

Scalable Resolution Display Walls

By JASON LEIGH, ANDREW JOHNSON, LUC RENAMBOT, TOM PETERKA, BYUNGIL JEONG,
DANIEL J. SANDIN, JONAS TALANDIS, RATKO JAGODIC, SUNGWON NAM,
HYEJUNG HUR, AND YIWEN SUN

ABSTRACT | This article will describe the progress since 2000 on research and development in 2-D and 3-D scalable resolution display walls that are built from tiling individual lower resolution flat panel displays. The article will describe approaches and trends in display hardware construction, middleware architecture, and user-interaction design. The article will also highlight examples of use cases and the benefits the technology has brought to their respective disciplines.

KEYWORDS | Computer graphics; stereo vision; virtual reality; visualization

I. INTRODUCTION

Global cyberinfrastructure, whether sensornets, supercomputers, or the Cloud, is transforming the way teams of scientists and engineers study and understand complex systems—whether physical, geological, biological, environmental, or atmospheric—from the micro- to the macroscale, in both time and space. These systems also produce greater volumes of data than ever before, which need to be correlated and interpreted in order for researchers to gain insight and knowledge. This has created

a major challenge for researchers who must now learn to manage the increased scale and complexity of their work.

Visualization is still the most effective means for researchers to gain insight—nearly one-third of the human brain is devoted to processing visual information, and 3-D visual cues, or *stereoscopy*, have been shown to be beneficial for disambiguating multidimensional information [27], [38]. In the scientific discovery process, visualization serves three important roles: it can help to quickly verify the correctness of a complex simulation model during initial development; it can make simulation results more immediately available when tightly integrated with the model; and it can help make solutions more easily understood by lay audiences, whether public policy decision makers or the general public.

In 2000, the Electronic Visualization Laboratory (EVL) began development of the Continuum (Fig. 1) to understand how people might be able to mitigate the problems of scale and complexity in data by working in project rooms where traditional paper-based media could be replaced with digital media [22]. Paper easels, poster boards and whiteboards were replaced with arrays of 50-in plasma displays equipped with a touch screen overlay, a 3-D GeoWall projection display, a 2×2 tiling of 20-in LCDs, and numerous wireless tablet computers. Users could write directly on a digital whiteboard, they could manipulate the documents with a wireless tablet computer, and they could communicate with distant colleagues using a multisite video conferencing system. The GeoWall system enabled them to view scientific data in 3-D while a tiling of 2×2 LCDs enabled them to view 2-D data such as high-resolution maps.

The target audiences for the Continuum were scientists and engineers who routinely had a need to collaborate with

Manuscript received July 15, 2011; revised November 9, 2011; accepted February 20, 2012. This work was supported in part by the National Science Foundation under Awards CNS-0935919 and 0959053, and by the King Abdullah University of Science and Technology (KAUST) under Award US-2008-107/SA-C0064. Work at the Center for Nanoscale Materials was supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract DE-AC02-06CH11357.

J. Leigh, A. Johnson, L. Renambot, D. J. Sandin, J. Talandis, R. Jagodic, S. Nam, H. Hur, and Y. Sun are with the Electronic Visualization Laboratory, University of Illinois at Chicago, Chicago, IL 60607 USA (e-mail: spiff@uic.edu).

T. Peterka is with the Argonne National Laboratory, Argonne, IL 60439 USA.

B. Jeong is with the Texas Advanced Computing Center, Austin, TX 78758 USA.

Digital Object Identifier: 10.1109/JPROC.2012.2191609



Fig. 1. *The Continuum—a project room at the Technology Research Education and Commercialization Center (TRECC) in Illinois (2000).*

one another to interpret large quantities of disparate types of data. This kind of workspace leverages the traditional benefits of war rooms, which have been shown to increase productivity by as much as twofold [35], and amplifies them with interactive computing connected to high-speed networks.

Based on this research EVL predicted that, as technologies become more affordable, seamless and infinitely tileable displays will be used in a broad range of areas including marketing and finance, entertainment, and education, to name a few. This paper will discuss some of the challenging technical issues of scalable resolution display walls (SRDWs). The paper will begin with a discussion of the computing and networking infrastructure needed to construct walls of nearly arbitrary size and resolution. Then it will describe the software systems needed to build applications for the wall. The paper will also describe approaches for interacting with SRDWs and also highlight case studies of how a variety of user communities have used them. Lastly this paper will discuss issues that arise when attempting to enable seamless stereoscopic 3-D on these walls.

II. HARDWARE

In 2000, a special issue of the IEEE COMPUTER GRAPHICS AND APPLICATIONS [4] called “...Onto the wall—Large displays” covered the early work in this space. Editors Thomas Funkhouser and Kai Li selected papers that mostly focused on how to build large-display systems using sets of screens, projectors, cluster of computers, and graphics software to drive them. Most of the papers addressed computer system design, calibration and blending of projectors, human–computer interaction, remote visualization, and content design. A few early adopters of these technologies generated use case studies, in the car industry and the “office of the future” initiative for example.

Five years later, Gordon Kurtenbach and George Fitzmaurice edited an issue of the magazine on the same topic [5]. They wrote: “the emergence of large displays holds the promise of basking us in rich and dynamic visual

landscapes of information, art, and entertainment.” The technologies described in the papers in this issue were still mostly tiling of projectors. However, a significant shift in the computer industry occurred: the emergence of powerful gaming graphics cards and PC clusters. This was a shift from specialized high-end computers (for computation and graphics) to off-the-shelf components driven by the consumer market. A single PC today can drive multiple displays using graphics processing units (GPUs) with several video outputs and/or several GPUs at the same time. Standard applications can now run unmodified on such systems with high-resolution. At the same time, HDTV has become widely available in the consumer market with high-definition (HD) plasma and LCD TVs. In the professional market 4-Mpixel LCD panels have been introduced (Apple 30-in LCD at 2560×1600 pixel resolution) and 4-K projectors have become available (Sony SXR4 4-K 4096×2160 pixel resolution). The introduction of these new technologies requires the reexamining of past assumptions about SRDW hardware design.

Graphical computer systems designed to drive large high-resolution displays have significantly evolved over the years, as a result of converging trends from the high-performance computing community and the consumer gaming market.

A. Computer Systems

Initial systems consisted of a large graphics supercomputer (often manufactured by SGI) with custom graphics cards driving cathode ray tube projectors. The CAVE (4-wall stereoscopic room) (Fig. 2) [7] and large visualization walls (often three to four projector systems) were designed this way: for instance, the CAVE first unveiled at the SIGGRAPH’92 conference consisted of four Electrohome video projectors each driven by a Silicon Graphics Onyx each with a Reality Engine graphics system. The advent of powerful personal computers (sustained by Moore’s law’s predicted increase in power and density) and the development of Beowulf clusters and Grid computing software packages offered new opportunities. Affordable PC clusters could be equipped with powerful



Fig. 2. The CAVE virtual reality environment (1992).

consumer graphics cards from the gaming market. This led to the design of large projection systems with commodity video projectors like the “Scalable Display Wall” at Princeton University and the “Office of the Future” at the University of North Carolina at Chapel Hill. As LCD panels also became commodity devices, they replaced the projectors. Though LCDs possessed large borders, they were overall easier to maintain than projectors.

A complementary trend is the component parallelization inside today’s computer: while the frequency race has slowed down in the industry, CPUs have become increasingly parallel utilizing multicore technology. A 2011 workstation has as many as 12 cores (or 24 hardware threads), with many more planned in the manufacturers roadmaps. Similarly, GPUs are designed as a parallel computing architecture with hundreds of processing units, and fast PCI-express communication lets system integrators build systems with multiple GPUs and network interfaces. Large memory capacity is also available with increasing chip density. Overall, yesterday’s cluster is

becoming today’s workstation, with tens of concurrent tasks accessing large amounts of memory, driving several GPUs and accessing high-bandwidth networks. It is therefore now possible, and more practical, to build a large high-resolution wall driven by a single computer. One such example is the “Cyber Commons” wall at EVL (see Fig. 3), an 18-Mpixel display designed for science and education built around 18 near-seamless LCD panels driven by one single PC (three GPUs output six dual-link DVI signals which are multiplexed over eighteen single-link DVI cables using six Matrox TripleHead2Go multi-display adapters).

B. Networking

While initial systems in the 1990s used shared-memory within a graphics supercomputer to drive large displays, networking technologies became essential to manage and control a cluster of PCs. Slow networks by today’s standards, 100-Mb LANs, were initially used to synchronize systems and share user interaction events with the rendering nodes. Gigabit Ethernet networks and then specialized networks (Myrinet, etc.) provided enough bandwidth (and in some cases lower latency) to share more data and change the programming models of these displays. Infiniband, popular in the high performance computing community, with its high-bandwidth and low-latency proved useful: high-bandwidth meant convenient access to large stored data sets to be rendered, while low-latency was useful for user interaction and synchronization. However, in 2003, technology specialists noticed a change in trends concerning optical networking and envisioned a world where wide-area bandwidth would become almost infinite (relative to local bandwidth) and would surpass the capacity of any single endpoint, with exponential capacity growth in the core of the network (with DWDM optical network technology) and falling price of 10-Gb/s network interface (at the endpoints). This




Fig. 3. Cyber-commons at the Electronic Visualization Laboratory—an 18-Mpixel scalable resolution display wall (2009).

is the premise of the National Science Foundation's OptIPuter project [33]. OptIPuter's goal has been to enable researchers to easily interact with large-scale data objects and communicate via HD video with colleagues, whether local or remote, as easily as the Web makes it easy to access small data. The OptIPuter network was thought of as an optical overlay to the standard shared Internet, so that individual researchers could tightly couple computational resources over dedicated multigigabit optical networks. Such data superhighways are needed because the shared Internet and the World Wide Web are engineered to interactively handle megabyte-sized objects, whereas today's scientific instruments generate gigabyte- to terabyte-sized data sets [33]. One outcome of the OptIPuter project was to scale up an end-user device from a single PC, appropriate for the shared internet, to a SRDW driven by low-cost commodity PC clusters and connected to an optical network (termed an OptIPortal)—to provide a greater pixel display area, storage, compute power, and I/O bandwidth, while maintaining personal interactivity (Fig. 4) [33]. To manage the information on OptIPortals, visualization middleware called the scalable adaptive graphics environment (SAGE) [19], was developed. SAGE is essentially an operating system that lets users launch distributed visualization applications on remote clusters whose outputs are then streamed directly to OptIPortals of variable size, where they can be viewed and manipulated [29].

Not only reserved to the scientists, this model is available today to everyone through the use of "cloud computing" where services are hosted by providers (such as Amazon, Google, Microsoft, etc.) on distributed data centers located around the globe. When interacting with colleagues and manipulating data sets, users leverage high-performance services (computation, storage, networking, rendering) while receiving only processed results (documents, graphs, images, videos) over low-bandwidth channels to his/her endpoints. Improvement in networking bandwidth to the endpoints will improve the quality of the user experience, from mobile devices to the largest high-resolution wall display.

C. The Compute Node

After the era of the graphics supercomputer (mainly shared-memory SGI machines), intense developments were focused on PC clusters equipped with high-frequency processors. This led the processor industry to a "power wall" where a single 3-GHz chip could not be air-cooled further with powerful fans. Simpler and smarter chips were designed which opened two avenues: on the one hand building systems with a very large number of nodes using simpler chips (like IBM with Blue-Gene) supported by very fast interconnection networks, and on the other hand, building systems with fewer nodes but each one including multiple simpler processors on a single chip (like Intel with multicore processors or Nvidia with hundreds of simple graphics processor cores onto a single chip). While



Fig. 4. The LambdaVision—a 100-Mpixel LCD OptIPortal was built from 55 20-in LCD panels and driven by a cluster of 30 computers (2005).

both approaches are explored by scientists to provide the highest-end systems in the world (toward peta- and exa-scale computing), the concentration of resources inside a single box has been very useful when designing and building a large display wall. Today, a single PC node can provide a sizeable amount of computing resources (multi-cpu multicore systems with hardware threading support), a combination of fast and large amount of storage (large spinning hard drives and fast SSDs devices), several high-end GPUs, and several multi-10-Gb/s network interfaces (QDR Infiniband or dual-10-G Ethernet). A growing trend is also to use the GPUs for both computation and graphics, using a hybrid programming approach leveraging both multicore processors and GPUs appropriately. The parallelism of such systems (many CPU cores, many GPU cores, deep memory hierarchies, several network interfaces, several level of storage) results in complex resource allocation problems for the programmer [36]. However, it presents many benefits for the end user, in this case a domain scientist, having to manage only a single computer (or just a few) instead of a large cluster.

D. The Displays

Initial wall displays were built around expensive and bulky CRT projectors. While low intensity, they were relatively high resolution and also highly configurable. Driven by a growing market for projectors in boardrooms to show "powerpoint" presentations, CRT projectors were eventually replaced by projectors that use LCD and DLP technology, which were lower in cost and smaller in form factor. These projectors presented, however, many challenges when used in groups (heat, maintenance, color accuracy, uniformity, control, high-contrast, etc.). Numerous efforts were made to use these inexpensive projectors. Although the initial cost of installation was relatively low, the high maintenance cost made them impractical for long-term use on large-scale walls.

The innovation in LCD technology expanded the design space, especially with desktop LCD panels (or consumer LCD HDTVs). LCD panels work reliably for tens of thousands of hours, are easily aligned in terms of color and geometry, and are easy to mount into tiled arrays. The only disadvantage was the thick borders/mullions around them. Early adopters found them to be invaluable for viewing high-resolution imagery such as satellite images, but found them problematic for displaying textual information as the thick borders tended to occlude text in ways that made it difficult to read. What was needed were flat panel displays with no borders or very small borders. To satisfy this demand display manufacturers (e.g., NEC, Sharp, Samsung, etc.) finally introduced in 2009 LCD panels with borders on the order of 2.5–3.5 mm per side, often referred to as seamless or near-seamless displays. While still expensive (from \$5000 to \$10 000 a panel), they provided the first opportunity to build SRDWs with nonintrusive borders. The first generation of these displays had a resolution of 1366×768 per panel. The next generation provided full 1920×1080 HD resolution.

III. MIDDLEWARE

In order to drive a collection of displays as if it were one unified display some form of middleware (either implemented as hardware or software) is needed to interface between the user's application and the graphics hardware that drives the displays. While an application could conceivably drive the graphics hardware directly, the use of middleware has the advantage that it enables a broader range of applications to operate on displays regardless of resolution or tiling configuration.

One can classify tiled display middleware as window managers such as Distributed Multihead X (DMX) [10] and SAGE [20], or as parallel graphics rendering frameworks such as Chromium [16], Equalizer [12], Cross-Platform Cluster Graphics Library (CGLX) [11], and Deep Computing Visualization (DCV) [17]. Tiled display window managers enable users to run and interact with multiple applications at the same time, as in a typical multitasking operating system, whereas parallel rendering frameworks are designed to enable an application to maximize the use of a display wall's resolution.

DMX provides a single unified X Window desktop by coupling multiple X servers on multiple machines. In DMX a master node distributes X Window primitives to other client nodes which render the graphics. This allows users to view and control general X applications on a tiled display in the same way as they do on a desktop computer environment. In contrast SAGE routes streams of pixels that may come from multiple disparate and concurrently running remote applications. SAGE provides a user-interface that enables multiple users to interact with the windows on the wall simultaneously, which makes better use of a large scale tiled display.

In Chromium, a server intercepts OpenGL command streams from an OpenGL application and routes them to the intended display client nodes for rendering. IBM's DCV is a proprietary solution that has a similar set of functionalities to Chromium. Equalizer and CGLX are similar in that they run a copy or part of the application codes on the display client nodes, but Equalizer offers a more scalable and flexible approach in which a user can combine various parallel rendering algorithms, such as sort-first or sort-last [26], and support multiscreen configurations such as CAVEs or display walls. ICE-T is a sort-last parallel rendering framework that can redistribute the composited image fragments for a tiled display.

The main advantage of Chromium and DMX over other approaches is that they enable conventional desktop graphics applications to run on a tiled display without modifying the underlying application code. However, this application transparency only extends as far as supporting OpenGL (Chromium) and X11 (DMX) applications. On the other hand, SAGE can support a wider range of applications with minimal code change, because it fetches and streams pixels from applications. Furthermore, SAGE permits multiple applications to appear on the tiled display wall at the same time. Equalizer also requires some application code modification. In its simplest mode Equalizer runs an Equalizer-enabled OpenGL application executable on each client node in a manner similar to CGLX. ICE-T provides API calls to perform sort-last image compositing and image data redistribution for a tiled display. ParaView and VisIt [37], two widely used parallel visualization tools, use ICE-T for image compositing.

Both Chromium and DMX have limited scalability because in each its single data source (server) can become a bottleneck as the number of display client nodes increases. Equalizer and CGLX show better scalability than Chromium because each rendering client directly reads data from the data source (e.g., the file system) in parallel. Chromium also provides a parallel application interface for sort-last rendering, but it loses its main advantage, application transparency, in this mode. Chromium generally shows good performance with fill-rate limited applications running in retained mode (i.e., graphics data are streamed once in the beginning and follow-on streams are mostly control messages). However, even in the best case, Chromium is outperformed by Equalizer as the number of rendering clients increases [12]. While the maximum tiled display dimension DMX supports is 4×4 (16 display client nodes), SAGE is designed to support much larger tiled displays (e.g., 15×5 tiles driven by 27 display client nodes). Because SAGE can route pixel streams from multiple application nodes to tiled display nodes, it does not have the single data source bottleneck.

SAGE's uniqueness lies in that it supports distant collaboration among multiple endpoints equipped with tiled displays connected by high-speed networks. It enables users to share their scientific visualization at an

extremely high-resolution with remote collaborators while communicating with each other via multipoint HD video and audio streaming to their tiled displays. Chromium has been integrated with VNC [30] to support remote visualization and collaboration [1] and DCV also supports remote visualization streaming using VNC. However, in their approaches, a remote client can only be single desktop computer rather than a tiled display.

Historically, computer clusters were used to drive scalable tiled display walls. With the advent of “multiheaded” graphics cards such as AMD’s Eyefinity, NVIDIA’s Scalable Visualization Solutions, and Matrox’s Display Wall Controller, it is possible to support over a dozen displays from a single PC. Driving a tiled display with a single computer greatly reduces the cost of ownership and maintenance of a tiled display wall. Furthermore, it enables applications to run natively, thereby eliminating the need to parallelize them. On the other hand driving a high-resolution wall with a single computer introduces significant challenges in resource management. In particular, the operating system must be able to intelligently align the cores of a CPU, the available network capacity, bus bandwidth, and GPU resources to respond to the real-time interactions of multiple users attempting to manipulate content on the wall at the same time. The prioritization of system resources would not necessarily be based on fastest job completion time such as in prior approaches [13], [14], [40], but would also consider which windows are the largest on the wall, which ones are least occluded by other windows, and which ones are being interacted with by a user. Work to develop such a multiuser interaction-based resource manager is currently underway in the design of the next generation of SAGE, called *SAGE-Next*. In order to adequately control systems resources, *SAGE-Next* uses the concept of Rails [36], which enables the alignment of specific cores and memory caches of a CPU, system interrupts, and the network interfaces, on an application basis—in essence providing them with the notion of “quality of service.”

IV. USER INTERACTION

The goal of constructing rooms equipped with SRDWs is to create information spaces that cover our field of view, allowing multiple remote or collocated people to work together to create a shared representation made up of a variety of interrelated digital artifacts that are easily created, manipulated, and persistently stored. SRDWs empower users to better cope with information intensive tasks because they enable more information to be juxtaposed simultaneously. They also have the potential to promote more natural physical navigation of the content, and improve spatial performance when users are analyzing high-resolution data sets. But perhaps most importantly, the intersection of large size and high resolution makes them excellent collaborative workspaces, not too dissimilar from traditional paper-based war rooms or project

rooms that have been found to provide as much as a twofold improvement in the productivity of its users [35].

However, the cited potential benefits of the SRDWs can only be achieved through the design of user-interfaces that are informed by the way people work in environments surrounded with content. Prior research has shown that previously well-established desktop interaction techniques generally do not scale to larger displays [2], [8], motivating the need for new approaches. For large display walls, prior efforts were directed at addressing the issues of target acquisition, cursor tracking, gesturing and pointing, mostly with single users in mind. However, given that vast quantities of content can potentially be posted on the display walls, there is a need for approaches to help multiple users, especially those working in teams, to easily and concurrently manipulate and organize the content. In this regard, the user-interface design in SAGE [18] embodies the most comprehensive understanding of the problem to date.

SAGE makes the following assumptions about the way users work in front of SRDWs.

- **Distributed**—User interfaces can be automatically displayed and interacted with across any number of displays, driven by any number of machines.
- **Scalable**—As the size and resolution of scalable displays varies significantly, it is imperative that the interaction system adapts to the target display size and resolution from the perspective of visibility and usability.
- **Multiuser**—Generally, empowering every user with control minimizes social barriers for participation by eliminating the need to request control or interfering with other users. This in turn promotes discussion and fosters collaboration. Additionally, individual control facilitates task division among group members, potentially increasing group performance through a greater degree of parallelism.
- **Multimodal**—Given the numerous potential applications of scalable displays, it is unlikely that a single input modality will satisfy interaction requirements of every use case. The choice of interaction modality depends on the number of potential users, the target display size and the mobility of users in front of the display.

A. Interaction Zones

Various devices such as touch input, Gyromouse, trackpad, six degree-of-freedom tracked wand, Nintendo’s Wiimote, Microsoft Kinect, and wireless laptops/tablets, have all been examined as possible candidates for interacting with SRDWs. Observations of approximately 30 users (students, staff and faculty at the University of Illinois) over a period of two years suggest that as the distance between the user and the wall increases, so too does the preference for indirect interaction increase




Fig. 5. Suitable modalities of interaction for four major interaction zones in front of a scalable resolution display wall.

(see Fig. 5). For example, when a user is up-close to the wall, it is most convenient to be able to reach out with one's hand and directly touch the image on the wall, whereas a user who is seated far from the wall would prefer to manipulate the information on the wall using the mouse pointer on their laptop to remotely control a pointer on the wall.

B. Display Space Management

Mostly overlapping windows on traditional desktop displays become mostly visible windows on SRDWs, meaning that the problem of window switching now gives way to the problem of window organization. Compounded by the fact that large screen real estate will inevitably entice users to juxtapose more information, one must examine better ways to help users organize this information into meaningful arrangements. SAGE provides a number of features toward this end. For example, SAGE provides resizable horizontal and vertical partitions for dividing the wall into

smaller workspaces that can be managed independently (see Fig. 6). A tiling feature will automatically keep windows within a specified workspace organized in a logical grid, or an ordering that minimizes windows crossing over physical display mullions. SAGE also enables the current content on the wall to be saved to enable users to continue working at a later time, or to restore the state of the collaboration in the event of a computer or power failure.

C. Supporting Collaborative Work

Work in SRDWs closely mirrors the kinds of activity found in traditional paper-based war rooms and project rooms. Collaboration is amplified due to two factors: users are able to collectively externalize and organize all the relevant information; and users maintain a constant spatial awareness of the information and the experts in the room. Observations of users in SAGE revealed the following.

- Partitioning the display space—Partitioning the display space was frequently employed to designate individual workspaces for each user and to spatially categorize information.
- Individual organization preferences—Users seemed to embrace the ability to organize their own display space without interfering with others, and if we allow them to exercise their individual preferences they are more likely to adopt wall displays as a collaborative medium.
- Sense of ownership—During collaborative work users seem to respect each other's window ownership and personal space, even if not explicitly defined. This suggests that social protocols may be sufficient for coordinating control among groups of users.
- Supporting reorganization—When users are engaged in analysis or sense-making tasks, viewing information arranged in different logical arrangements



Fig. 6. A user partitions on the work space in SAGE by drawing wall dividers with his finger.

can bring about different perspectives—a phenomenon often referred to as the Rashomon effect [15].

- Structure versus freedom—Although automatic layout features are very useful in structuring the environment, users will always want the freedom to fine-tune the layout to exactly fit their needs.
- Organization learning curve—Display space organization became a necessary part of the workflow, however it is evident that using the space effectively is a skill to be learned. Perhaps the system can monitor the state of the display and assist with organization given any metadata that might be available about the content being displayed.
- Offline preparation—Users also have a need to prepare for meetings ahead of time by scripting the presentation from their desktop or laptop computers using a virtual template of the SRDW.

D. Voice Command

A much less examined modality for interacting with SRDWs is through the use of voice commands. Voice commands are most effective when recognition accuracy is high, when they do not require the user to memorize a vocabulary or grammar, and when it takes less time to utter the verbal command than to perform the action through direct manipulation. *Articulate* is a system that attempts to interpret natural language queries to automatically produce appropriate visualizations of data. This is can potentially be very powerful for controlling SRDWs because it can provide a rapid means of creating data plots and other more advanced visualizations without requiring the user to struggle with yet another user-interaction device. The use of a natural language interpreter enables the user to issue commands without conforming to a rigid grammatical structure, and a graph reasoner enables the system to make an educated guess as to the type of visualization that best answers the user's query. The approach deemphasizes the physical manipulations needed to create a desired graph, so that the user can focus on asking questions about the data. Initial studies of a desktop system showed this indeed to be true. Users took less time to accurately produce a desired graph using *Articulate* than a more familiar tool, Microsoft Excel [34]. Research is needed to more deeply understand how to fully leverage this modality in SRDWs.

V. ACHIEVING 3-D

Thus far, SRDWs have primarily displayed 2-D content. With the advent of 3-D-capable LCD displays it is likely that all displays in the future will provide a 3-D display capability essentially for free in the same way that stereo and surround sound audio are a standard feature of all audio/visual systems today. In the early 1990s, stereoscopic head-mounted displays (HMDs) and immersive CAVEs [7] cost from tens of thousands to one million dollars; today a 3-D television costs as little as a thousand dollars.

The holy grail of stereoscopic computer displays is autostereoscopy—the ability to show stereoscopic images without requiring the user to wear any form of mediating glasses. Autostereo displays work by presenting the viewer with a vertically interleaved set of left and right eye stereoscopic images, and then using either vertical strips of plastic lenses or a physical line screen barrier separated by a small gap. This allows separate left and right eye images to be presented to the viewer without requiring them to wear any form of eyewear, such as those used in HMDs or the CAVE. One drawback of this technique is that when a viewer moves his/her head an opposite pair of images can potentially be presented to the viewer creating reversed stereo (also known as pseudostereo). To compensate for this, head tracking systems such as the ones used in the CAVE, or camera-based tracking systems, were used to track the viewers head to ensure the correct set of left/right eye images were presented to the viewer. To support multiple viewers, multiple views must be simultaneously projected from the display. Approaches for this include: spatial multiplexing—where the total resolution of the display is split between multiple views; multiprojector—where an array of projectors simultaneously project multiple views onto a special transmissive or reflective screen; or time-sequential—where a single very fast display device creates multiple views that are synchronized with a secondary optical component (such as a ferroelectric liquid crystal shutter) that directs the images to the appropriate zones in space. At the present time, the only practical approach is spatial multiplexing, as the approach can take advantage of new and higher density LCD displays.

One significant drawback of autostereo displays is that at least half the horizontal resolution of a display is lost in order to present both left and right eye images. This has led some researchers to develop high-resolution autostereo SRDWs using either a tiling of LCD panels or an array of horizontally placed projectors [25]. In the former case a barrier-screen-based system called Varrier used a 7×5 tiling of 21-in LCD panels, providing an autostereo resolution of 3000×6000 [32], pictured in Fig. 7. Shortly after the invention of the Varrier, a major breakthrough was achieved whereby a screen could display both autostereoscopic and monoscopic content within sub-windows on a single screen. Known as the Dynallax (for Dynamic Parallax barrier), this particular approach replaced the static Varrier barrier screen with an LCD panel that could be dynamically controlled to draw lines that are any width and pitch, as well as regions with and without the line screens (and hence regions with and without the stereo effect) [28].

Whereas Varrier was designed to provide autostereo for a single viewpoint, the Dynallax technique has been shown to be able to potentially support multiple views. The ability to present multiple views simultaneously allows more than one viewer to see autostereoscopic images at the same time. This is important if autostereo is to ever become a practical mainstream display device.




Fig. 7. The Varrier autostereo display wall (2005).

Commercial products are also beginning to emerge, albeit still at a very low resolution. For example Philips' WOWzone system consists of a tiling of 3×3 42-in 1920×1080 LCDs that provides nine views at an effective autostereo resolution of 640×3240 pixels per eye.

Whereas it will take many years for autostereo to be perfected to the point that they are practical commercial products, it is possible in the interim to develop high quality stereoscopic SRDWs using passive stereo techniques, such as in CAVE2 [9], [23], [24].

The goals in selecting a display technology for CAVE2 are scalable resolution, high brightness and contrast, thin borders, and low maintenance. These factors argue for using LCDs, especially given the emergence of ultrathin bezels (< 4 mm per side). Several 3-D technologies for LCDs were considered and passive micropolarized stereo was determined to offer the best combination of image quality and practicality. A micropolarizer is a thin (0.8 mm) overlay that polarizes each pixel row in alternate directions, as in Fig. 8. Its strengths are simple construction, compatibility with thin-bezel displays, commercial availability, and good image quality normal to the display surface as well as off-axis in the direction parallel to the micropolarizer lines. Its main drawback is image crosstalk in the direction perpendicular to the micropolarizer lines. Several configurations have been tested, including finished products as well as inhouse bench assembly of a micropolarizing sheet on top of LCD displays from two manufacturers. All displays were 46-in, 1920×1080 , with a pixel pitch of 0.530 mm.

The micropolarizing sheet consists of rows of alternating retarders, further separated by black lines called guard bands, on a glass substrate. When the micropolarizer is registered with the pixel grid, alternate pixel rows are polarized oppositely and are visible in each eye.

The percentage of crosstalk is the amount of light, intended for one eye, entering the other eye. While zero crosstalk is ideal, in practice 2%–5% is very good, and

5%–8% is acceptable. Three test patterns were used for quantitative evaluation (see Fig. 9) in addition to qualitatively evaluating virtual scenes of 3-D objects. The full-field pattern consists of a black image for one eye and white for the other and is used for registering the micropolarizer overlay with the LCD pixel rows. Both angular and positional registration are easily found by adjusting the orientation and location of the overlay in order to create full, uniform fields that maximize the contrast between the white and black channels.

The cross-hatch test pattern consists of white diagonal bars in opposite directions, angled up for one eye and down for the other eye, and is used to measure crosstalk photometrically by aiming a spot-meter at four different regions in the pattern. If region 0 is black in both channels, region 1 is black in the “signal” channel and white in the crosstalk channel, region 2 is white in the signal channel and black in the crosstalk channel, and region 3 is white in both channels, then the crosstalk is $(I_1 - I_0)/(I_2 - I_0)$, where I_n is the intensity in region n . The process is repeated for each eye, and the results are averaged.

The Weissman test pattern is a visual gradient consisting of varying gray-scale reference bars [39], and crosstalk is visually compared to a reference pattern. The percentage crosstalk for each set of reference lines has been precomputed and is read directly from the chart. Accurate results require color and contrast to be matched among displays, and there is also some subjectivity when different individuals view the pattern. In practice, the pattern is accurate to a range of values (for example, between 2%–4%), and the center of that range was recorded.

Table 1 shows the results of the crosstalk measurements on three LCD displays and demonstrates that overall crosstalk levels are acceptable. A good optical bond can improve crosstalk by a factor of two, as the first and second rows (bonded versus unbonded) of Fig. 9 show. It is




Fig. 8. Exploded view of the main components of a micropolarized LCD 3-D display.

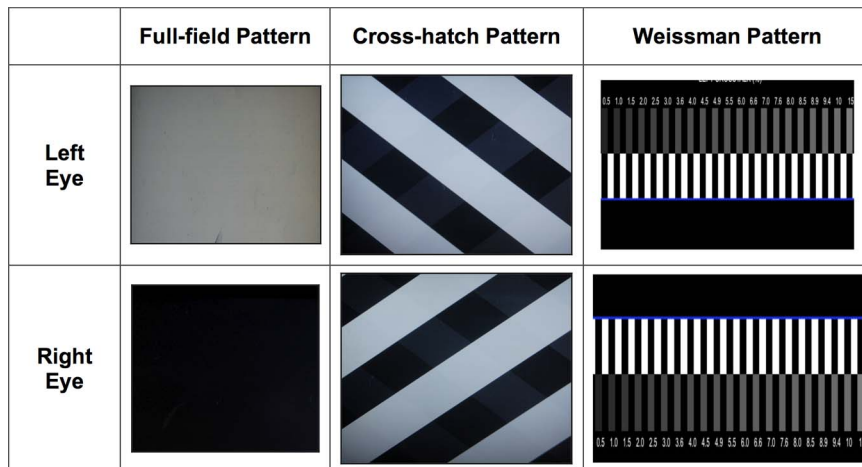


Fig. 9. Test pattern used to register micropolarizer overlay with pixel grid (full-field) and two test patterns to measure crosstalk (cross-hatch and Weissman).

believed that different pixel substructure, brightness and contrast account for the differences between the second and third rows (unbonded display models). Our tests indicate that thin-bezel micropolarized displays are a viable direction for the development of CAVE2 display technology. More commercial models have become available since these tests were conducted, and those will be evaluated as well.

One limitation of micropolarized displays is that crosstalk increases dramatically when viewing off-center in a direction perpendicular to the micropolarizer lines. Tests showed a cutoff at approximately 15° – 20° . This must be considered when designing tiled walls and CAVEs. For example, a wall spanning an 8-ft high space would need to be viewed from 12 ft away in order for the top and bottom of the display to have acceptable crosstalk. Taller spans or closer viewing distances can be mitigated by tilting the top and bottom rows. Another approach that is being investigated is to shift or bias the micropolarizer overlay

so that the micropolarizer rows are centered with pixel rows relative to the off-axis line of sight rather than relative to the normal direction of the screen. To the best of our knowledge, this approach has not appeared in prior literature. Several configurations of CAVE2s using these ideas are illustrated in Fig. 10.

VI. USE CASES

SRDWs are proliferating rapidly in universities, national laboratories as well as industry. For example in 2011 over





Fig. 10. Possible designs for CAVE2. The top designs use 46-in micropolarized LCDs, the middle designs use 55-in displays. The bottom design is an artist's conception of the final design.

Table 1 The LCD Panels Are Evaluated Using Two Test Patterns. Bonded/Unbonded Refers to Whether a Commercial Finished Product Was Evaluated Versus Our Bench Assembly of an Xpol Overlay on Top of the LCD, Without an Optical Bond

LCD Type	Cross-hatch Pattern Crosstalk	Weissmann Pattern Crosstalk	Average Crosstalk
Bonded	2%	2%	2%
Unbonded Brand X	5%	3%	4%
Unbonded Brand Y	10%	4%	7%

70 sites around the world had adopted the technology and are using SAGE to drive them [20], [31].

A. Antarctic Exploration

ENDURANCE is a NASA funded project that sent a robotic submersible under ice-covered West Lake Bonney in the Taylor Valley of Antarctica to gather data about the three-dimensional biochemical structure of the interior of the lake, and to create high-resolution sonar maps of the bottom of the lake and the face of Taylor Glacier at the end of the lake. Biochemical data, camera data, and sonar data from the robot was collected in fall 2008 and fall 2009. EVL's high-resolution walls were used for several activities during this project.

Mission Planning: The science and engineering team members viewed full-resolution Quickbird satellite photos of the Taylor Valley and Lake Bonney on the wall in the months before the mission. The members of the team who had been to the lake before were better able to brief the new members of the expedition on what to expect and discuss possible locations to place the camp and melt holes on the surface of the lake.

Data Validation: The science and engineering team members used the EVL cybercommons in April 2010 to look at all the data collected from both deployments (see Fig. 11). Validating the data required simultaneous viewing of tables of data, graphs of the data, and the geographic location and times of the data collection points, which was easy given the size and resolution of the cybercommons wall. Linking these representations together allowed aberrations in the graphs to be easily located in time and space and decisions could be made on whether the data was valid.

Research and Analysis: The science team used the high resolution walls to easily integrate a full-resolution Quickbird satellite photo of Lake Bonney taken during the mission with full-resolution photographs taken by the robot from beneath the surface of the lake looking upwards through the ice to determine the amount of small debris present at different points in the lake, which affects the amount of light that is transmitted through the ice, which affects the biochemical properties. The wall allowed the science team to see both context and detail simultaneously and made it easier to see overall patterns.

The processing of the sonar data for the bottom of the lake and the glacier face was complicated due to a concentrated salt layer in the lake. Various sets of parameters were used to generate a final sonar point cloud and to convert it into a polygonal mesh. The science team used the cybercommons in April 2011 to simultaneously display visualizations of the results from the different sets of parameters to better compare and contrast them and pick the most appropriate set of values.

In May 2011, the science team returned to analyze the data using the wall to simultaneously show multiple representations of the data to better discuss and understand the underlying phenomena in the lake. The wall helped externalize and share the mental models of the biochemical processes that individuals were using.

B. Classrooms of the Future

As with most universities, the University of Illinois at Chicago has multimedia classrooms with a projector and computer. These standard multimedia rooms are not well equipped for courses in areas such as visualization, visual analytics, computer graphics, interaction, or video game design. For these courses it is important to have multiple screens of notes, code, and examples available, and its



Fig. 11. The ENDURANCE teams used EVL's high-resolution walls to juxtapose multiple data sets for data validation and analysis.

impossible to do that with a single projection screen. These standard multimedia rooms also make it difficult to teach courses with remotely participating classrooms or speakers. They typically do not have enough bandwidth for high-quality video conferencing, and not enough screens to simultaneously show the remote site and the shared topic of discussion.

In order to address this need in 2004 EVL converted one of our large meeting rooms into a classroom space for 40 students, and it has been evolving since that time with increasing numbers of screens. In summer 2009, we installed a 6×3 tiled LCD wall of thin-border 46-in digital signage displays with a total resolution of 18 Mpixels driven by a single PC with 20 Gb of network capability. The room also contains ambient microphones for the class, dedicated mikes for speaker, and multiple high definition cameras for remote sessions. In the summer of 2010 we added a touch screen to the wall allowing interaction with content on the wall at the wall itself in addition to using a gyromouse, or a laptop. The wall acts like a traditional classroom that had multiple sliding blackboards or whiteboards for increased space to give context to the current discussion, but with the advantages of being able to bring up digital media on those screens. The single computer driving the display makes support simple. The small borders mean text is always readable.

In the last two years six faculty members have taught ten different courses in this room including courses in computer science, high-energy physics, and art and design,

and three of those courses were taught with remote participants. The wall is typically divided into sections for classwork, and different courses divide the wall in different ways. Local courses typically break the wall into three parts with 1/3 for the current notes in powerpoint or html format, 1/3 for a variety of related images, movies, or pdf documents to give context, and 1/3 for a Wacom tablet streamed to wall as a whiteboard. Students can add to the discussion by dragging and dropping images, pdf files, or movies onto the main screen and pointing at media on the wall via their laptops. Social constraints are effective in limiting the potential anarchy. During student project presentations the wall is typically broken into two parts with the students presenting from their laptop on half, while the other half hold images of the different project solutions the class has come up with allowing the rest of the class to compare and contrast them. Remote courses typically break the wall into two parts with half for the current notes shared between the sites and half for the remote site in high-resolution video.

There is also considerable ad hoc use of the wall. Groups of students in the Visual Analytics and Video Game Programming courses use the wall to brainstorm different visual and interaction ideas by sketching with the streaming Wacom whiteboard on part of the wall, and streaming their laptop screens to windows on the wall to display previous solutions or related material from the web. Similarly groups of graduate students brainstormed ideas for their research projects using the streaming whiteboard for

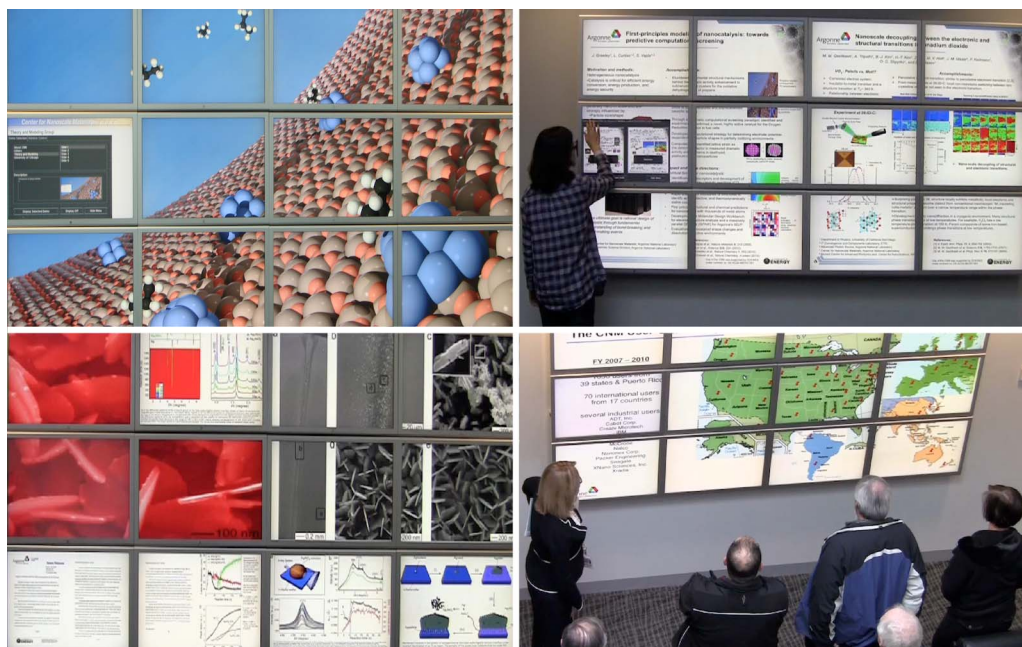


Fig. 12. From top left to bottom right, four snapshots depicting the use of the Science Information Wall at the Center for Nanoscale Materials: (a) display of dynamic molecular simulations (propane reaction on supported platinum nanoclusters); (b) presentation of research posters; (c) presentation of research highlights (growth process of nanoparticles); and (d) introductory presentations for visitors of the facility.



Fig. 13. Argonne National Laboratory's virtual poster gallery.

drawing while posting related pdf papers on the wall. We also saw groups of doctoral students studying for the computer science qualifying exam by putting multiple exams and answers along with related web pages on the screens for discussion.

Surveys of 30 students showed that 2/3 liked it better than a single computer projector, and half thought it much better than one or more whiteboards or blackboards. Two-thirds thought visibility of information was much better than in traditional classrooms, and half thought it was much better in helping them follow their lectures.

C. Nanoscale Materials Science Research

The Center for Nanoscale Materials (CNM) at Argonne National Laboratory is a premier user facility, providing expertise, instruments and infrastructure for interdisciplinary nanoscience and nanotechnology research. There are 1050 users from 39 states and Puerto Rico, 70 international users from 17 countries, and several industrial users. Two SRDWs were installed (see Fig. 12 and 13) to enable the examination of experimental as well as simulation data from research experiments, the presentation of research findings, and the observation of clean room activities. The Science Information Wall (SIW) is a 4×3 SRDW with

44 Mpixels of resolution. This made it particularly useful for arbitrarily accessing and juxtaposing a wide variety of data types that include: scanning transmission electron microscope imagery, reflected light optical microscopy images, hydrogen bonds in water and polymer, nanomagnetic toggle switching of vortex cores, dynamic evolution of Molecular Dynamics and Kinetic Monte Carlo simulations, platinum atoms deposited in aluminum oxide membrane. The Virtual Poster Gallery (VPG) is a 12×1 SRDW with 24 Mpixels of resolution that is used to display numerous scientific posters as well as real-time camera views of the nanotech clean rooms. Users interact with the SIW using a touch screen mounted over one of the LCD panels to serve as a master control interface. For the VPG, users can use an Apple Magic Trackpad and its associated multifinger gesture controls to view and manipulate the content. As the use of the walls gained popularity when they were integrated with CNM's Enterprise Content Management system to make it a part of their research and presentation workflow.

VII. CONCLUSION

Looking out to the next ten years, the adoption of scalable resolution displays will likely expand beyond academic institutions and research laboratories to include offices, and eventually homes. Their resolution will increase to approach print quality thereby realizing the possibility of digital wallpaper. Rather than being driven by a large cluster of computers or a single multiheaded graphics computer, small-scale computing will be embedded directly into the support structure of the wall paper and will form a seamless network on which image data can be streamed for display. Each pixel in the display will also act as an image sensor in a high-resolution camera that can be used for video conferencing, as well as multitouch, and touchless gestural interaction. Stereoscopic 3-D will be built into these systems and will cost virtually nothing. All these capabilities will provide a visual and interaction experience that will enable a sense of presence and immersion currently unparalleled. ■

REFERENCES

- [1] B. Paul, S. Ahern, E. W. Bethel, E. Brugger, R. Cook, J. Daniel, K. Lewis, J. Owen, and D. Southard, "Chromium renderserver: Scalable and open remote rendering infrastructure," *IEEE Trans. Visualiz. Comput. Graph.*, vol. 14, no. 3, pp. 627–639, May/Jun. 2008.
- [2] R. Ball and C. North, "Analysis of user behavior on high-resolution tiled displays," in *Human-Computer Interaction—Interact*, 2005, pp. 350–363, DOI: 10.1007/11555261_30.
- [3] A. Cedilnik, B. Geveci, K. Moreland, J. Ahrens, and J. Favre, "Remote large data visualization in the ParaView framework," in *Proc. Eurograph. Workshop Parallel Graph. Visualization*, 2006, pp. 162–170.
- [4] *IEEE Comput. Graph. Appl.*, vol. 20, no. 4, 2000.
- [5] *IEEE Comput. Graph. Appl.*, vol. 25, no. 4, 2005.
- [6] A. Christopher, E. Alex, and N. Chris, "Space to think: Large, high-resolution displays for sensemaking," in *Proc. ACM Human Comput. Interaction*, 2010.
- [7] C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti, "Surround-screen projection-based virtual reality: The design and implementation of the CAVE," in *Proc. ACM Computer Graph. Conf.*, Aug. 1993, pp. 135–142.
- [8] M. Czerwinski et al., "Large display research overview," in *Proc. ACM Human Comput. Interaction*, 2006.
- [9] T. A. DeFanti, D. Acevedo, R. A. Ainsworth, M. D. Brown, S. Cutchin, G. Dawe, K.-U. Doerr, A. Johnson, D. 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 Eur. J. Eng. Online First*, Nov. 2, 2010.
- [10] Xdmx: Distributed Multi-Head X Project. [Online]. Available: <http://en.wikipedia.org/wiki/Xdmx>
- [11] K. Doerr and F. Kuester, "CGLX: A scalable, high-performance visualization framework for networked display environments," *IEEE Trans. Vis. Comput. Graph.*, vol. 17, no. 3, pp. 320–332, May 2011.
- [12] S. Eilemann, M. Makhinya, and R. Pajarola, "Equalizer: A scalable parallel rendering framework," *IEEE Trans. Vis. Comput. Graph.*, vol. 15, no. 3, pp. 436–452, May 2009.
- [13] Y. Etsion, D. Tsafir, and D. G. Feitelson, "Desktop scheduling: How can we know what the user wants?" in *Proc. 14th Int. Workshop*

- Netw. Oper. Syst. Support Digital Audio Video, 2004, pp. 110–115.
- [14] S. Evans, K. Clarke, D. Singleton, and B. Smaalders, “Optimizing unix resource scheduling for user interaction,” in *Proc. USENIX Summer Techn. Conf.*, Cincinnati, OH, 1993, vol. 1.
 - [15] K. Heider, “The Rashomon effect: When ethnographers disagree,” *Amer. Anthropol.*, vol. 90, no. 1, pp. 73–81, Mar. 1988.
 - [16] G. Humphreys, M. Houston, Y. Ng, R. Frank, S. Ahern, P. Kirchner, and J. Klosowski, “Chromium: A stream-processing framework for interactive rendering on clusters,” in *Proc. ACM Comput. Graph. Conf.*, 2002, pp. 693–702.
 - [17] IBM Systems and Technology Group, IBM deep computing visualization, Technical White Paper, 2005. [Online]. Available: <http://www-03.ibm.com/servers/deepcomputing/visualization/downloads/dcv%techw.pdf>
 - [18] R. Jagodic, L. Renambot, A. Johnson, J. Leigh, and S. Deshpande, “Enabling multi-user interaction in large high-resolution distributed environments,” *Future Generat. Comput. Syst.*, vol. 27, no. 7, pp. 914–923, Jul. 2011.
 - [19] B. Jeong, L. Renambot, R. Jagodic, R. Singh, R. J. Aguilera, A. Johnson, and J. Leigh, “High-performance dynamic graphics streaming for scalable adaptive graphics environment,” in *Proc. Supercomput.*, Nov. 2006.
 - [20] B. Jeong, J. Leigh, A. Johnson, L. Renambot, M. Brown, R. S. Nam, and H. Hur, “Ultrascale collaborative visualization using a display-rich global cyberinfrastructure,” *IEEE Comput. Graph. Appl.*, vol. 30, no. 3, pp. 71–83, May/Jun. 2010.
 - [21] D. Kirsh, “The intelligent use of space,” *Artif. Intell.*, vol. 73, no. 1–2, pp. 31–68, 1995.
 - [22] J. Leigh, A. Johnson, K. Park, A. Nayak, R. Singh, and V. Chowdry, “Amplified collaboration environments,” in *Proc. VizGrid Symp.*, Tokyo, Japan, Nov. 2002. [Online]. Available: www.vizgrid.org
 - [23] J. Leigh, A. Johnson, L. Renambot, S. Dan, T. Defanti, M. Brown, B. Jeong, R. Jagodic, C. Krumbholz, D. Svistula, H. Hur, R. Kooima, T. Peterka, J. Ge, and C. Falk, “Emerging from the CAVE: Collaboration in ultra high resolution environments,” in *Proc. 1st Int. Symp. Universal Commun.*, Kyoto, Japan, Jun. 14–15, 2007, pp. 96–99.
 - [24] J. Leigh, A. Johnson, and L. Renambot, “Advances in computer displays,” *Adv. Comput.*, vol. 77, pp. 58–79, 2009.
 - [25] G. J. Martin, A. L. Smeyne, J. R. Moore, S. R. Lang, and N. A. Dodgson, “Three-dimensional visualization without glasses: A large-screen autostereoscopic display,” in *Proc. SPIE—Cockpit Displays VII: Displays for Defense Applications*, vol. 4022. [Online]. Available: <http://dx.doi.org/10.1117/12.397765>
 - [26] S. Molnar, M. Cox, D. Ellsworth, and H. Fuchs, “A sorting classification of parallel rendering,” *IEEE Comput. Graph. Appl.*, vol. 14, no. 4, pp. 23–32, 1994.
 - [27] K. Nakayama and G. H. Silverman, “Serial and parallel processing of visual feature conjunctions,” *Nature*, vol. 320, pp. 264–265, 1986.
 - [28] T. Peterka, R. L. Kooima, J. I. Girado, J. Ge, D. J. Sandin, A. Johnson, J. Leigh, J. Schulze, and T. A. DeFanti, “Dynamall: Solid state dynamic parallax barrier autostereoscopic VR display,” in *Proc. IEEE Virtual Reality Conf.*, Charlotte, NC, 2007.
 - [29] L. Renambot, B. Jeong, H. Hur, A. Johnson, and J. Leigh, “Enabling high resolution collaborative visualization in display rich virtual organizations,” *Future Generat. Comput. Syst.*, vol. 25, no. 2, pp. 161–168, Feb. 2009.
 - [30] T. Richardson, Q. Stafford-Fraser, K. Wood, and A. Hopper, “Virtual network computing,” *IEEE Internet Comput.*, vol. 2, no. 1, pp. 33–38, Jan./Feb. 1998.
 - [31] Scalable Adaptive Graphics Environment. [Online]. Available: <http://www.sagecommons.org>
 - [32] D. Sandin, T. Margolis, J. Ge, J. Girado, T. Peterka, and T. A. DeFanti, “The Varrier autostereoscopic virtual reality display,” in *Proc. ACM Comput. Graph. Conf.*, Jul. 30, 2005.
 - [33] L. Smarr, M. Brown, and C. de Laat, “Special section: OptiPlanet—The OptiPutter global collaboratory,” *Future Generat. Comput. Syst.*, vol. 25, no. 2, pp. 109–113, Feb. 2009.
 - [34] Y. Sun, J. Leigh, A. Johnson, and S. Lee, “Articulate: A semi-automated model for translating natural language queries into meaningful visualizations,” in *Proc. 10th Int. Symp. Smart Graph.*, Banff, Canada, 2010, pp. 184–195.
 - [35] S. Teasley, L. Covi, M. Krishnan, and J. Olson, “How does radical collocation help a team succeed?” in *Proc. Comput. Supported Coop. Work*, New York, 2000, pp. 339–346.
 - [36] V. Vishwanath, J. Leigh, T. Shimizu, S. Nam, L. Renambot, H. Takahashi, M. Takizawa, and O. Kamatani, “The Rails Toolkit (RTK)—Enabling end-system topology-aware high end computing,” in *Proc. IEEE 4th Int. Conf. e-Science*, 2008, pp. 309–316.
 - [37] VisIt. [Online]. Available: <https://wci.llnl.gov/codes/visit/home.html>
 - [38] C. Ware and G. Franck, “Viewing a graph in a virtual reality display is three times as good as a 2D diagram,” in *Proc. IEEE Visual Lang.*, 1994, pp. 182–183.
 - [39] M. A. Weissman and A. J. Woods, “A simple method for measuring crosstalk in stereoscopic displays,” in *Proc. SPIE Stereoscopic Displays Appl.*, 2011, vol. 7863, 786310.
 - [40] H. Zheng and J. Nieh, “RSIO: Automatic user interaction detection and scheduling,” in *Proc. ACM Int. Conf. Meas. Model Comput. Syst.*, New York, pp. 263–274.

ABOUT THE AUTHORS

Jason Leigh received the Ph.D. degree from the University of Illinois at Chicago, Chicago.

He is a Professor of Computer Science and Director of the Electronic Visualization Laboratory and the Software Technologies Research Center at the University of Illinois at Chicago, Chicago. He is an internationally recognized pioneer in collaborative virtual reality. His prior projects and research for which he is best known include: the OptiPutter, GeoWall, CoreWall, LambdaVision, Tele-Immersion, and Reliable Blast UDP. His research for the past ten years focused on ultra-resolution display-rich collaboration environments amplified by high-performance computing and networking. His current research focuses on human augmentics—the development of technologies for expanding the capabilities and characteristics of humans.



Andrew Johnson received the Ph.D. degree from Wayne State University, Detroit, MI.

He is an Associate Professor in the Department of Computer Science and a member of the Electronic Visualization Laboratory, University of Illinois at Chicago, Chicago. His research and teaching focus on interaction and collaboration using advanced visualization displays and the application of those displays to enhance discovery and learning.



Luc Renambot received the Ph.D. degree from the University of Rennes-1, France, in 2000, conducting research on parallel rendering algorithms for illumination simulation.

Holding a postdoctoral position at the Free University of Amsterdam, Amsterdam, The Netherlands, until 2002, he worked on bringing education and scientific visualization to virtual reality environments. Since 2003, he has been with the Electronic Visualization Laboratory, University of Illinois at Chicago, Chicago, first as a Postdoctoral Researcher and now as Research Assistant Professor, where his research topics include high-resolution displays, computer graphics, parallel computing, and high-speed networking.



Tom Peterka received the Ph.D. degree in computer science from the University of Illinois at Chicago, Chicago, in 2007, where he studied the use of virtual environments for scientific visualization at the Electronic Visualization Laboratory.

He joined the Mathematics and Computer Science Division at Argonne National Laboratory, Argonne, IL, in 2007, where he researches analysis and visualization on leadership-class architectures such as the IBM Blue Gene/P and develops scalable algorithms for the maximum number of available processors.



Byungil Jeong received the B.S. and M.S. degrees in electrical engineering from Seoul National University, Seoul, Korea and the Ph.D. degree in computer science from the University of Illinois at Chicago, Chicago.

He worked as a primary Scalable Adaptive Graphics Environment (SAGE) Architect at the Electronic Visualization Laboratory, University of Illinois at Chicago. He was a member of the Technical Staff at Sun Advanced Visualization group before joining the Texas Advanced Computing Center, Austin.



Ratko Jagodic received the Ph.D. degree from the University of Illinois at Chicago (UIC), Chicago.

While at UIC, he developed user-interaction methodologies for ultrahigh-resolution display walls. He is currently at Intel Corporation.



Daniel J. Sandin received the B.A. degree in natural sciences from Shimer College, Chicago, IL, in 1964 and the M.S. degree in physics from the University of Wisconsin—Madison, Madison, in 1967.

He is a Professor Emeritus of the School of Art & Design, University of Illinois at Chicago, Chicago, and Co-Founder and Co-Director of the Electronic Visualization Laboratory at the University of Illinois at Chicago. He is an internationally recognized pioneer in computer graphics, electronic art, and visualization. He became interested in video in 1967, while helping to organize student demonstrations at the University of Illinois. In 1969, he moved from the University of Wisconsin—Madison to the University of Illinois at Chicago after his presentation at the Smithsonian of “Glowflow,” a computer controlled light and sound system, created with Myron Krueger. His major achievements were a series of projects Glowflow (1969), Sandin Image Processor (IP) (1971–1973), Sayre Glove (1977), PHSColograms (1988), CAVE (1991), and ImmersaDesk and Infinity Wall.



Sungwon Nam received the M.S. degree from University of Southern California, Los Angeles. Currently, he is working toward the Ph.D. degree at the University of Illinois at Chicago, Chicago.

His research examines scheduling algorithms for interactive media on ultrahigh-resolution display walls.



Jonas Talandis received the Ph.D. degree in bioengineering from the University of Wyoming, Laramie, WY.

He is an Engineer and Product Creator currently contributing design and development expertise to advanced, highly interactive visualization systems as Senior Research Scientist at the Electronic Visualization Laboratory (EVL), University of Illinois at Chicago, Chicago. His recent work involves 3-D display and environment development for EVL’s NG-CAVE, Cyber-Commons 3-D, and OmegaDesk. He was a founding developer responsible for commercialization and tech development at the University of Illinois’ global-distributed-technology-center *TRECC*, and was instrumental in evolving that center into the Argonne National Laboratory managed *TRACC*. Prior to academia, his experience includes several years of business and product development in industry, specializing in analytical instrumentation for materials testing, process monitoring and medical applications.



Hyejung Hur received the B.S. degree in computer science from Duksung Women’s University, Korea, in 1999 and the M.S. degree in 2001. She is currently working toward the Ph.D. degree in the Department of Computer Science, University of Illinois at Chicago (UIC), Chicago.

Her research interest lies on the area of interactive visualization and collaborative environment.



Yiwen Sun received the B.S. degree in computer science from Shandong University, China, in 2003, and the Ph.D. degree from the University of Illinois at Chicago, Chicago, in 2012.

Her research interest lies in the area of interactive visualization and visual analytics. More specifically, her Ph.D. focused on developing a methodology for automatically translating natural language into visualizations. She currently works for Microsoft Corporation.

