fakespace Systems [33] and is installed at the Advanced Simulation and Computing (ASC) Program's site at Los Alamos National Laboratory. The immersive viewing environment provides a 43 million pixel display across five rear-projected screens (three walls, floor, and ceiling). Nine tiles compose the front wall and six tiles for each remaining screen. A total of 33 stereoscopic digital projectors are seamlessly tiled to produce continuous images across the display environment which measures 15 ft wide by 10 ft deep by 12 ft high.

2.2 Multi-Monitor Desktop

Multi-monitor (or "multimon" in other literature) is an increasingly popular configuration in businesses and homes to extend a standard desktop PC with more screen real estate. Modern operating systems such as Windows, Mac OS X and Linux offer plug-and-play capability, and dual-head video output is common on even a modest graphics card. With PCI[®] and PCI Express[®] SLI[™] graphics cards, having multiple heads is also possible. Multi-monitor configurations are easy to configure, without demanding expertise in computer science or related subjects. Figure 1a shows an example of a multi-monitor setup for desktop computing. Figure 1b and Colorplate Figure C2 show a tiled-LCD multi-monitor desktop possibility.




Figure 1: Desktop configurations: (a) dual monitor and (b) tiled-LCD multi-monitor desktop display (Images courtesy Ball [4]).

Multi-monitor setups are receiving more attention in the research community. For example, bezel issues (i.e., edges of the tiles breaking up the continuity of the large display) have been examined in [70, 114]. Usability issues of multi-monitor setups for multi-tasking have been explored in [25, 4, 51]. Researchers also attempted to construct projection-based multi-monitor displays. Starkweather et al. created a curved desktop called DSHARP [109] using DLP projectors and parabolic mirrors. Bishop and Welch [12] created a "desktop" environment that used projections on the wall to alleviate bezel and ergonomic issues.

2.3 Tiled LCD Panels

Tiled LCD panels are sets of LCD displays arranged in a 2D array. The array can be arranged flat like a wall display (Figure 2a and

Colorplate Figure C3), flat like a table (Figure 2b and Colorplate Figure C4), curved, or in other configurations. The combined pixel count across the arrays can reach into the 100 million pixel range. For example, the Electronic Visualization Laboratory, University of Illinois (EVL-UIC) has developed a large 100 MPixel display called LambdaVision (Figure 2a and Colorplate Figure C3) [94, 95] and NASA has developed the Hyperwall which contains 49 LCD panels tiled in a 7 \times 7 array [103].

Some advantages for tiled LCD panels are: (1) they are easier to align and color correct than projectors; (2) they are less expensive than projectors, which use expensive bulbs with relatively short lifetimes; and (3) they take less space (no throw distance needed). A disadvantage is the array has borders between each tile (it affects displaying of text, but is not so bad for imagery).

The large tiled LCD panel displays offer a variety of uses due to their high pixel counts. Ultra-high resolution imagery such as geospatial data can be shown (and interacted with) in one contiguous display. Multiple display content (e.g., graphics, desktop content, video, and imagery) can be displayed all at once on different parts of the display.

EVL-UIC has developed a table-top tiled LCD display called the LambdaTable [62, 93] (Figure 2b and Colorplate Figure C4). The advantage of the table-top format is that it is familiar to many users who prefer a sandbox as a metaphor: every user can see the physical shape in a sandbox and can move to a location to modify the surface. The tangible nature of the interaction with the data is a big advantage. The table-top nature makes the system practical for many users to collaborate in an application.




Figure 2: Tiled LCD panels: (a) LambdaVision, a 100-MPixel wall and (b) LambdaTable, a horizontal tiled display and interface (Images provided by Luc Renambot, EVL-UIC [94, 93]).

2.4 Projector Arrays

Projector arrays can consist of CRT, light-valve, or LCD projectors. CRT projectors provide the most flexibility in terms of geometry control, but have limited size and brightness constraints. The light-valve projectors are bright and very flexible for expanding the overall size of the array. Since several light-valve imaging schemes are driven by scanning CRTs, they also have good geometry controls for the output image. LCD projectors offer a cost-effective, low-maintenance solution for arrayed projection; however, they have virtually no geometry controls.

Projector arrays are becoming popular due to their lack of bezels and the constantly-improving seamless integration of multiple tiles. A large amount of research has been done in the area of perfecting both the overall resultant display and the rendering algorithms. For example, research has been performed on color gamut matching [129], seams [110], misalignment [47, 19], luminance matching [71], and image blending [48]. Projectors also offer what CRT-based, LCD-based, and plasma monitors do not: a separation between the device size and the size of the displayed image. A small projector can be used to create a very large display or to create a very small display. The possible range of image size is limited by lumens, lens configurations, and available space. The resolution of the projectors, continually improving, is also a factor to consider. The highest resolution projection technology we are aware of is Sony's high resolution liquid crystal device, Sony 4K SXRD [108], which

can produce an image resolution of 4096×2160 . The technology will be used in high-end projectors. Additionally, arbitrary physical shapes can be used as the display surface. The result is that projector arrays enable reconfigurable and flexible display designs with (theoretically) little bezel distortion [90, 89].

Several examples of commercial high-resolution projection-based tiled walls are Fakespace's PowerWall™ [33], VisBox's stackable, reconfigurable projector array called VisBlock™ [124] (Figures 3a, 3b and Colorplate Figures C5), Cyviz's Vizwall™ [24] (Figure 4a and Colorplate Figure C6), and Barco's control rooms [6]. Example installations are shown at Lawrence Livermore National Laboratory [87, 104] and UC-Davis [121]. In meeting room settings, large projected displays have been tiled enabling increased display bandwidth (e.g., iLand [111], Alias's Visualization Studio [35], DIII-D National Fusion Facility's control room [1], and AT&T's Global Network Operations Center [130, 20]).




Figure 3: Stackable, reconfigurable projector arrays: (a) front and (b) side views of VisBlock™ from VisBox Innovative Display and Interaction Technologies (Images courtesy VisBox.com [124]).

2.5 Stereoscopic Displays




Figure 4: Stereoscopic displays: (a) Cyviz's stereoscopic projector array called Cyviz Vizwall™ (Image courtesy Cyviz.com [24]) and (b) EVL-UIC's autostereoscopic display called Varrier™ (Image courtesy EVL-UIC [102]).

Stereoscopic displays show two sets of pixels for an image, making one set visible to the user's left eye and the other to the right eye. Typically the user is required to wear special glasses or viewing aids to see the 3D effects. However, a recent development—autostereoscopic displays—eliminates the need for special glasses. High-resolution stereoscopic displays are typically physically larger than their low-resolution counterparts and must account for the following factors: (1) larger space for head tracking (recommended to properly view the 3D effects) and (2) larger number of pixels to be displayed (twice as many). These factors are especially important if the user desires to interact in real-time, operate closely to the display surface, and move around the space in front of the display.

One of successful examples of high-resolution stereoscopic displays is Cyviz's Vizwall [24] (Figure 4a and Colorplate Figure C6). It is based on Cyviz Viz3D™ passive stereo capable display technology [24], which includes two projectors mounted on a positioning system. Vizwall tiles several Viz3D modules in certain arrays (3×2 in Figure 4a), leading to a high-resolution passive stereoscopic display. EVL-UIC has developed an autostereoscopic display called Varrier™ [102] (Figure 4b and Colorplate Figure C7), which does not require users to wear any stereo glasses to view the 3D effects. Their approach involves a curved LCD tiled display with a parallax barrier affixed to the front. The user is free to move within an

area of approximately $32 \text{ in} \times 48 \text{ in}$ ($81.3 \text{ cm} \times 121.9 \text{ cm}$). In addition, Liao et al. [67] have developed a high-resolution display using Integral Videography technology and 9 XGA projectors arranged in a 3×3 array, leading to a total resolution of 2872×2150 pixels. Their system generates geometrically accurate high-quality autostereoscopic images, and reproduces motion parallax in 3D space without any special viewing glasses and head trackers.

Another innovative high-resolution stereoscopic display is the D-vision from Tokyo Institute of Technology [46]. D-vision uses 24 projectors to provide stereoscopic projections on a hybrid screen. A combination of rear and front projection provides high-resolution in the flat central region in front of the user and lower resolution on curved screens around the periphery.

2.6 Volumetric Displays

Another class of high-resolution displays is volumetric displays. Rather than placing the pixels on a single surface, they “stack” voxels (3D pixels) via different technologies to show depth. One type of volumetric display is the swept volume display, which uses a fast rotating surface to which images from one or more projectors (or laser sources) are projected. The rotating display surface can be opaque or semi-transparent and is rotated at sufficient speeds to render it mostly invisible to the viewer. Swept volume displays are available or being developed commercially from companies such as Actuality Systems, Felix 3D-Display, Genex Technologies, and Hitachi [83].

Another type of volumetric display is LightSpace Technologies' DepthCube™ [112] (Figure 5a, 5b and Colorplate Figure C8). The DepthCube uses a single projector with twenty liquid-crystal projection screens stacked up in front of the projector. The screens are at five-millimeter intervals. At any given time, nineteen of the LCDs are transparent and only one is scattering the projected light. The projector displays 1000 images per second so the total volume is refreshed at $1000/20 = 50 \text{ Hz}$ refresh rate.




Figure 5: Volumetric displays: (a) external view and (b) cutaway mockup of DepthCube™ display from LightSpace Technologies (Image courtesy LightSpace Technologies [112]).

The XenoVision Dynamic Matrix Display [133], which may also be considered a volumetric display, is a table-top display that moves the physical display surface within a six-inch vertical range according to an elevation file. Special rendering algorithms correctly project a full-color image file onto the irregular surface. The $3 \text{ ft} \times 4 \text{ ft}$ display area takes about two minutes to reach the desired elevations and 30 seconds to erase.

3 RENDERING AND STREAMING

3.1 Architectures and Data Distribution

Cluster rendering can be described as the use of a set of computers connected via a network for rendering purposes. Rendering may take many forms: distributed video streaming, non-photorealistic volume rendering, ray tracing and radiosity-based rendering, or interactive rendering using application programming interfaces (APIs) like OpenGL. We distinguish between display data streaming software and distributed rendering software. Data streaming toolkits enable streaming of any type of data for large display systems.

Molnar et al. [74] classified parallel 3D rendering algorithms into three general classes based on when the sorting of the primitives occurs in the transition from object to screen space. The three classes are sort-first, sort-middle, and sort-last.

Sort-first In sort-first algorithms, the display is partitioned into discrete, disjoint tiles. Each rendering node of the cluster is then assigned one or more of these tiles and is responsible for the complete rendering of only those primitives that lie within one of its tiles.

Sort-middle Sort-middle algorithms begin by distributing each graphics primitive to exactly one “processor” for geometry processing. After the primitive has been transformed into screen space, it is forwarded to another processor for rendering.

Sort-last In sort-last approaches, each primitive is sent to exactly one node for rendering. After all primitives have been rendered, the nodes must composite the images to form the final image. This usually requires a large amount of bandwidth because each node must send the entire image to a compositor.

This taxonomy is generally conceived as considering a single, multi-processor machine, but it applies equally well to a cluster of independent machines that produce a single conceptual display. Tiled displays lead naturally toward a sort-first approach. The screen is already partitioned into tiles, with each tile driven by a single cluster node. Having multiple machines with separate memories complicates the task of keeping the data consistent, however.

Cluster rendering systems vary widely with respect to the way data is distributed among the cluster nodes. Chen et al. [17, 18] first looked at the problem of data distribution. Two general models have emerged: client–server and master–slave.

Client-server In the client–server model, a user interacts with a single instance of the application that runs on a client node. This client generates the geometry and distributes it to the render servers (Figure 6a). Graphics processors offer two rendering modes – immediate mode and retained mode. In immediate mode, the client sends the primitives over the network every frame. In retained mode, each render server stores and reuses primitives it receives. The client then needs to send only changes to the geometry. This method is usually accomplished through a scene graph.

Master-slave In the master–slave model, the application executes on every cluster node. Execution of the application on all nodes must be synchronized to ensure consistency among all application instances. Typically, a master node handles all user interaction and synchronizes state changes between all other nodes (Figure 6b).




Figure 6: Cluster rendering distribution schemes: (a) client–server, (b) master–slave.

The master–slave approach usually requires the least amount of bandwidth. The results of user interactions and other state changes are sporadic and relatively simple to transmit over a network. This approach, however, is not transparent, as everything affecting program execution must be considered as input. Timers, random number generation, system calls, or any variables influencing program execution need to be distributed and synchronized among the nodes. The client–server approach is usually transparent to the programmer. The program can be implemented as if it were running on a single machine and the system will handle the rest.

3.2 Display Data Streaming Software

There are several software toolkits that enable streaming, splitting, and displaying many types of data (e.g., 2D imagery, videos, and desktop content) for large display systems.

TeraVision The TeraVision system captures and distributes visual imagery from any graphics platform over a high-speed network [107]. TeraVision has been demonstrated to successfully stream real-time microscopy images at 2000×2000 pixel resolution from the National Center for Microscopy and Imaging Research (NCMIR) to Hawaii, to provide researchers at the University of Hawaii access to NCMIR’s microscope.

SAGE EVL-UIC has developed the Scalable Adaptive Graphics Environment (SAGE), a graphics streaming architecture that supports seamless display of networked applications across high-resolution displays. The data from the applications may include 3D renderings, video streams, 2D geospatial imagery, etc. SAGE supports hundreds of megapixels of contiguous display resolution and provides dynamic pixel routing capability, which allows users to freely move and resize an application’s imagery over tiled displays in real time [56, 96]. SAGE relies on TeraVision to distribute visual imagery across different platforms [107].

EVL-UIC VNC Viewer EVL-UIC developed a Virtual Network Computer (VNC) protocol client that enables users to bring desktop content to the SAGE environment. EVL-UIC’s VNC Viewer is a standard VNC client program modified to serve as a proxy between a VNC server (of any size and pixel depth) and SAGE. Once the pixels are retrieved from the VNC server, the same pixels are given to the SAGE API for immediate display. SAGE supports any number of simultaneous VNC applications, making use of the large real estate offered by high-resolution tiled displays.

IBM’s Scalable Graphics Engine Perrine and Jones [85] developed a parallel rendering environment for the IBM Scalable Graphics Engine (SGE). Their toolkit supports tunneling, which allows graphics applications to communicate directly with the SGE, SMP rendering, and includes an OpenGL implementation that utilizes the SGE. Based on this toolkit, they implemented a parallel MPEG video player.

VLC Media Player The VLC media player [125] is a cross-platform video player that supports a large number of video and audio codecs. VLC support media streaming and includes a “wall” video filter that splits the output in several tiles.

OptiStore and LambdaRAM OptiStore is a high-performance, low-latency data retrieval system that filters raw 2D and 3D volumetric data and produces a sequence of visual objects. OptiStore uses LambdaRAM to achieve low-latency data access by aggressively using high bandwidth networks and caching.

JuxtaView JuxtaView is an application for visualizing extremely high-resolution montage images on scalable tiled displays, which are able to deal with the largest montages produced by biologists [61]. JuxtaView uses OptiStore to enable rapid panning and zooming through enormous images accessed over wide-area high-speed networks.

Scalable Visualization Consumer The Scalable Visualization Consumer (SVC) receives MPEG2 data through a Firewire® (IEEE1394) interface, files on disk, or network interface and decompresses it for streaming to a tiled display. The MPEG2 data is decompressed, split as sub-images, and streamed to the appropriate display nodes.

NCSA Pixel Blaster The National Center for Supercomputing Applications (NCSA) developed Pixel Blaster [79], a distributed high-definition movie player, which reads raw-format HD images and distributes them to display nodes. The application was ported by EVL-UIC to operate with SAGE.

3.3 Distributed Rendering Software

Raffin and Soares [88] present a review of common software toolkits supporting PC clusters for Virtual Reality systems, including CAVELib, VR Juggler, Syzygy, OpenSG, Chromium, and others. Although they discuss distributed rendering software in the context of VR and parallelism, these tools are naturally applicable to large high-resolution display systems. For example, Chromium [50] has been widely used to support interactive parallel visualization applications displaying to tiled displays. OpenSG [127] implements rendering on tiled displays by dividing the screen into $M \times N$ uniformly spaced tiles with a render server assigned to each of them. We summarize several additional toolkits that specifically handle rendering for large high-resolution display systems.

Distributed Multihead X (DMX) Typical X-Windows servers provide support for multiple displays connected to the same machine via the *Xinerama* extension. The DMX project [29] provides a proxy X server that is a front end to X servers running on each rendering node in a cluster. The X client application will connect to the front-end server; rendering requests will be broken down as needed and sent to the appropriate back-end servers via X11 library calls. DMX is transparent to the application and supports standard mouse and keyboard input through the XInput extension.

Aura Aura [122] is a multi-platform API designed for scientific visualization on tiled displays. In “broadcast” mode it implements a client-server model. It provides the user with a scene graph interface to take advantage of frame-to-frame coherence. Aura also provides a master-slave configuration called “multiple copies”.

Virtual Immersive Reality Program Interface (VIRPI) Germans et al. [40] developed VIRPI on top of Aura [122] to provide a high-level user interface toolkit. VIRPI provides standard widgets, such as menus and radio buttons, as well as a framework for measuring in virtual spaces.

Blue-c Distributed Scene Graph The blue-cTM distributed scene graph by Näf et al. [78] is based on OpenGL Performer [100]. To support collaboration and cluster rendering, it has been enhanced with node serialization and state update interfaces. Additional custom nodes and attribute objects are integrated to support multimedia elements [76, 77].

Virtual Graphics Platform (VGP) ModViz, Inc. developed the Virtual Graphics Platform [123], a cluster rendering toolkit for Linux environments. Similar to Chromium, it is transparent to the application and intercepts OpenGL function calls, which are distributed to rendering servers running on each cluster node. VGP supports large-scale multi-screen projection displays and image compositing.

Renderizer ModViz, Inc.’s Renderizer [123] enables multi-display cluster rendering for OpenGL Performer-based applications. It replaces a small number of OpenGL Performer function calls and eliminates the need to parallelize existing applications explicitly.

EVL-UIC OpenGL Wrapper The EVL-UIC OpenGL Wrapper allows easy porting of native OpenGL applications to SAGE. The wrapper operates similarly to WireGL by intercepting the “glSwapBuffer” command, and then it streams the OpenGL data to SAGE.

Vol-a-Tile Vol-a-Tile [106] is a volume visualization tool for large-scale, time-series scientific datasets rendered on high-resolution

scalable displays. These large-scale datasets can be dynamically processed and retrieved from remote data stores over optical networks using OptiStore. Vol-a-Tile utilizes the fast OpenGL 3D texturing and fragment shaders.

NCSA TerraServer Blaster The NCSA TerraServer Blaster [80] is an OpenGL-based interface to the terraserver (www.terraservice.net), a 3.3-terabyte online database of high resolution USGS aerial imagery for all of the United States. It can be used together with Chromium to display on high-resolution display walls.

4 APPLICATIONS OF LARGE HIGH-RESOLUTION DISPLAYS

4.1 Command and Control

Large high-resolution displays have been widely installed in command and control centers for a variety of applications including military, aerospace, and telecommunications. The Air Force Research Laboratory developed the Interactive DataWall, which is an ultra-high-resolution large screen display that has an interface using wireless interaction devices. Jedrysik et al. [55] use the Interactive DataWall for situational awareness and collaborative decision-making tasks involving battlefield data.

A large four-wall immersive room has been deployed at the Naval Research Laboratory (NRL). The 3D Virtual & Mixed Environment Laboratory developed a submarine command application for detecting target submarines that operates in the immersive room [105]. A sonar operator is immersed in an oceanographic view, where a set of tracking sonar buoys are dropped, and can visualize the output of several different tracking algorithms that have been developed (Figure 7a and Colorplate Figure C9).

Christie constructed a large, high-resolution, 198,000 ft² communication command and control center for AT&T in 1999 [130, 20]. The control center, AT&T’s Global Network Operations Center, has given AT&T the unparalleled capability to manage the flow of communications traffic across its network anywhere in the world from one location. The center contains over 75 high-resolution projected displays, each used to visualize computer-generated data and graphics associated with over 250 million voice calls on a typical weekday.




Figure 7: Applications: (a) NRL’s immersive room demonstrating a submarine command & control application, and (b) Deere & Company’s VR testing facility for virtual prototyping vehicle designs and testing their effectiveness (Photo courtesy Deere & Company, Moline, Illinois [27])

4.2 Vehicle Design

It has been a fundamental requirement of the automotive design industry to display and interact with vehicle models at 1:1 scale [15]. Therefore, automotive design studios have explored the use of a variety of large-format digital displays in their design workflow, including tape drawing, electronic drafting tables, ImmersaDeskTM VR systems, CAVEs, and PowerWalls. Deere & Company is developing applications that simulate vehicle operations using large VR display facilities like the one shown in Figure 7b and Colorplate Fig-

ure C10. They utilize such applications to evaluate human factors and ergonomics, analyze complex engineering data, and build capabilities in vehicle manufacturing process development [27].

4.3 Geospatial Imagery and Video

Large high-resolution displays offer the sense of scale needed for geospatial imaging and large film-quality video applications. The ability to obtain realistic terrain representations, zoom across scales, and create fly-through animations certainly benefits geospatial visualization. Large displays allow users to see critical details in complex dynamic phenomena, such as subtle eddies that are critical to understanding global ocean circulation models [22].

High-resolution display systems are used by several major oil and gas companies for geospatial exploration and engineering, 3D mapping, and geophysical analysis [31]. Evans et al. [32] used high-resolution displays to develop interactive prototypical spatial models of forest stands from LIDAR and multi-spectral data. They use the large displays to visualize a spatially true models of the stands. MacEachren [69] studied how the large displays impact geospatial analysis. They revealed that iconic data representations, interaction methods, and remote collaboration techniques need to be considered when using the displays for geospatial analysis.

4.4 Scientific Visualization

Large high-resolution displays have been one of the favorite choices for scientific visualization applications because they offer (1) viewing of data at true-to-life or human-scale physical sizes and (2) viewing of large amounts of data simultaneously with the increased number of pixels available. One organization using large high-resolution display for scientific visualization is Oak Ridge National Laboratory. They are performing a variety of scientific visualization research including scientific simulations of galactic supernovae, protein expression, nanostructure models, and fusion energy devices [81]. They have a 30 ft \times 8 ft display wall with a resolution of more than 11,000 \times 3,000 pixels (35 million pixels) installed at their Science Visualization Facility. Another organization is EVL-UIC, which is developing applications for visualizing geoscience data like seismic activity, flow of water through the earth, and physical processes of the inner earth. The applications have been designed for their high-resolution tile wall called GeoWall2 (a 5 \times 3 array of LCD panels) and other displays [66].

4.5 Collaboration and Tele-immersion

An integral part of collaborative work is a public display surface that serves as a medium for presenting, capturing, and exchanging ideas. Several examples of display surfaces and collaborative research associated with them are itemized.

Liveboard Xerox PARC's Elrod et al. [30] pioneered a large interactive display system called Liveboard, which provides computer-supported group meetings, presentations, and remote collaboration. The Liveboard is a stylus-based, interactive large-format display, which can be used as a meeting-support tool providing whiteboard-like functionality, allowing users to write down ideas and retrieve documents.

UNC's Office of the Future UNC [91] proposed the technology of "Office of the Future", combining wall-sized high-resolution displays, cameras, and several types of interaction techniques. The goal of the system is to provide very compelling tele-collaboration among distant individuals.

Dynamo Izadi et al. [54] designed and implemented Dynamo, a communal multi-user interactive surface with which users can easily access and interact. This device allows people meeting in public

spaces to share and manipulate digital information, such as documents, video, and images, in a manner that emulates physical documents.

Alias Visualization Studio The Alias Visualization Studio [35, 2] houses a number of state-of-the-art corporate meeting facilities, including large display walls (Colorplate Figure C11). The designers indicated the importance of good interior design, drawing on theatrical and cinematic principles, and social concerns, such as privacy, aesthetics of large displays, and attention management (including distractions on secondary displays).

DIII-D National Fusion Facility Control Room The DIII-D National Fusion Facility control room [1] has installed a large, shared display wall composed of three 50-inch Toshiba P500DK data wall cubes arranged horizontally creating a nearly seamless 3840 \times 1240 pixel tiled display. The US National Fusion Collaboratory Project is utilizing the large shared display wall for control room collaborations in the following ways: (1) presenting up-to-date information about experiment status and group activity; (2) sharing data between personal desktop screens and the large display wall; and 3) video-conferencing of core team members from remote sites.

Tele-immersion is a new paradigm of collaboration. It holds a promise of creating collaborative visualization and VR applications over national and global high-speed networks. A large high-resolution display is an ideal facility for tele-immersion applications, since collaborative exploration of massive scientific data sets requires a large screen real estate. In the late 1990s, the National Tele-immersion Initiative (including UNC-Chapel Hill, Brown, and Univ. of Pennsylvania) was formed to experiment with realistic human representation at a distance (over Internet2) in order to facilitate tele-collaboration [101]. The blue-c project expanded those results and unified them with lessons learned from CAVE-related research [41, 64, 77].

The Air Force Office of Scientific Research is supporting work at Iowa State University to develop new technology for monitoring and control of unmanned aerial vehicles (UAVs). Iowa State is developing a system that provides real-time data in a tele-immersive large-scale display environment and allows operators better control of the UAVs [128].

4.6 Education and Training

Large high-resolution displays are a great tool for education and training in astronomy, bioinformatics, medical imaging, urban planning, and geographic information. Researchers at EVL-UIC have built an advanced tiled-visualization system called GeoWall2 to develop informal science education applications for museums, galleries, and other public events [66]. UC-Santa Cruz [73] has developed a collaborative learning environment for the classroom by using a large shared tile-wall display. The display space is shared by the instructor and students. The large display space provides the primary means of presentation of lecture material, allowing the lecturer to keep multiple screens of material in view for the students. The students can download any portion of the display, enabling them to go back at any time to previous screens of the presentation. The display also permits students to show their work, ideas, and even pose questions.

4.7 Immersive Applications

Iowa State University's Virtual Reality Applications Centers (VRAC) has been developing immersive applications with their large-display environments over the last several years [128]. They have projects for virtual prototyping thermal systems in a power plant and developing virtual engineering tools for livestock production. The Naval Research Laboratory [119] and others [128] have

developed fire training tools using an immersive room that enable users to practice fire safety procedures inside of buildings, ships, and other environments.

4.8 Public Information Displays

Advances in flat panel display technology and projector research have led to consumer grade devices and the ability to use them in many novel ways, such as the use of high-resolution imagery in pervasive public displays. Where there once might have been a printed image or a low grade digital image, we are now starting to see large tiled public displays [36]. No longer are the displays limited to flat billboards either—most any surface or collection of surfaces has the potential to receive digital enhancement [28].

Many of the concepts of pervasive displays have roots in ambient spaces research [132]. In traditional computer applications, including VR, user actions are explicit. That is, users consciously direct the computer to perform some desired task. In ambient and pervasive applications, actions could be implicit. The computer could recognize an action and respond appropriately, even when there is no intended computer interaction or an awareness that user actions could trigger a response from the computer. Inherent in this effort to separate the display from the computing hardware are many of the ideals expressed in ubiquitous computing [131].

5 VISUAL EFFECTS AND HUMAN PERFORMANCE

When large interactive displays were constructed, researchers realized that the changing visual effects afforded by the increased display landscape would have a profound impact on how users work with computing workspaces. As Swaminathan and Sato [113] observed, “when a display exceeds a certain size, it becomes qualitatively different.” Prior research not only qualitatively observes in detail user behavior on such displays, but also quantitatively explores visual effects of large high-resolution displays on users’ task performance. Practical usability issues of both hardware and software have been uncovered, and guidelines for design and presentation of information systems on large high-resolution displays have been formed.

5.1 Qualitative Evaluations

Large displays were initially applied in building collaborative workspaces, and benefits of large-format displays for groups working together have been demonstrated. Despite the Liveboard project’s focus on underlying hardware and a pen-based user interface, an informal survey found it was most frequently used for group meeting facilitation. Guimbretiere et al. [44] also explored pen-based interaction with high-resolution wall-size displays, and have tested their design with professional product design groups engaged in brainstorming tasks.

Ball et al. [4] reported an observational analysis of the use of a large tiled display consisting of 9 LCD monitors in 3×3 array over the course of six months. Although it is not a controlled experimental evaluation in its nature, it provides insightful feedback concerning common usage of how users do and do not use a large high-resolution display to perform ordinary tasks, such as reading papers, surfing web pages, viewing images, programming, and entertaining. For example, they found a bezel adaptation strategy employed by most users when working with tiled LCD surfaces, showing that bezels tend to help users quickly separate multiple applications and tasks, which significantly decreases context switching. Based on their observations, they summarized the advantages and disadvantages of using tiled high-resolution displays and formed design recommendations and guidelines for application designers.

5.2 Quantitative Tests of Spatial Tasks

With LCD panels becoming less expensive, it is an emerging trend to have personal multi-monitor configurations in offices. Also, it is not uncommon to install individual large displays with affordable projectors. It therefore opens up numerous research opportunities with respect to individual gains and user behavior on large displays.

Tan et al. [117, 116] conducted two studies to quantify benefits of physically large displays for individual users. They took an approach of designing controlled experiments, identifying independent and dependent variables explicitly while holding constant other factors. This approach, compared to previous practical experiments, allows precise statistical analysis of the results to identify the causes of any observed difference in performance or usability data [14]. They maintained a constant viewing angle for each of two displays in their experimental design. In the first study, participants performed spatial orientation tasks involving static 2D scenes. A significant performance gain was observed on the large display, even though the two displays cast identical-size retinal images. In the second study, they designed triangle completion tasks to examine how display size affected path integration performance in interactive 3D virtual environments. Not surprisingly, users were more effective on the physically large display. They suggested that large displays may afford a greater sense of presence, leading users to adopt an egocentric rather than an exocentric strategy in spatial tasks. Further studies [117] found the benefits of being able to interact with the environment were independent of the physical size of the display.

Researchers have also been interested in examining how large displays reduce gender bias in navigating virtual environments. Existing reports have suggested that males significantly outperformed females in VE navigation. However, Czerwinski et al. [26] undertook controlled experiments, uncovering that women were able to achieve similar VE navigation performance to men if they were exposed to a large display coupled with a wide field of view. A follow-up study by Tan et al. [115] attempted to identify what factors were driving 3D navigation gains for females. They indicated that the gender-specific navigation benefits come from the presence of optical flow cues, which were better afforded by a wider field of view on large displays.

While recommendations and guidelines for design and presentation of interactive 3D virtual environments on large displays have been made, we lack insights in the effectiveness of large high-resolution displays for basic low-level data visualization and navigation tasks. Ball et al. [5] described an exploratory study comparing basic visualization task performance on a large tiled display with a resolution of 3840×3072 and two smaller displays (1560×2048 and 1280×1024). They concluded that with finely detailed data, large high-resolution displays result in more physical navigation, which is preferable to virtual zoom-and-pan navigation on smaller displays.

In addition, researchers have started to explore the possible usage of large-format displays in creating semi-immersive environments. Patrick et al. [84] conducted an empirical study to examine how users acquired spatial knowledge in a VE presented on a head-mounted display (HMD), a large projection display, and a desktop monitor. An intriguing finding was there was no observed significant difference in reproduction of VE-survey knowledge when users viewed through an HMD versus a large screen, implying that large displays may be an effective low-cost alternative for HMDs for VE applications.

5.3 Quantitative Tests for Information Display

Polys et al. [86] considered how varied display size affects information layout interface design in an information-rich virtual environment (IRVE). Two information layout techniques to support search and comparison tasks were designed and evaluated. The Object

Space technique was to associate textual labels relative to their referent virtual objects, and the Viewport Space technique was to display labels on an image plane workspace. They altered two levels of software field of view (SFOV) on a single monitor and a 3×3 tiled LCD display, respectively. Search and comparison tasks were designed, and both accuracy and time were measured. They identified a significant advantage provided by the Viewport Space technique combined with a wide SFOV. Since an IRVE is a spatial, perceptual VE enriched by abstract information, it is useful in many application domains where pure perceptual information is not sufficient, such as scientific visualization, military simulation, architecture walkthrough, education, and entertainment. Intuitively, IRVE applications may benefit greatly from large high-resolution displays. Polys et al.'s work represents a trend toward designing and evaluating original techniques for new displays to increase users' productivity in IRVE applications.

The Infocockpit [118] showed that the memory retention of a user was improved 56% by using multiple monitors to spread the information around the user's space, creating a sense of presence. Lin et al. [68] found a correlation between memory retention of a virtual environment, the sense of presence gained from within that environment, and the size of the field of view (FOV). They showed that increased FOV improved performance on memory-based tasks and correlated with a better sense of presence. They did not attempt to establish evidence for the cause of this correlation.

6 USER INTERFACES AND INTERACTION TECHNIQUES

While user interfaces for standard desktop displays have been developed over a few decades now, there has been relatively little work on interfaces for large format displays. We identify large display usability issues and interaction challenges, and then examine the work in the field by dividing the efforts into two categories: work aimed at 2D displays and work aimed at 3D displays. While techniques for the former tend to emerge as extensions of desktop interfaces, techniques for the latter share more with virtual reality (VR) interfaces.

6.1 Usability Issues and Interaction Challenges

To design usable and useful interface and interaction for large high-resolution displays, we must understand which factors have left conventional interface techniques awkward to use on large displays. We identify five categories of large display usability issues.

1. Reaching distant objects. As screen real estate grows, it is increasingly difficult for users to access objects scattered around on a wall-sized display [98], especially when they tend to stay relatively close to a large display. For example, if a user seated in front of a large screen tries to drag a file icon near the lower right corner to the recycle bin icon on the left, it will be a terrible experience to use a traditional drag-and-drop interaction paradigm. It is also a common problem with more modest multi-monitor configurations, since accessing an icon or window at a distance requires more cursor moving, which takes time and raises cursor-tracking issues [98]. Reaching distant information becomes even harder with heterogeneous monitor configurations (e.g. a SmartBoard combined with a regular LCD panel plus a PDA).

2. Tracking the cursor. With increased physical screen size, users employ higher mouse acceleration to traverse large displays [8]. The faster the mouse cursor moves, however, the more difficulty users have keeping track of it. In addition, during meetings or presentations with large-format display facilities, a speaker is likely to point a cursor at a target to direct audience's attention. Locating a stationary cursor, however, becomes increasingly problematic on large displays [59].

3. Crossing bezels. Using multiple monitors is still a popular configuration to gain extra working space. Multi-monitors display bezels are beneficial in allowing users to organize multiple tasks onto different monitors [98]. Problems occur, however, when users cross bezels. A window or an image may be sufficiently large to occupy several monitors, creating visual discontinuity at bezels. When a cursor traverses across a bezel, there is normally a discrepancy between its actual traveling course and what users may expect, since there is no virtual space underneath the bezels.

4. Managing space and layout. Interaction with large high-resolution displays imposes many space and layout management issues, especially when windowing systems are used. On desktop displays, various window or task management systems exist, such as Apple Exposé [3]. Effectively handling space and layout on large-format displays, however, is by no means trivial.

5. Transitioning between interactions. Based on tasks a user is likely to perform, interaction with large displays can be categorized into two broad paradigms. For tasks involving dealing with detailed information, working up close to a large display is reasonable. There are other tasks, however, that are best performed from a distance, such as sorting photos and pages or presenting large drawings to a group [126]. Also, large displays often feature touch screen capacity. Consequently, techniques allowing a smooth transition from up-close interaction to interaction at a distance and vice versa are needed.

6.2 Extending Existing GUIs and 2D Metaphors

Traditional desktop metaphors such as Windows, Icons, Menus, and Pointing (WIMP) do not always scale well to large format displays. Even for large displays that are still situated on physical desktops, the amount of screen real estate can make pointing to or even finding windows, icons, and menus difficult. Baudisch et al. [8] found that a high density of cursor positions seen on the screen could assist users in visualizing the mouse path, which resulted in a small improvement of performance in moving to distant targets and higher user satisfaction. Robertson et al. [98] propose a set of extensions to the traditional mouse actions and window capabilities that aim to enable the user to find, move, and manage desktop objects more easily. They also address issues such as the confusion caused by the bezel between physical displays, which limits usability for systems that simply place multiple displays in proximity.

A number of techniques allow users to share portions of windows [11, 98], which can be quite useful for moving tools near areas in large-format displays where they can be helpful, or for briefly viewing windows at remote locations of the desktop without having to move to a new location and lose context. Another set of techniques helps draw the user's attention to specific areas of the display by changing the relative illumination of items of interest [9]. In a search-and-identify task, Spotlight [59] was found to significantly increase the user's speed in identifying the desired information. The application may also give users the ability to manually move windows large distances across the display [39].

Drag-and-drop is a popular technique for passing information between applications (including the operating system and an application) or collecting items in desktop interfaces. Since this technique requires the pointing action in traditional WIMP interfaces, it suffers on large-format displays. Several attempts have been made to alleviate the difficulties, notably moving potential targets near the selected object [7, 21]. Such techniques improve user reach and time over drag-and-drop and have competitive or lower error and drop rates. It is interesting to note that users sometimes performed faster with traditional drag-and-drop over short distances, reflecting the familiarity with that technique and the time for the user to orient to the assembled candidate targets.

A technique called Vacuum [10] pulls objects towards the user's invocation point so that they may enter the user's physical reach. Upon release, the desktop returns to its previous (pre-vacuumed) state. This technique increased speed in tasks that consisted of selecting multiple targets, but not for the task of selecting a single target, with an increase in error rate of under three times, mostly due to missing small representations of objects that should have been selected.

Another novel GUI widget and interaction technique, known as a "Frisbee" [58], addresses the problem of arm's length direct manipulation of inaccessible regions of large-format displays in an application independent way. A frisbee consists of a local "telescope" and a remote "target", acting as a portal to another part of a display. Interactions performed within the telescope are applied on the remote data, which are surrounded by the target and drawn in the telescope. This technique satisfies major design principles such as minimizing physical travel and visual disruption, maintaining visual consistency, and supporting concurrent multi-user interactions, and proves to be advantageous over conventional "click-walk-click" interaction style for wall-sized displays.

Direct drawing on 2D displays is not yet a standard interface tool, but it can be a powerful metaphor for a wall-sized display. The Interactive Mural used a pen-based drawing tool [44] with which the user could draw, write text, or manipulate parameters such as the scale and position of the display space.

6.3 3D User Interfaces for Large Displays

General VR tracking systems, whether based on magnetic, mechanical, optical, ultrasonic, or other technologies, have often been used as 3D input devices for large display spaces. Through their ability to enable pointing via ray-casting and similar techniques, they can be used to select objects at arbitrary distances from the user in a 3D space. With their six DOF input, they can be mapped to navigation operations in 3D; many such interfaces have been inspired by helicopter controls.

The Interaction Table [45] provides six DOF input with a series of devices that together can provide an interface similar to a 3D wheel for navigation or manipulation. This can provide a natural interface and has the benefit of providing passive feedback to the user.

Grossman et al. [42] developed a 3D modeling interface for large-scale displays. The interface integrates a variety of methods that work well for large-scale interaction. Some of these include 2D construction planes spatially organized in a 3D volume, tape drawing for curve and line creation, and continuous two-handed interaction.

The VisionWand [16] represents a compromise between the difficult task of tracking hands and the ease of users forming gestures. Using a wand with colored endpoints creates a device that is easy to track, but requires little hardware in the user's hands. The user can then gesture with such methods as tapping, tilting, rotating, or moving the wand to control parameters. Gestures can control standard user interface items such as menus and widgets to enable the user to manipulate more parameters.

LaViola et al. [65] developed a set of novel input devices for CAVE-based virtual environments. They employ hand gestures, such as pointing with a tracked, finger-worn sleeve or foot gestures, such as tapping toes or heels on a map with a foot-worn slipper for navigation. They enable object selection with flexing and pinching while wearing a glove with buttons added to engage actions. Similar user actions can be measured by computer vision systems [52, 60], but the robustness of such systems still presents difficulties such as incorrectly recognized gestures.

Malik et al. [72] use multi-hand gestures to enable the user to control an object and the workspace simultaneously, thus allowing the user to bridge the distance between objects, or to offer the user a wider range of gestures by allowing the two hands to work together.

They activate widgets under the hands to similarly add to the power of gesture-based interactions.

Vogel et al. [126] use hand gestures to indicate typical user interface actions such as point-and-click when working at a distance from the display surface. In order to achieve a stable pointing operation, they filter the detected finger position. They found that closing the thumb to the index finger does not appear to have any speed or accuracy benefit over tapping in mid-air, despite the kinesthetic feedback offered by the former. They found that a ray-based pointing operation was faster but inaccurate compared to a technique that combined a clutching gesture to engage control and hand motion to move. The latter technique enabled similar performance to a hybrid approach that used approximate pointing.

The Interactive Workspaces Project [57] explored interface possibilities for people working together using large displays. They integrated a variety of interaction devices and techniques including wireless multimodal devices. Other similar work combining interaction devices with display walls can be found in [97, 92, 75, 53, 38, 82].

6.4 User Interface Evaluation

While many believe that large-format displays automatically provide benefits, the evidence is not quite so clear. Tyndiuk et al. [120] found that navigation to objects that are already in view does not benefit from a large display over a standard desktop display, whereas navigation to unseen objects does (presumably due to benefit in the search portion of the navigation). They also found that performance on a manipulation task improved with the large display over the standard desktop. They further found a correlation between visual attention and performance with large displays compared to small displays.

7 TOP TEN RESEARCH CHALLENGES

After consideration of the issues associated with large high-resolution displays, we have come up with a list of what we believe are the top ten research challenges faced by this community. We hope the challenges inspire future research projects involving large high-resolution displays.

1. Truly seamless tiled displays. Tiling projected images to form a "seamless" large high-resolution display has been a popular approach, and we have witnessed a lot of work done on image blending and geometric registration. However, unresolved technological problems exist, such as variations of color and luminosity, which may easily break the illusion of a *single seamless display*. Calibration is also reported to be a headache in practice.

2. Stereoscopic large high-resolution displays. Building large-scale, high-resolution headtracked stereoscopic displays is a key challenge to producing a high-resolution immersive virtual environment experience. The Varrier autostereoscopic display developed at EVL-UIC [102] and a few projector-based high-resolution solutions (e.g., Cyviz's Vizwall [24]) are achievements heading in the right direction toward solving this challenge.

3. Easily reconfigurable large high-resolution displays. Creating displays that can easily be reconfigured and support several form factors (e.g., flat, curved, and other representations) is desirable. Today's reconfigurable displays often require tedious hours of realignment after any shape reconfiguration. Future reconfigurable displays should be easy to pack, move, align, and color calibrate.

4. High-performance cluster rendering. There is a growing number of software APIs and toolkits for cluster rendering with support for high-resolution displays. A number of toolkits support only master-slave data distribution. These schemes are easy to implement and allow legacy applications that can't be parallelized to run

in a cluster environment. However, they do not fully utilize cluster resources since every node has to run a copy of the application and there is no performance increase compared with running the application on a single machine. Some toolkits intercept OpenGL function calls and distribute rendering data to cluster nodes. While this is transparent to the application, these toolkits require high-bandwidth and low-latency connections between cluster nodes. There is no ultimate solution for cluster rendering today, and the choice for a particular toolkit largely depends on the application.

5. Scalability. The majority of tiled-display installations today is limited to fewer than twenty tiles. Existing cluster rendering and display data streaming software provides adequate support for these types of systems. However, we expect new challenges with the advance of massively-tiled displays such as EVL-UIC's 55-tile LambdaVision system.

6. Design and evaluate large high-resolution display groupware. Although large-format displays are appealing technologies to support collaborative interaction, it remains challenging to design and evaluate groupware applications that fully exploit their capacity and potential [49]. Grudin outlined research challenges for creating desktop groupware applications in [43], most of which still hold up in designing groupware applications for large high-resolution displays. However, unique characteristics and requirements of large-scale display groupware systems present new challenges. For example, users perceive large high-resolution displays in radically different ways due to their form factors, which may affect the groupware interface design. Also, large display groupware applications have to accommodate semi-public or public contexts, which may affect the visibility and privacy of interactions [49].

7. Effective interaction techniques. Traditional mouse and keyboard and associated standard desktop interaction techniques are not sufficient for interaction with large-format displays. A number of interaction techniques have been investigated for large displays including natural gestures, voice recognition, multi-handed interaction techniques, and methods to improve the reach of the user. Also, several user interface metaphors for facilitating specific tasks such as windows management and distal target access have been explored. These techniques show promise, but need to be evaluated for specific tasks in order to gain a better understanding of how effective they are for interacting with large displays.

8. Perceptually valid ways of presenting information on the large displays. The field of view has a demonstrable effect on the perception of the user; this is the most obvious but not the only feature of perceptual significance for large displays. Factors such as apparent brightness, contrast, and resolution heavily affect the user's understanding of information on a display. Non-visual effects, such as ergonomic comfort, may also play a role.

9. Empirical evidence for the benefits of large high-resolution displays. A taxonomy that matches low-level tasks, perceptual factors, and high-level applications with different form-factor display types (e.g., resolution, physical size, configuration) would enable laboratories to gain the most benefit for their applications without having to risk investing significant resources on less appropriate systems. We should analyze the benefits and limitations of large high resolution for a range of tasks. We have presented some work in this area, but much more is needed.

10. Integrating large high-resolution displays into a seamless computing environment. It is an increasing trend to bring portable computing devices into workspaces. A seamless computing environment is an infrastructure supporting a smooth integration of portable devices and pre-installed computing facilities, making it possible to fluidly exchange, share, and store information within a collaborative

group. While large high-resolution displays are becoming prevalent to support collaborative tasks, a challenging question arises: is there an effective way to seamlessly integrate emerging display technologies with existing heterogeneous devices? Associated are many fundamental research issues to address including low level networking and hardware interfaces, supporting software toolkits, user interfaces, and social interaction.

8 CONCLUSIONS

We have presented a comprehensive survey of prior research on large high-resolution displays, identified major unresolved problems, and from these made a list of what we feel are the top ten challenges in the field. Our expectation is that the survey and top ten challenges will bring attention to what has been done and what needs to be done for research involving large high-resolution displays. We hope to inspire the display community into having more discussion and debates about research in this field, and more importantly, spark efforts to maintain an updated list of achievements and current research challenges in the field. Scientific research for large high-resolution displays has been advancing tremendously in the past decade, and we believe it is far from running out of steam. We feel it also will evolve into an interdisciplinary research area with a stimulating and challenging agenda in the future.

ACKNOWLEDGEMENTS

We wish to thank Luc Renambot and Jason Leigh (EVL-UIC), Mike Eisenhard (CYVIZ), Paul Rajlich (VisBox, Inc.), Jeff Brum and Don Garwood (Fakespace Systems), Doug Bowman (Virginia Tech), Azam Khan (Alias), Presley Salaz (LANL), Simon Julier and Dennis Brown (NRL), and Mary Leonard (Deere & Company) for their contributions.

REFERENCES

- [1] G. Abela, G. Wallace, D. Schissel, and et al. Shared display wall based collaboration environment in the control room of the DIII-D national fusion facility. In *WACE*, 2005.
- [2] Alias. www.alias.com/eng/about/visualization_studio/index.shtml. Last accessed January 4, 2006.
- [3] Apple's exposé. www.apple.com/macosx/features/expose. Last accessed December 21, 2005.
- [4] R. Ball and C. North. An analysis of user behavior on high-resolution tiled displays. In *Interact*, 2005.
- [5] R. Ball and C. North. Effects of tiled high-resolution display on basic visualization and navigation tasks. In *CHI*, 2005.
- [6] Barco. www.barco.com. Last accessed December 21, 2005.
- [7] P. Baudisch, E. Cutrell, D. Robbins, and et al. Drag-and-pop and drag-and-pick: techniques for accessing remote screen content on touch-and pen-operated systems. In *Interact*, 2003.
- [8] P. Baudisch, E. Cutrell, and G. Robertson. High-density cursor: a visualization technique that helps users keep track of fast-moving mouse cursors. In *Interact*, 2003.
- [9] P. Baudisch and R. Rosenholtz. Halo: a technique for visualizing off-screen locations. In *CHI*, 2003.
- [10] A. Bezerianos and R. Balakrishnan. The vacuum: Facilitating the manipulation of distant objects. In *CHI*, 2005.
- [11] A. Bezerianos and R. Balakrishnan. View and space management on large displays. *IEEE CG&A*, 25(4), 2005.
- [12] G. Bishop and G. Welch. Working in the office of "real soon now". *IEEE CG&A*, 20(4), 2000.
- [13] D. Bowman, E. Kruijff, J. LaViola, and et al. *3D User Interfaces: Theory and Practice*. Pearson Education, 2005.
- [14] D. Bowman and D. Raja. A method for quantifying the benefits of immersion using the cave. *Presence-Connect*, 5, 2004.
- [15] W. Buxton, G. Fitzmaurice, R. Balakrishnan, and et al. Large displays in automotive design. *IEEE CG&A*, 20(4), 2000.

- [16] X. Cao and R. Balakrishnan. Visionwand: interaction techniques for large displays using a passive wand tracked in 3d. In *UIST*, 2003.
- [17] H. Chen, Y. Chen, A. Finkelstein, and et al. Data distribution strategies for high resolution displays. *Computers and Graphics*, 25, 2001.
- [18] Y. Chen, H. Chen, D. Clark, and et al. Software environments for cluster-based display systems. In *IEEE/ACM International Symposium on Cluster Computing and the Grid*, 2001.
- [19] Y. Chen, D. W. Clark, A. Finkelstein, and et al. Automatic alignment of high-resolution multi-projector displays using an un-calibrated camera. In *IEEE Visualization*, 2000.
- [20] Christie. www.christiedigital.com/recentinstallations/networkoperation/att/att_noc.asp. Last accessed January 5, 2006.
- [21] M. Collomb, M. Hascoët, P. Baudisch, and et al. Improving drag-and-drop on wall-size displays. In *Graphics Interface*, 2005.
- [22] Computer Science and Telecommunications Board. *IT Roadmap to a Geospatial Future*. National Academies Press, 2003.
- [23] C. Cruz-Neira, D. Sandin, and T. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the cave. In *ACM SIGGRAPH*, 1993.
- [24] Cyviz. www.cyziv.com. Last accessed December 21, 2005.
- [25] M. Czerwinski, G. Smith, T. Regan, and et al. Toward characterizing the productivity benefits of very large displays. In *Interact*, 2003.
- [26] M. Czerwinski, D. Tan, and G. Robertson. Women take a wider view. In *CHI*, 2002.
- [27] John deere engineering. www.deere.com/en_US/deerecom/india/engineering.html. Last accessed December 15, 2005.
- [28] P. Dietz, R. Raskar, S. Booth, and et al. Multi-projectors and implicit interaction in persuasive public displays. In *Working Conference on Advanced Visual Interfaces*, 2004.
- [29] Distributed multihead x project. URL: <http://dmx.sourceforge.net>. Last accessed December 20, 2005.
- [30] S. Elrod, R. Bruce, R. Gold, and et al. Liveboard: A large interactive display supporting group meetings, presentation and remote collaboration. In *CHI*, 1992.
- [31] J. Engelhardt. The 3d future is now. *Geospatial Solutions*, Oct. 2004.
- [32] D. Evans, R. Moorhead, S. Roberts, and et al. Forest visualization through development of prototypical models of stands from lidar and multi-spectral data. In *Geospatial Information in Agriculture and Forestry*, 2001.
- [33] Fakespace systems. www.fakespace.com. Last accessed December 21, 2005.
- [34] Fakespace delivers world's highest resolution visualization room to los alamos national laboratory (article dated 11/08/2005). www.fakespace.com/pressNewsReleases.htm. Last accessed January 3, 2006.
- [35] G. Fitzmaurice, A. Khan, G. Kurtenbach, and et al. Cinematic meeting facilities using large displays. *IEEE CG&A*, 25(4), 2005.
- [36] B. Fogg. *Persuasive Technology: Using Computers to Change What we Think and Do*. Morgan Kaufmann, 2002.
- [37] T. Funkhouser and K. Li. Large format displays. *IEEE CG&A*, 20(4), 2000.
- [38] J. Geissler. Gedrics: the next generation of icons. In *Interact*, 1995.
- [39] J. Geissler. Shuffle, throw or take it! working efficiently with an interactive wall. In *CHI-Summary*, 1998.
- [40] D. Germans, H. Spoelder, L. Renamot, and et al. VIRPI: A high-level toolkit for interactive scientific visualization in virtual reality. In *IPT/EGVE*, 2001.
- [41] M. Gross, E. Lamboray, M. Näf, and et al. blue-c: A spatially immersive display and 3d video portal for telepresence. *ACM TOG*, 22(3), 2003.
- [42] T. Grossman, R. Balakrishnan, G. Kurtenbach, and et al. Interaction techniques for 3d modeling on large displays. In *Interactive 3D Graphics*, 2001.
- [43] J. Grudin. Groupware and social dynamics: Eight challenges for developers. *Communications of the ACM*, 37(1), 1994.
- [44] F. Guimbretière, M. Stone, and et al. Fluid interaction with high-resolution wall-size displays. In *UIST*, 2001.
- [45] M. Hachet and P. Guitton. The interaction table – a new input device designed for interaction in immersive large display environments. In *EGVE*, 2002.
- [46] N. Hashimoto, S. Jeong, T. Takeyama, and et al. Immersive multi-projector display on hybrid screens with human-scale haptic and locomotion interfaces. In *CyberWorlds*, 2004.
- [47] M. Hereld, I. R. Judson, and R. Stevens. Dotytoto: A measurement engine for aligning multi-projector display systems. In *ANL/MCS-P958-0502*, 2002.
- [48] M. Hereld, I. R. Judson, and R. L. Stevens. Introduction to building projection-based tiled display systems. *IEEE CG&A*, 20, 2000.
- [49] E. Huang. Design and analysis of groupware for large displays. In *CHI*, 2005.
- [50] G. Humphreys, M. Houston, Y. Ng, and et al. Chromium: A stream processing framework for interactive rendering on clusters. *ACM TOG*, 21(3), 2002.
- [51] D. Hutchings, M. Czerwinski, B. Meyers, and et al. Exploring the use and affordances of multiple display environments. In *Workshop on Ubiquitous Display Environments at UbiComp*, 2004.
- [52] N. Interaction. www.naturalinteraction.org. Last accessed December 21, 2005.
- [53] H. Ishii, M. Kobayashi, and K. Arita. Iterative design of seamless collaboration media. *CACM*, 37(8), 1994.
- [54] S. Izadi, H. Brignull, T. Rodden, and et al. Dynamo: A public interactive surface supporting the cooperative sharing and exchange of media. In *UIST*, 2003.
- [55] P. Jedrysik, J. Moore, L. Brykowsky, and et al. Interactive displays for command and control. In *Aerospace Conference*, 2000.
- [56] B. Jeong, R. Jagodic, L. Renamot, and et al. Scalable graphics architecture for high-resolution displays. In *IEEE Information Visualization Workshop*, 2005.
- [57] B. Johanson, A. Fox, and T. Winograd. The interactive workspaces project: Experiences with ubiquitous computing rooms. *IEEE Pervasive Computing*, 1(2), 2002.
- [58] A. Khan, G. Fitzmaurice, D. Almeida, and et al. A remote control interface for large displays. In *UIST*, 2004.
- [59] A. Khan, J. Matejka, G. Fitzmaurice, and et al. Spotlight: Directing users attention on large displays. In *CHI*, 2005.
- [60] H. Kim and D. W. Fellner. Interaction with hand gesture for a back-projection wall. In *Computer Graphics International*, 2004.
- [61] N. Krishnaprasad, V. Vishwanath, S. Venkataraman, and et al. Juxtaview: A tool for interactive visualization of large imagery on scalable tiled displays. In *IEEE Cluster*, 2004.
- [62] C. Krumbholz, J. Leigh, A. Johnson, and et al. Lambda table: High resolution tiled display table for interacting with large visualizations paper. In *Workshop on Advanced Collaborative Environments*, 2005.
- [63] G. Kurtenbach and G. Fitzmaurice. Applications of large displays. *IEEE CG&A*, 2005.
- [64] E. Lamboray, A. Zollinger, O. G. Staadt, and et al. Interactive multimedia streams in distributed applications. *Computers and Graphics*, 27(5), 2003.
- [65] J. LaViola, D. Keefe, R. Zeleznik, and et al. Case studies in building custom input devices for virtual environment interaction. In *VR 2004 Workshop: Beyond Glove and Wand Based Interaction*, 2004.
- [66] J. Leigh, A. Johnson, P. Hamilton, and et al. Geowall2. www.evl.uic.edu/index2.php. Last accessed December 23, 2005.
- [67] H. Liao, M. Iwahara, N. Hata, and et al. High-resolution integral videography autostereoscopic display using multi-projector. In *International Display Workshop*, 2002.
- [68] J. Lin, H. Duh, D. Parker, and et al. Effects of view on presence enjoyment, memory, and simulator sickness in a virtual environment. In *IEEE Virtual Reality*, 2002.
- [69] A. MacEachren, R. Edsall, D. Haug, and et al. Virtual environments for geographic visualization: Potential and challenges. In *NPIVM*, 1999.
- [70] J. Mackinlay and J. Heer. Wideband displays: Mitigating multiple monitor seams. In *CHI*, 2004.
- [71] A. Majumder and R. Stevens. Lam: Luminance attenuation map for photometric uniformity in projection based displays. In *VRST*, 2002.
- [72] S. Malik, A. Ranjan, and R. Balakrishnan. Interacting with large displays from a distance with vision-tracked multi-finger gestural input. In *UIST*, 2005.
- [73] P. Mantey and J. Richards. Using large shared displays to create a

- collaborative classroom learning environment. In *WACE*, 2005.
- [74] S. Molnar, M. Cox, D. Ellsworth, and et al. A sorting classification of parallel rendering. *IEEE CG&A*, 14(4), 1994.
- [75] B. Myers. Using multiple devices simultaneously for display and control. *IEEE Personal Communications*, 7(5), 2000.
- [76] M. Näf, O. Staadt, and M. Gross. blue-c api: A multimedia and 3d video enhanced toolkit for collaborative vr and telepresence. In *VR-CAI*, 2004.
- [77] M. Näf, O. G. Staadt, and M. Gross. Multimedia integration into the blue-c api. *Computers and Graphics*, 29(1), 2005.
- [78] M. Näf, O. G. Staadt, and M. H. Gross. The blue-c distributed scene graph. In *IPTEGVE*, 2003.
- [79] NCSA pixel blaster. brighton.ncsa.uiuc.edu/prajlich/wall/ppb.html. Last accessed December 21, 2005.
- [80] NCSA terraserver blaster. brighton.ncsa.uiuc.edu/prajlich/wall/tsb.html. Last accessed December 21, 2005.
- [81] Oak ridge national laboratory review. www.ornl.gov/info/ornlreview/v37_2_04/article14.shtml. Last accessed December 23, 2005.
- [82] D. Olsen and T. Nielsen. Laser pointer interaction. In *CHI*, 2001.
- [83] R. Otsuka, T. Hoshino, and Y. Horry. Transpost: All-around display system for 3d solid modeling. In *VRST*, 2005.
- [84] E. Patrick, D. Gosgrove, A. Slavkovic, and et al. Using a large projection screen as an alternative to head-mounted displays for virtual environments. In *CHI*, 2000.
- [85] K. A. Perrine and D. R. Jones. Parallel graphics and interactivity with the scalable graphics engine. In *Supercomputing*, 2001.
- [86] N. Polys, S. Kim, and D. Bowman. Effects of information layout, screen size, and field of view on user performance in information-rich virtual environments. In *VRST*, 2005.
- [87] L. L. N. L. V. Project. www.llnl.gov/icc/sdd/img/displaywall.shtml. Last accessed January 5, 2006.
- [88] B. Raffin and L. Soares. Pc clusters for virtual reality. In *IEEE Virtual Reality*, 2006.
- [89] D. Raja, D. Bowman, J. Lucas, and et al. Exploring the benefits of immersion in abstract information visualization. In *IPTE*, 2004.
- [90] R. Raskar, J. van Baar, P. Beardsley, and et al. ilamps: Geometrically aware and self-configuring projectors. *ACM TOG*, 22(3), 2003.
- [91] R. Raskar, G. Welch, M. Cutts, and et al. The office of the future: a unified approach to image-based modeling and spatially immersive displays. In *SIGGRAPH*, 1998.
- [92] J. Rekimoto. A multiple-device approach for supporting whiteboard-based interactions. In *CHI*, 1998.
- [93] L. Renambot. Lambdatable. www.evl.uic.edu/cavern/sage/gallery/lambdatable09.jpg. Last accessed December 21, 2005.
- [94] L. Renambot. Lambdavisioin 100mpixel display at uic. www.evl.uic.edu/cavern/sage/gallery/PICT0032.jpg. Last accessed December 21, 2005.
- [95] L. Renambot, A. Johnson, and J. Leigh. Lambdavisioin: Building a 100 megapixel display. www.evl.uic.edu/cavern/sage/pubs/LambdaVision-light.pdf. Last accessed December 18, 2005.
- [96] L. Renambot, A. Rao, R. Singh, and et al. Sage: the scalable adaptive graphics environment. In *WACE*, 2004.
- [97] M. Ringel, H. Berg, Y. Jin, and et al. Barehands: Implement-free interaction with a wall-mounted display. In *CHI Extended Abstracts*, 2001.
- [98] G. Robertson, M. Czerwinski, P. Baudisch, and et al. The large-display user experience. *IEEE CG&A*, 25(4), 2005.
- [99] M. Robinson, J. Laurence, A. Hogue, and et al. Ivy: Basic design and contruction details. In *ICAT*, 2002.
- [100] J. Rohlf and J. Helman. Iris performer: a high performance multiprocesing toolkit for real-time 3d graphics. In *SIGGRAPH*, 1994.
- [101] A. Sadagic, H. Towles, J. Lanier, and et al. National tele-immersion initiative: Towards compelling tele-immersive collaborative environments, 2001.
- [102] D. Sandin, T. Margolis, J. Ge, and et al. The varrier autostereoscopic virtual reality display. In *SIGGRAPH*, 2005.
- [103] T. Sandstrom, C. Henze, and C. Levit. The hyperwall. In *CMV*, 2003.
- [104] D. Schikore, R. Fischer, R. Frank, and et al. High-resolution multi-projector display walls. *IEEE CG&A*, 20(4), 2000.
- [105] G. Schmidt, S. Chen, A. Bryden, and et al. Multi-dimensional visual representations for underwater environmental uncertainty. *IEEE CG&A*, 24(5), 2004.
- [106] N. Schwarz, S. Venkataraman, L. Renambot, and et al. Vol-a-tilde: A tool for interactive exploration of large volumetric data on scalable tiled displays. In *IEEE Visualization*, 2004.
- [107] R. Singh, J. Byungil, L. Renambot, and et al. TeraVision: A distributed, scalable, high resolution graphics streaming system. In *IEEE Cluster Computer*, 2004.
- [108] Sony develops new 4k sxrd high-resolution liquid crystal display panel (article dated 05/18/2004). www.sony.net/SonyInfo/News/Press_Archive/200405/04-027E. Last accessed January 3, 2006.
- [109] G. Starkweather. Dsharp: A wide-screen multi-projector display. *Journal of optics: pure and applied optics*, 5, 2003.
- [110] M. C. Stone. Color and brightness appearance issues in tiled displays. *IEEE CG&A*, 21(5), 2001.
- [111] N. Streitz, J. Geissler, T. Holmer, and et al. i-land: An interactive landscape for creativitiy and innovation. In *CHI*, 1999.
- [112] A. Sullivan. 3 deep. *IEEE Spectrum*, 42(4), 2005.
- [113] N. Swaminathan and S. Sato. Interaction design for large displays. *Interactions*, 4(1), 1997.
- [114] D. Tan and M. Czerwinski. Effects of visual separation and physical discontinuities when distributing information across multiple displays. In *Australian CHI*, 2003.
- [115] D. Tan, M. Czerwinski, and G. Robertson. Women go with the (optical) flow. In *CHI*, 2003.
- [116] D. Tan, D. Gergle, P. Scupelli, and et al. With similar visual angles, large displays improve spatial performance. In *CHI*, 2003.
- [117] D. Tan, D. Gergle, P. Scupelli, and et al. Physically large displays improve path integration in 3d virtual navigation tasks. In *CHI*, 2004.
- [118] D. Tan, J. K. Stefanucci, D. Proffitt, and et al. The infocockpit: providing location and place to aid human memory. In *PUI*, 2001.
- [119] D. Tate, L. Sibert, and T. King. Using virtual environments for fire-fighter training. *IEEE CG&A*, 17(6), 1997.
- [120] F. Tyndiuk, G. Thomas, V. Lespinet-Najib, and et al. Cognitive comparison of 3d interaction in front of large vs. small displays. In *Virtual Reality Software and Technology*, 2005.
- [121] I. V. UC Davis and G. R. Group. graphics.idav.ucdavis.edu/research. Last accessed January 5, 2006.
- [122] T. van der Schaaf, L. Renambot, D. Germans, and et al. Retained mode parallel rendering for scalable tiled displays. In *IPTE*, 2002.
- [123] ModViz Renderizer. www.modviz.com/products/. Last accessed December 22, 2005.
- [124] Visbox innovative display and interaction technologies. visbox.com. Last accessed December 18, 2005.
- [125] VideoLAN VLC media player. URL: <http://www.videolan.org/vlc/>. Last accessed December 22, 2005.
- [126] D. Vogel and R. Balakrishnan. Distant freehand pointing and clicking on very large, high resolution displays. In *UIST*, 2005.
- [127] G. Voss, J. Behr, D. Reiners, and et al. A multi-thread safe foundation for scene graphs and its extension to clusters. In *EGPGV*, 2002.
- [128] Iowa state university, virtual reality applications center. www.vrac.iastate.edu. Last accessed December 15, 2005.
- [129] G. Wallace, H. Chen, and K. Li. Color gamut matching for tiled display walls. In *EGVE*, 2003.
- [130] B. Wei, C. Silva, E. Koutsofios, and et al. Visualization research with large displays. *IEEE CG&A*, 20(4), 2000.
- [131] M. Weiser. The computer for the 21st century. *Scientific American*, 256(3), 1999.
- [132] C. Wisneski, H. Ishii, H. Dahley, and et al. Ambient displays: Turning architectural space into an interface between people and digital information. In *CoBuild*, 1998.
- [133] Xenovision dynamic matrix display. www.xenotran.com/products.htm. Last accessed December 23, 2005.
- [134] T. Yamada, M. Hirose, and Y. Isda. Development of complete immersive display: Cosmos. In *VSM*, 1998.



(C1)



(C2)



(C3)

Figure C1: Los Alamos National Laboratory's ultra high-resolution immersive room—La Cueva Grande [34]—developed by Fakespace Systems [33] (Image courtesy of Los Alamos National Laboratory, reference LA-UR-05-7498. Photographed by Presley Salaz, LANL).

Figure C2: Tiled-LCD multi-monitor desktop display (Image courtesy of Ball [4]).

Figure C3: LambdaVision, a 100-MPixel wall (Image provided by Luc Renambot, Electronic Visualization Laboratory, University of Illinois at Chicago [94, 93]).



(C4)



(C5)



(C6)

Figure C4: LambdaTable, a horizontal tiled display and interface (Image provided by Luc Renambot, EVL-UIC [94, 93]).

Figure C5: Stackable, reconfigurable projector array, VisBlock™ from VisBox Innovative Display and Interaction Technologies (Image courtesy VisBox.com [124]).


Figure C6: Cyviz's stereoscopic projector array called Cyviz Vizwall™ (Image courtesy Cyviz.com [24]).



(C7)



(C8)



(C9)

Figure C7: EVL-UIC's autostereoscopic display called Varrier™ (Image courtesy EVL-UIC [102]).

Figure C8: Cutaway mockup of DepthCube™ display from LightSpace Technologies (Image courtesy LightSpace Technologies [112]).

Figure C9: NRL's immersive room demonstrating a submarine command & control application.



(C10)



(C11)

Figure C10: Deere & Company's VR testing facility for virtual prototyping vehicle designs and testing their effectiveness (Photo courtesy Deere & Company, Moline, Illinois [27]).

Figure C11: State-of-the-art meeting facility at Alias Visualization Studio (Image courtesy of Alias [2]).