

# CAVE2: A Hybrid Reality Environment for Immersive Simulation and Information Analysis

Alessandro Febretti<sup>a</sup>, Arthur Nishimoto<sup>a</sup>, Terrance Thigpen<sup>a</sup>, Jonas Talandis<sup>a</sup>, Lance Long<sup>a</sup>, JD Pirtle<sup>a</sup>, Tom Peterka<sup>a</sup>, Alan Verlo<sup>a</sup>, Maxine Brown<sup>a</sup>, Dana Plepys<sup>a</sup>, Dan Sandin<sup>a</sup>, Luc Renambot<sup>a</sup>, Andrew Johnson<sup>a</sup>, Jason Leigh<sup>a</sup>

<sup>a</sup>Electronic Visualization Laboratory, University of Illinois at Chicago (UIC)

## ABSTRACT

Hybrid Reality Environments represent a new kind of visualization spaces that blur the line between virtual environments and high resolution tiled display walls. This paper outlines the design and implementation of the CAVE2<sup>TM</sup> Hybrid Reality Environment. CAVE2 is the world's first near-seamless flat-panel-based, surround-screen immersive system. Unique to CAVE2 is that it will enable users to simultaneously view both 2D and 3D information, providing more flexibility for mixed media applications.

CAVE2 is a cylindrical system of 24 feet in diameter and 8 feet tall, and consists of 72 near-seamless, off-axis-optimized passive stereo LCD panels, creating an approximately 320 degree panoramic environment for displaying information at 37 Megapixels (in stereoscopic 3D) or 74 Megapixels in 2D and at a horizontal visual acuity of 20/20. Custom LCD panels with shifted polarizers were built so the images in the top and bottom rows of LCDs are optimized for vertical off-center viewing- allowing viewers to come closer to the displays while minimizing ghosting.

CAVE2 is designed to support multiple operating modes. In the Fully Immersive mode, the entire room can be dedicated to one virtual simulation. In 2D model, the room can operate like a traditional tiled display wall enabling users to work with large numbers of documents at the same time. In the Hybrid mode, a mixture of both 2D and 3D applications can be simultaneously supported. The ability to treat immersive work spaces in this Hybrid way has never been achieved before, and leverages the special abilities of CAVE2 to enable researchers to seamlessly interact with large collections of 2D and 3D data. To realize this hybrid ability, we merged the Scalable Adaptive Graphics Environment (SAGE) - a system for supporting 2D tiled displays, with Omegalib - a virtual reality middleware supporting OpenGL, OpenSceneGraph and Vtk applications.

**Keywords:** Hybrid Reality, Display Wall, Passive Stereoscopic, Immersive Systems

## 1. INTRODUCTION

In recent years the scientific community has faced an exponential increase in the amount of data gathered through the observation of natural phenomena or generated by complex supercomputer simulations. All of this data is collected at ever-increasing resolutions, stored digitally, and analyzed using computer-based visualization instruments: these visualization instruments (such as CAVE2) represent the modern version of traditional telescope and microscope lenses that have brought scientific phenomena into focus for centuries. As the complexity of data continues to grow, visualization instruments become increasingly essential to researchers, letting them transform raw data into discovery. In particular, immersive systems are an attractive option for exploring 3D spatial data such as molecules, astrophysical phenomena, and geoscience datasets<sup>1</sup>.

For decades, immersive systems were based on bulky Head Mounted Displays, with limited resolution and field of view. In 1992, the original CAVE (Cave Automatic Virtual Environment) created a paradigm shift in Virtual Reality<sup>2</sup>: CAVE consisted of a cube measuring 10 feet at each side, with projector-based stereoscopic graphics on five of its sides, and it allowed users to experience stereo 3D wearing much lighter shutter stereo glasses. The system was large enough to fit multiple individuals who could simultaneously experience the visualization. Additionally, head tracking enabled scientists to explore complex datasets using embodied interaction such as walking around in the CAVE or using a tracked 'wand' to fly through the virtual world.

While the original CAVE provided limited resolution (4 Megapixels) and poor contrast ratio, recent evolutions of the system utilize advanced projectors, yielding resolutions of up to 200 Megapixels<sup>3,4</sup>. However, both original and modern CAVE systems require specially designed rooms and low environment light levels to operate, making them difficult to integrate in office environments. Furthermore, high resolution projectors can be extremely expensive, and require regular maintenance (calibration and bulb replacement). This held up their adoption in everyday scientific workflows and limited them to opportunistic use.

More recently, tiled LCD display walls have emerged as a practical platform for large-scale data visualization<sup>5</sup>. Constructed by tiling multiple LCD monitors to form a contiguous display surface, these environments often span entire walls. When compared to projector-based setups, LCD walls provide superior image quality and resolution, often reaching 100 to 300 Megapixels, and require low maintenance. Large LCD display walls have been successfully integrated in everyday workflows: Scientists can use them to visualize very large datasets at their native resolution, providing both detail and context, or they can use the surface to lay out a variety of correlated visual data and documents for collaborative analysis<sup>6-9</sup>.

To make effective use of Display walls, a set of APIs and toolkits have been developed. One of the more successful among them is the Scalable Adaptive Graphics Environment (SAGE)<sup>10</sup>. SAGE's major innovation is that it can launch, display and manage the layout of multiple visualization: prior to SAGE, software to drive tiled displays operated like DOS in the 1980s, running only a single application at a time, and occupying the entire wall. SAGE turns a tiled display wall into a true multitasking environment, which better matches the way groups of users would ordinarily want to use a large wall.

## 2. TOWARDS THE HYBRID REALITY ENVIRONMENT

The qualities of CAVEs and display walls make them effective at visualizing different classes of data. CAVEs are extremely effective for visualizing 3D spatial datasets but are far less suited for 2D information. Large, high-resolution display walls on the other hand excel at visualizing multiple, big, 2D datasets, but are less suited for 3D data, as they do not provide the same degree of immersion and often lack the ability to show stereoscopic 3D.

Current technology trends lead to the productions of larger, affordable thin-bezel LCD monitors with built-in support for stereoscopic 3D<sup>3,11</sup>. Such recent advancements made it conceivable to merge the best-in-class capabilities of immersive Virtual Reality systems, such as CAVEs, with the best-in-class capabilities of Ultra-high-resolution Display Environments, thus creating conceptually new immersive environments which we refer to as Hybrid Reality Environments (HREs)

HREs have five main characteristics:

1. A large, high-resolution display that approaches the sphere of influence and perception of a human<sup>12</sup>, while providing a pixel density that comes close to matching the visual acuity of the human visual system.
2. Support for stereoscopic rendering for 3D datasets. The benefit of stereoscopy is not limited to 3D datasets however. 2D representations can greatly benefit from stereoscopy, providing an extra perceptual channel to encode additional data (such as time).
3. Support for naturalistic interactions: to offer a truly hybrid 2D/3D visualization platform, such environments should support a wide variety of interactions, including keyboard/mouse, 6 degrees-of-freedom joysticks, head tracking for a viewer-centered perspective, and voice-activated interfaces. The appropriate configuration is, of course, application dependent.
4. A space to encourage multiple, co-located scientists to collaborate. The ability to solve complex problems involving big data often requires a variety of scientific expertise. Therefore, hybrid environments should provide an inviting space where scientists can comfortably sit together, analyze, and interpret data.
5. A software layer that is able to leverage the hardware to simultaneously display multiple related datasets and utilize hybrid 2D/3D visualization and interaction modalities.

A few existing immersive environments follow some of the HRE guidelines named above. However, the CAVE2 system presented in this paper is, to date, the first full implementation of a true Hybrid Reality Environment.

### 3. DISPLAY TECHNOLOGY EVALUATION

The goals in selecting a display technology for CAVE2 are scalable resolution, high brightness and contrast, thin borders, and low maintenance. Several 3D technologies for LCDs were considered: auto, passive and active stereo. Autostereo was the initial goal in the 2007 conceptualization of CAVE2, but the state of the art in autostereo display technologies did not provide sufficient resolution, near-seamless tileability, and ability to switch on and off autostereo modes, to be deemed practical for large-scale deployment. This is still the case in 2013. Active stereo was also considered. Whereas active stereo LCD TVs are widespread, the only way to properly synchronize multiples of them was to use very costly G-sync-capable Nvidia Quadro graphics cards. Furthermore there is no commercially available thin-bezel active stereo LCD technology available. LG had demonstrated the viability of such a technology with a tiling of 3x3 but ultimately we decided to focus on passive stereo instead.

#### 3.1 Micropolarized passive stereo

We determined micropolarized stereo to be the passive stereo technology offering 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 shown in Error! Reference source not found.. 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. 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: 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.

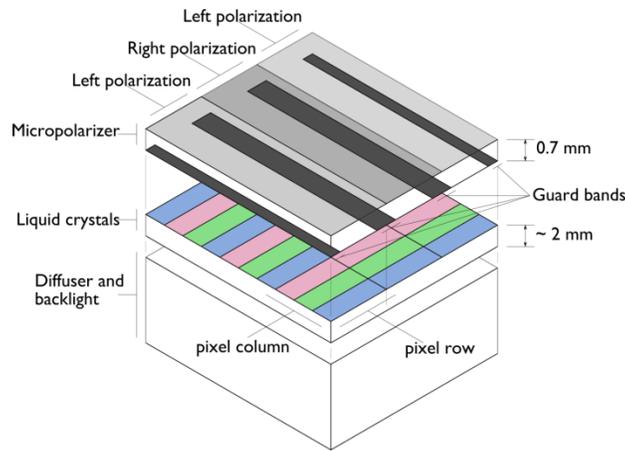
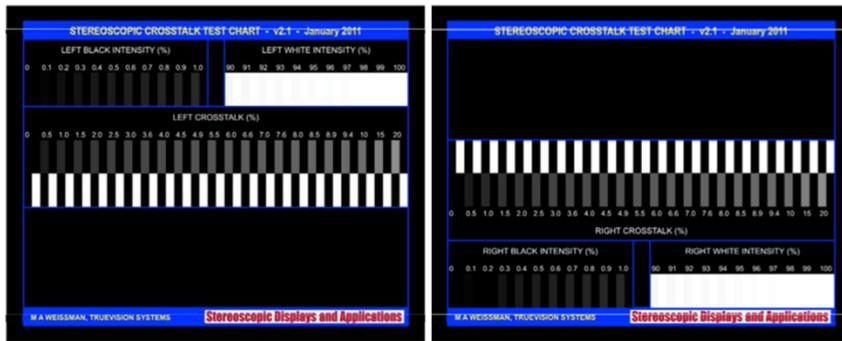


Figure 1. Exploded view of the main components of a micropolarized LCD 3D display

#### 3.2 Evaluating off-axis performance

Off-axis performance of displays under consideration for CAVE2 was determined empirically using the Weissman test pattern<sup>13</sup>, which is a simple method of viewing a test chart without using luminosity measurements or optical instruments. The viewer simply reads the crosstalk % from the interleaved scale at the place where the intensities of the two scales match while viewing with either the left or right eye. The opposite scale is used to determine crosstalk for the opposite eye. 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.



Observed Ghost (%) vs. View Angle (Deg)

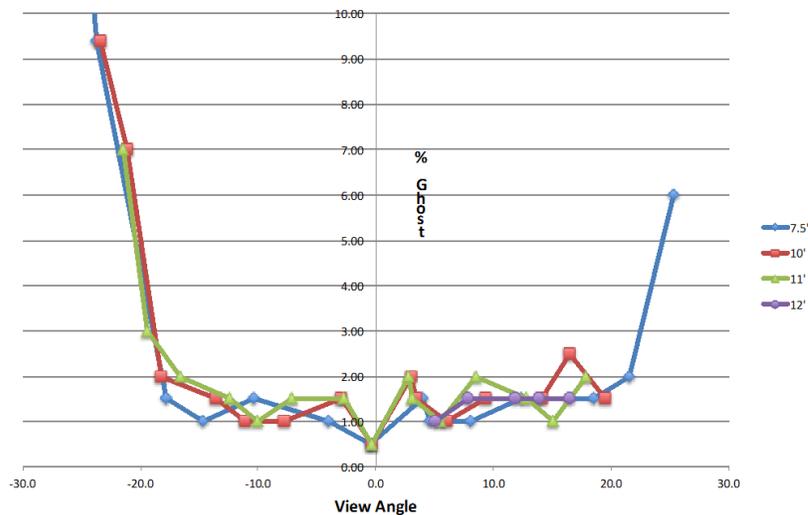


Figure 2. Top: left and right Weissman images appear as interleaved images when viewed in stereo. Bottom: observed ghost levels at varying view angles as observed using the Weissman pattern

Viewers of varying physical height (view height) observed the Weissman patterns located at specific target heights at varying view distances of 7.5', 10', 11' and 12'. The corresponding view angle for each measurement was calculated from the trigonometry of the viewing triangle. The results are plotted in Figure 2. Ghost levels were very consistent and low (<2%) throughout the viewing angles to approximately +/- 20 degrees. At that point, ghost level increased dramatically at a steep slope to uncomfortable and unusable levels (>10%).

### 3.3 Improving vertical off-axis performance

It was of utmost importance to consider the limits imposed by vertical off-axis crosstalk while designing CAVE2. For example, a setup based on an 8-foot high LCD wall would need to be viewed from 12 feet 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, or experimenting with other display setups that minimize off-axis viewing. Several configurations of CAVE2s using these ideas were considered, and are illustrated in Figure 3 (top).

Another approach 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. Tackling the off-axis issue through polarizer shifting instead of display tilting also allowed us to design CAVE2 using a pure cylindrical form factor shown in Figure 3 (bottom). This design is particularly desirable due to good immersion levels, active enclosed space, ease of assembly, and a seamless aesthetic design.

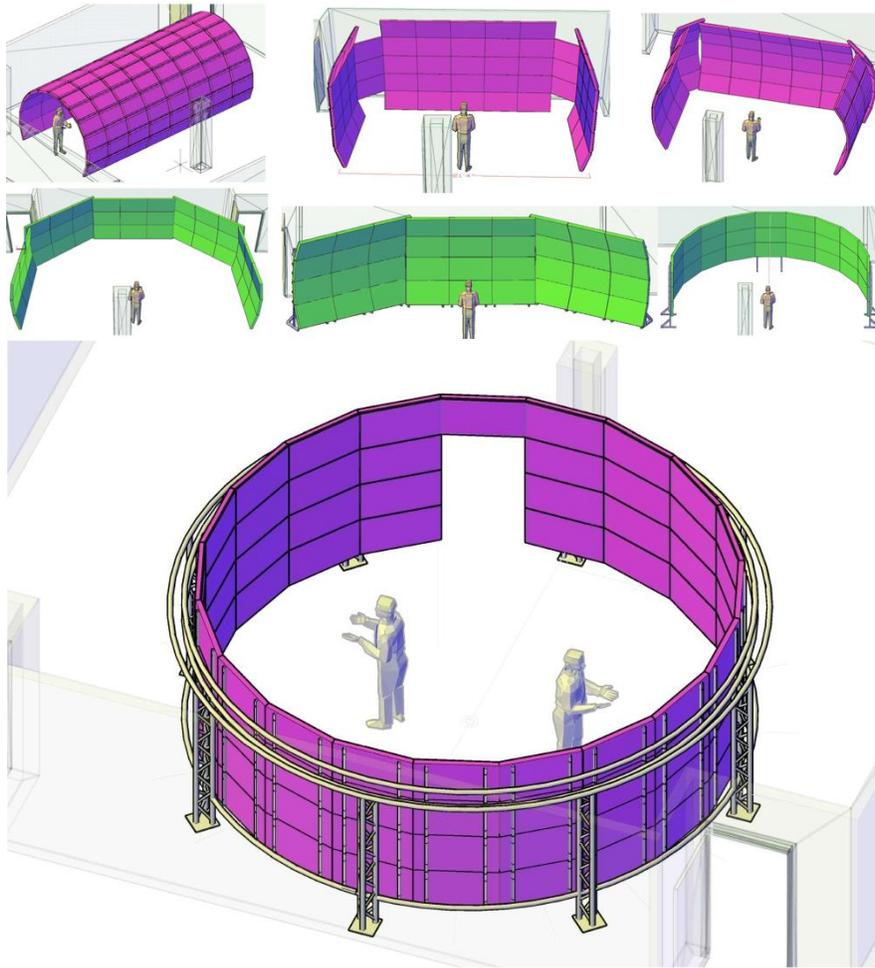


Figure 3. Top: possible designs for CAVE2. The top designs use 46-inch micropolarized LCDs, the middle designs use 55-inch displays. Bottom: concept drawing of the full-cylinder CAVE2

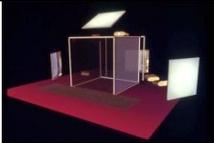
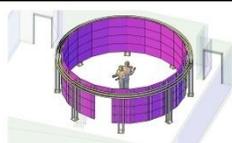
To accommodate off-axis vertical viewing angles of the very top and very bottom screens in a column of CAVE2, custom panels were developed in collaboration with the display provider, to shift the polarization filters. For example, for a positive off-axis view angle (viewer's eye lower than target) the polarization overlay would be shifted downward. Conversely, the polarizer would be shifted upward for negative viewing angles. These fixed shifts in registration would allow the viewer to see more of the "correct" pixel row and less of the "incorrect" or ghost producing adjacent pixel row. Experimental results and technical discussions determined the target shift angle specification to be +/- 9 degrees.

Expected luminance variations at the shifted display borders that were predicted by our initial geometric analyses were not evident in practice with our prototype. Measured light variations were on the order of 0.1 - 0.2 f-stops. Further evaluation determined that cumulative pitch variations in the micropolarizer registration measured from screen center to the edges actually worked in our favor in offsetting any visible luminance variations at the display seams. The net result was increased low-ghost viewing range, allowing a user to move closer to the display before exceeding the critical ghost angle.



Figure 4. A picture of the implemented CAVE2 Hybrid Reality Environment

Table 1. Comparison of the CAVE™ and CAVE2™ Systems

	CAVE	CAVE2	Improvement
<b>Image</b>			
<b>Year</b>	1992	2012	20 years
<b>Virtual Reality Environment– Cubic feet</b>	1,000 cu.ft. (10L x 10W x 10H)	3,167 cu.ft. ( $\pi \times \text{radius } 12\text{ft.}^2 \times 7\text{H}$ )	3X
<b>Projection vs 3D LCD</b>	4 Projectors	72 LCDs	n/a
<b>Stereo Resolution (Megapixels)</b>	2.6 (1280 x 512 x 4)	36 (1,366 x 736 x 72 / 2)	13X
<b>Visual Acuity</b>	20/110	20/20	–
<b>Brightness (lumens)</b>	4000 (1000 x 4)	266,400 (3,700 lumen x 72)	66.6X
<b>Contrast Ratio</b>	< 500:1	3,000:1	6X
<b>Bulb life (hours)</b>	2000 per projector	50,000 per LCD	25X
<b>Display cost per stereo Megapixel</b>	\$35,000 (Electrohome Marquee 8000)	\$14,000	2.5X
<b>Processor</b>	4x100 MHz MIPS R4000	36 x 2.9 GHz 16-core Xeon E5-2690	4,176X
<b>Graphics</b>	SGI Crimson VGXT	Nvidia GTX 680 2GB RAM	
<b>Memory</b>	256M	36 x 64GB	9000X
<b>Storage</b>	3.2GB	36 x 2TB	22,500X
<b>3D Tracking</b>	Cabled	Wireless	
<b>Networking</b>	10Mb/s	2x10Gb/s	2000X
<b>Total Cost</b>	\$2M in today's dollars	\$926K	50% cost

## 4. IMPLEMENTATION

The CAVE2 system is based on a cylindrical setup composed by 18 columns of 4 displays each. This arrangement provides a panoramic view of 320 degrees. Each display pair is driven by a separate computer, for a total of a 36 computer cluster, plus a head node. The 72 near-seamless passive displays have a resolution of 1366x768 pixels each and provide an aggregate resolution of 36 Megapixels per eye (74 MP total), almost 10 times the 3D resolution of the original CAVE. To an observer at the center of the system, this design provides a horizontal visual acuity of 20/20. A full comparison of CAVE and CAVE2 is summarized in Table 1.

### 4.1 Tracking System

CAVE2 uses an optical motion tracking system consisting of ten Vicon Bonita infrared cameras arranged in a circular configuration above the displays. Retroreflective markers are used to track the position and orientation of the primary viewer's head and the navigation controllers.

Our initial camera configuration had all ten cameras focused on the center of CAVE2. This allowed for very accurate tracking within a six foot radius from the center of CAVE2. (Figure 5 left) Since there is nearly an eleven foot radius of space within CAVE2, the tracked viewer would often walk within two feet of the displays and outside the trackable region. To compensate for this, we reoriented the cameras such that two cameras are focused on one of four sides of CAVE2 (Figure 5 right). This setup allows for tracking within two feet of the display without the need for additional cameras, at the cost of slightly decreased reliability at close range.

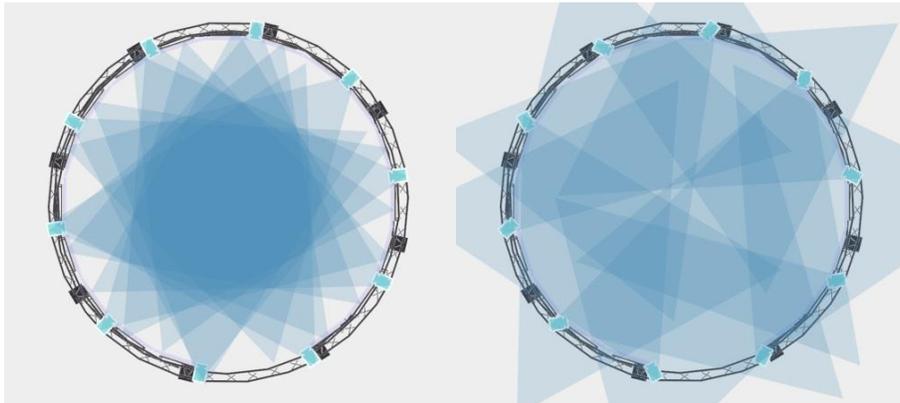


Figure 5. Left: original CAVE2 tracking camera coverage. A few areas next to the displays are not covered by any camera. Right: enhanced tracking camera coverage

The Vicon tracker uses the Virtual-Reality Peripheral Network (VRPN) protocol to send tracker data to an input management service. This service may aggregate data from other available input sources (like controllers, touch surfaces or other tracking technologies like the Kinect) before broadcasting input events to CAVE2 applications using TCP/UDP socket communication.

One of these applications is the CAVE2 System Display which shows the position and orientation of trackable objects, the status of the controllers, and the position of virtual sound objects currently playing inside CAVE2. This application also records when trackable objects are out of the camera view which is useful for calibrating the tracking camera configuration.

### 4.2 Audio system

The CAVE2 audio system is composed of 20 Genelec 6010A speakers and two 7050B subwoofers, for a total of 22 separate audio channels. Each speaker is mounted at the top of each column of displays, tilted at approximately 45 degrees toward the floor. The result is a full ring of speakers along the top of CAVE2. The target is a hypothetical cylinder located at the center of CAVE2, with a diameter roughly matching the ideal viewing area for stereo 3d, and with a height corresponding to average viewer height. The subwoofers are located just outside CAVE2, offset 90 degrees from the center of the visual environment. All 22 channels are connected to a digital-to-analog converter controlled by a

standalone audio machine through a MADI interface. MADI (Multichannel Audio Digital Interface) is a communications protocol that allows the transmission of 64 channels of 24-bit audio at a sample rate of up to 48 kHz, via either coaxial or fiber optic cable. This protocol is particularly desirable due to the support of relatively long cable lengths.

The audio computer runs a custom sound server written in Supercollider, an open source programming language and environment for real-time audio synthesis and algorithmic composition. One of the primary goals at the onset of software development was to provide a simple set of commands to interact with the CAVE2 Sound Server, which would facilitate playback and positioning of sound objects in virtual space. The sound server can be controlled through Open Sound Control (OSC) messages.

The physical realities of CAVE2, particularly the detrimental acoustic effects of a large polygonal glass environment, necessitated the installation of both acoustic ceiling tiles and carpet tiles to mitigate the reflection of acoustic waves.



Figure 6. CAVE2 running in full 3D mode. The visualization shows reconstructed brain vessels and cortical tissue gathered from MRI data, supplemented with artificially generated microvasculature. Effective display of the finer details of this dataset is possible thanks to CAVE2 resolution.

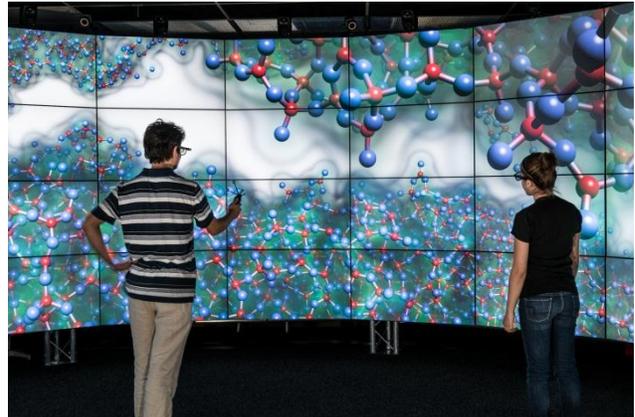


Figure 7. Another example of full 3D mode. Volume rendering of a glass fissure computed in a 5-Million atom molecular dynamics simulation.



Figure 8. CAVE2 can run in Hybrid mode, to juxtapose high-resolution 2D and 3D datasets side-by-side. The right side shows a stereoscopic, head-tracked rendering of planetary-scale Martian terrain data. The left side shows high-resolution pictures taken by the Curiosity rover along with 2D topographical maps of Mars.



Figure 9. CAVE2 running in full 2D mode through SAGE. With a diameter of more than 10 feet, CAVE2 is big enough to fit about 15 standing users, and can be easily turned into a small meeting room.

### 4.3 Operating modes

Due to its hybrid design, CAVE2 has multiple possible operating modes. In the Fully Immersive mode, the entire room can be dedicated to one virtual simulation (Figure 6, Figure 7). In 2D mode, the room can operate like a standard display wall, enabling users to work with large numbers of documents at the same time (Figure 9). In the Hybrid mode, some of the documents on the wall could represent immersive 3D windows that can be controlled with 3D user interfaces (Figure 8). In fact multiple 3D windows can be independently controlled by different users in the room if desired. The ability to treat information work spaces in this Hybrid way has never been achieved before, and leverages the special abilities of CAVE2 to enable researchers to seamlessly interact with large collections of 2D and 3D data.

## 5. HYBRID REALITY SOFTWARE

Although the CAVE2 system can run legacy / custom VR software libraries we also stated how HREs need a middleware to assist the development of applications on a hybrid 2D/3D system. Other than supporting 2D and 3D graphics, some reasons motivating the need for a custom software layer are:

- Support groups of observers within the system. The software should allow multiple users to perceive correct or nearly correct stereo. If multiple applications are running, different users can be associated to different applications.
- Simplify application development: Classic VR development toolkits are usually accessible technical users only. Users with limited programming experience should be provided with a simplified development path to the system. Good middleware should also simplify porting applications across different hybrid reality systems.
- Support a variety of visualization tools and interfaces without requiring developers to reinvent the wheel (transforms, navigation, messaging) for each one of them.
- Modularize interaction metaphors & technologies (navigation, selection, touch support, tracking, etc.), letting developers choose the most appropriate on a per-system or per-application basis.

Realizing no existing VR middleware would support all of these requirements we decided to develop our own, incorporating new capabilities into the Omegalib middleware<sup>14</sup>, and integrating it with SAGE. Omegalib has been designed as the driving software for the OmegaDesk<sup>15</sup>, a hybrid 2D/3D work desk. It was therefore a natural and convenient choice to adapt it to CAVE2.

As Figure 10 shows, Omegalib acts as an abstraction layer between the hybrid reality hardware and user applications. At the back end, Omegalib uses Equalizer to drive the display system. Equalizer is a toolkit for scalable parallel rendering based on OpenGL. It allows users to develop scalable graphics applications for a wide range of systems ranging from large distributed visualization clusters and multi-processor multipipe graphics systems to single-processor single-pipe desktop machines<sup>16</sup>. To allow for full flexibility, Equalizer uses an advanced and verbose syntax, making it complex and error prone to write or maintain configuration files. Omegalib solves this issue by describing display system geometry, cluster configuration and input capabilities through a much compact system description file. When applications launch, the Omegalib runtime system transparently generates the required extended configuration files and forwards them to Equalizer, which then proceeds to initialize and setup the display system.

The Equalizer rendering layer can also be configured to render offscreen, and forward pixels to a SAGE instance running in CAVE2. This way, Omegalib provides two rendering paths: a standard, fullscreen mode controlling the entire CAVE2 (or a fixed subset of it), and a windowed mode forwarding mono or stereo pixel streams to SAGE, which then acts as a compositing window manager.

At the front end, Omegalib integrates several APIs and visualization tools through the use of an abstracted scene graph and pluggable render passes. The approach is similar to the one used by VR Juggler<sup>17</sup> but operates at a higher level, allowing to modularize several common interaction tasks like navigation, picking and object manipulation, turning them into *interaction objects*, that act on the abstracted scene graph. This is analog to the concept of *essential geometry* in graphical user interfaces. Essential geometry offers a link between concrete graphics and higher level interaction semantics.

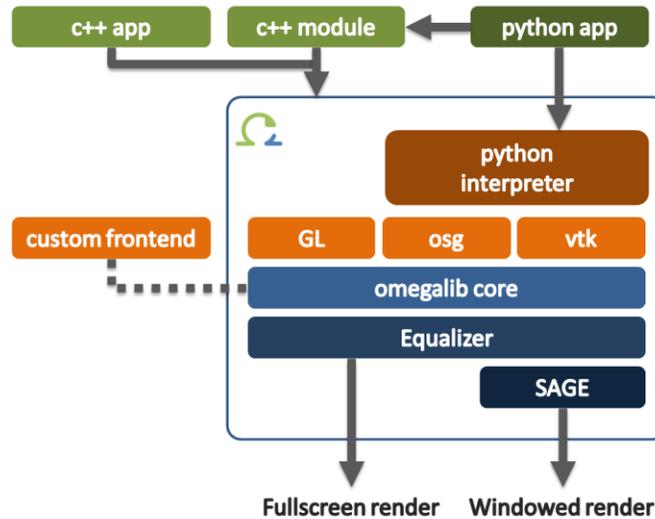


Figure 10. The Omegalib middleware stack.

### 5.1 Python interface

CAVE2 application developers can create native Omegalib applications using a C++ API, or they can use its embedded python interpreter to execute python scripts. Applications are easier to implement through scripting, and require less technical knowledge than C++ programming. An added advantage of an embedded interpreter is that CAVE2 users can execute script statements at runtime, modifying the virtual world on the fly or testing changes to their application without having to restart it. Python commands can also be received from web interfaces or remote Omegalib instances, speeding up the development of collaboration tools.

### 5.2 Available frontends

Omegalib applications can be developed using pure OpenGL: this is particularly useful when porting legacy code to CAVE2. Alternatively, developers can use standard APIs like Open Scene Graph<sup>18</sup> or the Visualization Toolkit (VTK)<sup>19</sup>. This makes it possible to load VTK pipelines as python scripts into CAVE2. It is also possible to create VTK programs for CAVE2 natively, using the C++ VTK API directly. This makes it extremely easy to build VTK programs for CAVE2 or port legacy pipelines to the system.

### 5.3 Panoptic stereo

As mentioned previously, a hybrid reality middleware should be designed to support groups of co-located users collaborating within the system. Because traditional CAVEs use viewer-centered perspective, one effect that typically occurs is that when the tracked user looks 180 degrees away from a given stereo screen, stereo reverses. While this is acceptable for a single user, given they are no longer looking at the screen, it poses a problem for other audience members. To solve this issue, Omegalib supports *Panoptic stereo*. Techniques similar to Panoptic stereo have been implemented in the past to support multi-user stereo without generating multiple views<sup>20</sup>. When Panoptic stereo is enabled, the system tracks the main viewer's position in CAVE2, but generates head orientation information based on each display plane normal. This way, the stereo frustum is projected outward to each display panel, preserving stereo separation for each user in the system. An additional benefit of Panoptic stereo is its independence from frame latency: When users move their head around, stereo will look correct regardless of the application frame rate, leading to a much more comfortable viewing experience. This comes at the price of a slightly incorrect overall projection, particularly when users rotate their head sideways. Panoptic stereo can be turned on and off at runtime if needed: in practice, we found Panoptic stereo to work very well in CAVE2, for both large audiences and single users.

### 5.4 Interaction

As the main interaction device within CAVE2, we are using multiple tracked wands, built using off-the-shelf game controllers (PS3 Move, shown in Figure 11). The wands provide 8 digital and 3 analog input channels, plus 6-DOF tracking. As mentioned in section 4.1, input and tracking data from the wands is processed and aggregated by a

dedicated input server machine, and streamed to CAVE2 applications using VRPN or the omicron protocol<sup>21</sup>. We experimented with additional interaction device designs: for instance, we successfully prototyped an intuitive object manipulation interface called CAVESphere, based on a plexiglass sphere with embedded tracking markers (Figure 12). The sphere also contains a camera and microphone, which we can use to detect finger taps on the sphere surface, providing simple touch input. We also integrated a dynamic HTML5 web interface (called Porthole<sup>22</sup>) to CAVE2. This interface lets users control applications using a laptop, tablet or smartphone. Moreover, the web interface can receive pixel streams from the CAVE2 display cluster, and present additional data on the user device, as shown in Figure 13.



Figure 11. The CAVE2 wand controller



Figure 12. Our prototype CAVESphere interaction device.

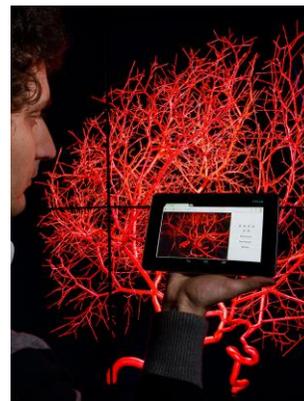


Figure 13. The Porthole interface running on a tablet device, displaying vasculature data rendered on the CAVE2 cluster.

## 6. CONCLUSION

This paper presented the CAVE2 Hybrid Reality Environment. Hybrid Reality Environments represent the intersection of best-in-class display walls and best-in-class virtual reality environments, enabling researchers to analyze high resolution 2D and 3D data through a unified visualization instrument.

To implement CAVE2, we successfully addressed several challenges related to display system geometry, off-axis stereo performance, tracking and audio system design. We also envisioned the need of a custom software layer to drive HREs. Our implementation of such software layer is based on improving and merging Omegalib and SAGE.

One aim of future research is to further improve the software driving CAVE2. Simplifying interaction transitions between 2D and 3D content and efficiently dealing with multiple 3D applications and users represent significant objectives of this research.

We also plan adapt popular game development toolkits like Unity3d to CAVE2, to further extend and simplify access to the system.

As a further research direction, we intend add a ring of Kinects around the bottom or from the top-down of CAVE2 to enable us to explore purely gestural interactions. This coupled with our recent work in natural language translation for automatically converting speech to visualizations<sup>23</sup>, could provide a highly naturalistic way for users to create and manipulate complex visualizations in CAVE2 without being encumbered by buttons, controllers and menuing interfaces.

### 6.1 Acknowledgements

The authors would like to thank Khairi Reda and Thomas Marrinan for their contributions toward CAVE2 application development.

The construction of the CAVE2 system was supported by the National Science Foundation (award # CNS-0959053) and the Department of Energy (award # DE-SC005067). SAGE was supported by the National Science Foundation (award # OCI-0943559)

## REFERENCES

- [1] Cruz-Neira, C. *et al.* Scientists in wonderland: A report on visualization applications in the CAVE virtual reality environment. in *Proceedings of 1993 IEEE Research Properties in Virtual Reality Symposium* 59–66 (IEEE Comput. Soc. Press, 1993). doi:10.1109/VRAIS.1993.378262
- [2] Cruz-Neira, C., Sandin, D. J., DeFanti, T. A., Kenyon, R. V. & Hart, J. C. The CAVE: audio visual experience automatic virtual environment. *Communications of the ACM* **35**, 64–72 (1992).
- [3] DeFanti, T. a. *et al.* The future of the CAVE. *Central European Journal of Engineering* **1**, 16–37 (2010).
- [4] DeFanti, T. a. *et al.* The StarCAVE, a third-generation CAVE and virtual reality OptIPortal. *Future Generation Computer Systems* **25**, 169–178 (2009).
- [5] Leigh, J. *et al.* Scalable Resolution Display Walls. *Proceedings of the IEEE* 1–15 (2012). doi:10.1109/JPROC.2012.2191609
- [6] Yost, B., Haciahetoglu, Y. & North, C. Beyond visual acuity: the perceptual scalability of information visualizations for large displays. *Proceedings of the SIGCHI conference on Human Factors* (2007).
- [7] Ball, R. & North, C. Effects of tiled high-resolution display on basic visualization and navigation tasks. *CHI'05 extended abstracts on Human factors in Computing* (2005).
- [8] Andrews, C., Endert, A. & North, C. Space to think: large high-resolution displays for sensemaking. *Proceedings of the 28th international conference on Human factors in computing systems* (2010).
- [9] Renambot, L., Jeong, B. & Jagodic, R. Collaborative visualization using high-resolution tiled displays. *ACM CHI Workshop on Information Visualization Interaction Techniques for Collaboration Across Multiple Displays* (2006).
- [10] Jeong, B. *et al.* Scalable Graphics Architecture for High-Resolution Displays. *Image (Rochester, N.Y.)* 1–4 (2005).
- [11] Toennis, M., Benzina, A. & Klinker, G. Utilizing Consumer 3D TV Hardware for Flexible and Reconfigurable Visualization Systems. *ar.in.tum.de* at <<http://ar.in.tum.de/pub/toennis2011frave/toennis2011frave.pdf>>
- [12] Andrews, C., Endert, A., Yost, B. & North, C. Information visualization on large, high-resolution displays: Issues, challenges, and opportunities. *Information Visualization* (2011). at <<http://ivi.sagepub.com/content/10/4/341.short>>
- [13] Weissman, M. & Woods, A. A simple method for measuring crosstalk in stereoscopic displays. *IS&T/SPIE* ... (2011). at <<http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=730525>>
- [14] Febretti, A., Nishimoto, A., Mateevitsi, V. & Chau, D. omegalib. at <<https://code.google.com/p/omegalib/>>
- [15] Febretti, A. *et al.* The OmegaDesk : Towards A Hybrid 2D & 3D Work Desk. *Virtual Reality* 13–23 (2011). at <<http://www.springerlink.com/index/P06341621767696T.pdf>>
- [16] Eilemann, S., Makhinya, M. & Pajarola, R. Equalizer: a scalable parallel rendering framework. *IEEE transactions on visualization and computer graphics* **15**, 436–52 (2009).
- [17] Bierbaum, A. & Just, C. VR Juggler: A virtual platform for virtual reality application development. *Virtual Reality, 2001.* ... (2001). at <[http://ieeexplore.ieee.org/xpls/abs\\_all.jsp?arnumber=913774](http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=913774)>
- [18] Osfield, R. & Burns, D. Open Scene Graph. (2004). at <<http://www.openscenegraph.com>>
- [19] The Visualization Toolkit. at <<http://www.vtk.org/>>
- [20] Holtkämper, T. Co-located collaborative use of virtual environments. *Proceedings AAPG Annual Convention and Exhibition.* (2007).
- [21] Febretti, A., Nishimoto, A. & Mateevitsi, V. The omicron SDK. at <<https://code.google.com/p/omicron-sdk/>>
- [22] Leigh, J. & Johnson, A. Porthole : A Decoupled HTML5 Interface Generator For Virtual Environments. (2012).
- [23] Sun, Y., Leigh, J., Johnson, A. & Lee, S. Articulate: A semi-automated model for translating Natural Language queries into meaningful visualizations. *Proceedings of 10th International Symposium of Smart Graphics (SG), Lecture Notes in Computer Science* (2010). at <<http://www.springerlink.com/index/A1R074M271472574.pdf>>