The AlloSphere: Immersive Multimedia for Scientific Discovery and Artistic Exploration

Xavier Amatriain, JoAnn Kuchera-Morin, Tobias Hollerer, and Stephen Travis Pope
University of California, Santa Barbara

The AlloSphere is a spherical space in which immersive, virtual environments allow users to explore large-scale data sets through multimodal, interactive media.

We designed the AlloSphere a novel environment that allows for synthesis, manipulation, and analysis of large-scale data sets to enable research in science and art. Scientifically, the AlloSphere can help provide insight on environments into which the body cannot venture. Artistically, the AlloSphere can serve as an instrument for creating and performing new works and developing new modes of entertainment, fusing art, architecture, science, music, media, games, and cinema.

The AlloSphere is situated at one corner of the California Nanosystems Institute building at the University of California, Santa Barbara (see Figure 1), and is surrounded by several associated labs for visual and audio computing, robotics, interactive visualization, world modeling, and media post-production. The building, which represents the culmination of five years of research, design, and construction, is a three-story-high cube.

The AlloSphere space contains a spherical screen that is 10 meters in diameter (see Figure 2). The sphere environment integrates several visual, audio, interactive, and immersive components and is one of the largest immersive instruments in the world, capable of accommodating up to 30 people on a bridge suspended across the middle. Once fully equipped, the AlloSphere will have several additional features, such as true 3D projection of video and audio data, in addition to interactive-sensing and camera-tracking capabilities.

The AlloSphere consists of an empty cube that is treated with extensive sound-absorption material, making it one of the largest near-to-anechoic chambers in the world. In a perfect anechoic space, sound waves aren't reflected in any of its surfaces, yielding a neutral or dead space from an acoustic perspective. Standing inside this chamber are two hemispheres constructed of perforated aluminum designed to be optically opaque and acoustically transparent. Figure 3 shows a detailed drawing of the AlloSphere.

Currently, we are equipping the AlloSphere with 14 high-resolution video projectors mounted below the bridge and around the seam between the two hemispheres to project video on the entire inner surface. The AlloSphere's loudspeaker array is suspended behind the screen, hung from the steel infrastructure in rings of varying density.

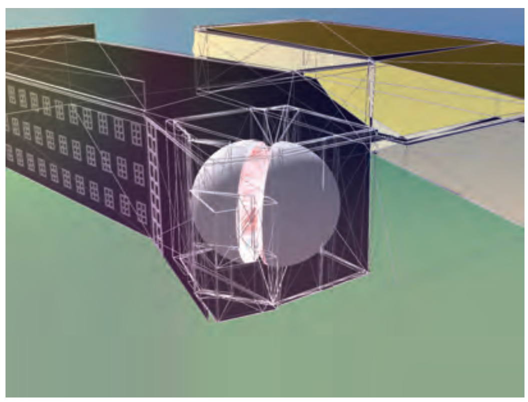

Figure 1. Virtual model of the AlloSphere space in the California Nanosystems Institute building at the University of California, Santa Barbara. (Image used with permission of Springer Science+Business Media.)

1070-986X/09/$25.00 © 2009 IEEE

Published by the IEEE Computer Society

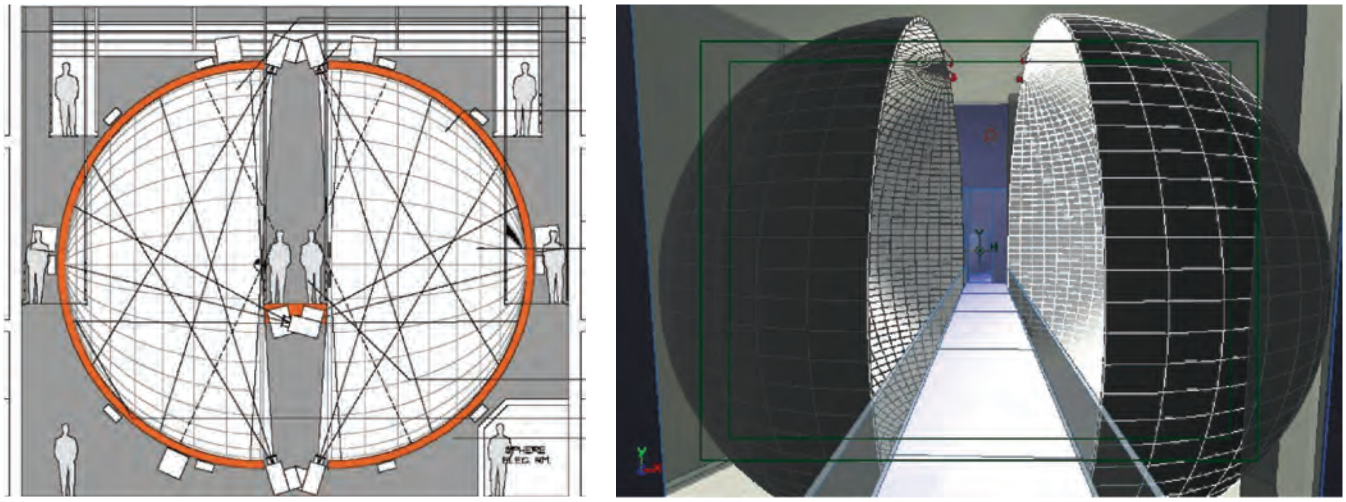

(a) (b)

Figure 2. The AlloSphere: (a) schematic view and (b) view provided by our own simulation. For visibility, we have omitted the screen segments above the bridge.

Once fully equipped and operational, the AlloSphere will be one of the largest immersive instruments in existence, offering several features that make it unique.

Beyond 3D immersion

The AlloSphere adds a new data point to the list of the world's largest and most precise immersive 3D environments, such as the newly upgraded Virtual Reality Applications Center at Iowa State University, the Fakespace Flext installation at Los Alamos National Laboratory, the Samuel Oschin Planetarium at the Griffith Observatory in Los Angeles, the Denver Museum of Nature & Science Gates Planetarium dome, and the Louisiana Immersive Technologies Enterprise center.

With its unique spherical shape, its high resolution, and its immersive multimodal capabilities, the AlloSphere represents a step beyond several capabilities of existing virtual environments, such as CAVE.[1,2] For example, the AlloSphere enables seamless stereo-optic 3D projection and doesn't distort the projected content due to room geometry. Stereo-optic 3D is possible for a large set of AlloSphere users because the audio and stereovision sweet spot area is large, but is restricted to the bridge.

There are several other technical and functional innovations of the instrument in comparison to existing immersive environments. The AlloSphere is a spherical environment with a full 4π steradians of stereo visual information. In this sense, it resembles state-of-the-art visual systems such as the CyberDome[3] but on a different scale. The AlloSphere surround-view design provides a sense of immersion with little encumbrance and limited distortion away from the center of projection or tracked user. Generally speaking, spherical systems enhance subjective feelings of immersion, naturalness, depth, and realism.[4]

In addition, its size makes it possible for several users (up to 30 people on the bridge) to be

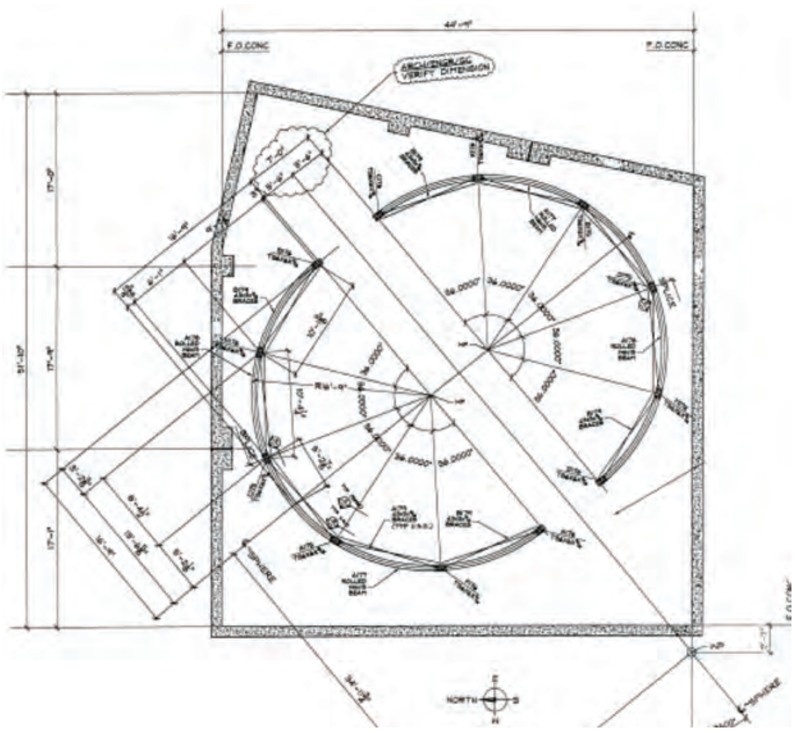

Figure 3. Horizontal section of the AlloSphere showing the nonparallel and acoustically treated surfaces surrounding the sphere. The AlloSphere is not a perfect sphere but rather two hemispheres separated by the bridge.

Figure 4. A large number of users can fit into the AlloSphere bridge.

collaborating in the environment (see Figure 4). For certain content, there is no need for viewpoint adaptation because of the separation of users from the projection screen, as long as users located on one of the bridge's ends don't focus on the screen surface closest to them. This phenomenon is similar to the Imax effect, in which a large number of users can view good-quality 3D images as long as they are far enough from the screen.

Moreover, the AlloSphere combines state-of-the-art techniques both on virtual audio and visual data spatialization. The spherical screen is placed in a carefully designed near-to-anechoic chamber and is perforated to enable spatialized audio from a speaker system behind it. There is extensive evidence of how combined audio-visual information can help information understanding,[5] although most existing immersive environments focus on presenting visual data.

Lastly, the space was designed not only as a multimodal interaction environment[6] consisting of camera-tracking systems, audio recognition, and sensor networks, but also as a pristine scientific instrument. Although the space is not fully equipped at this point, we have been experimenting with different equipment, system configurations, and applications. Because the AlloSphere is a research instrument rather than a performance space, it will be an evolving prototype rather than a fixed installation. We envision the instrument as an open framework that undergoes constant refinement with major releases signaling major increments in functionality.

Multimodal design

The diagram depicted in Figure 5, which illustrates the main subsystems and components in the AlloSphere, is a simplified view of the integrated multimedia and multimodal design. A typical multimodal AlloSphere application integrates services running on multiple hosts on the LAN. These hosts implement a distributed system consisting of the following elements:

- input sensing (camera, sensor, and microphone);
- gesture recognition and control mapping;

Figure 5. The AlloSphere components and subsystems: visual, audio, control, and sensing.

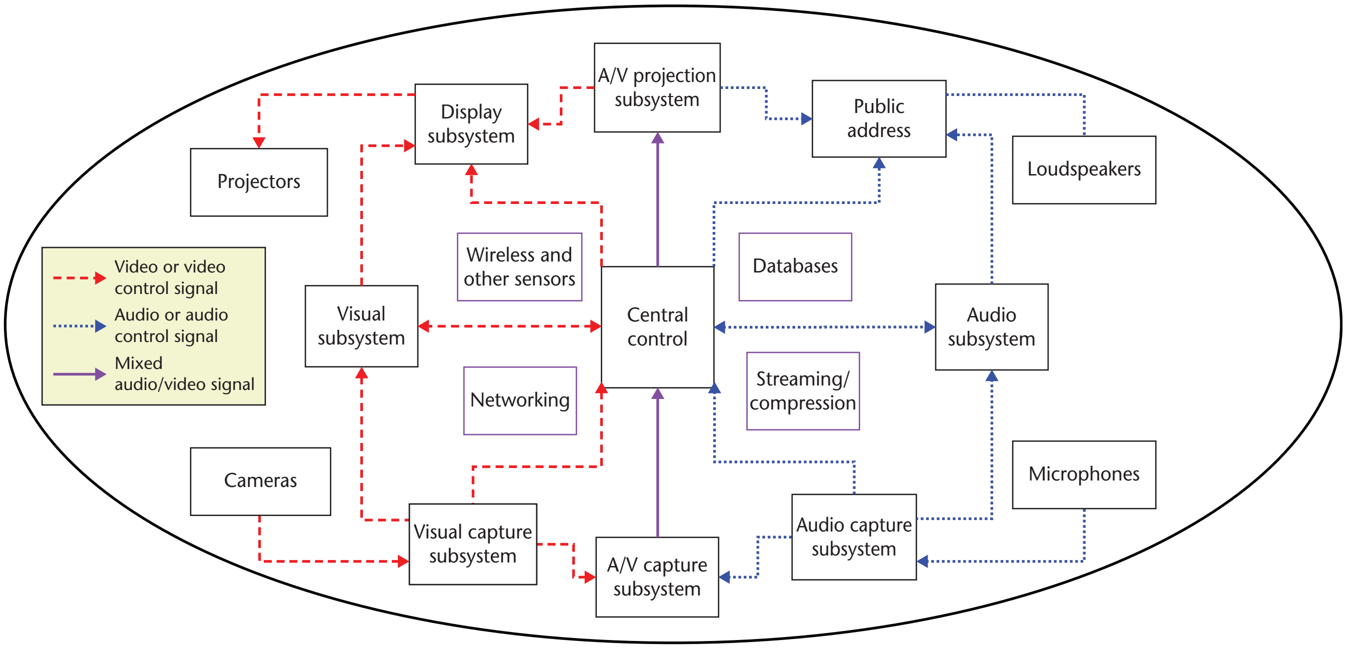

- interface to a remote application (scientific, numerical, simulation, and data mining);
- back-end processing (data and content access);
- output media mapping (visualization and sonification); and
- audiovisual rendering and projection management.

These requirements confirm that off-the-shelf computing and interface solutions are inadequate.

AlloSphere applications require not only a server farm dedicated to video and audio processing, but also a low-latency interconnection fabric so that data can be processed on multiple computers in real time. In addition AlloSphere applications require integration middleware and an application server that lets users manipulate the system and their data flexibly and meaningfully.

Input sensing is an important component of AlloSphere applications. Currently, users can interact with the AlloSphere through custom-built devices, camera-based infrared tracking, game controllers, and touch sensors on the bridge's rails. The coupling of infrared tracking with control devices allows users' positions to be monitored as they traverse the bridge while also allowing them to manipulate virtual objects in 3D space. We use the Precision Position Tracker system from WorldViz to determine user position, and Logitech game controllers and Wiimotes for user interaction.

In the final design, we plan to have a multimodal human–computer interaction subsystem with real-time vision and camera tracking, real-time audio capture and tracking, and a sensor network consisting of wireless sensors and input devices as well as presence and activity detectors.

The computation system will consist of a network of distributed computational nodes, with communication between processes accomplished through standards such as the Message-Passing Interface (MPI)[7] and the Open Sound Control (OSC)[8] protocol. The AlloSphere network must host this kind of standard message-passing along with multimedia, multichannel streaming.

In light of these requirements, we are still discussing the suitability of Gigabit Ethernet or Myrinet versus other proprietary technologies. In our first prototypes that used Chromium[9] to distribute rendering primitives, Gigabit Ethernet proved sufficient, but our projections show that the limitations of Gigabit Ethernet will become a bottleneck for the complete system, especially when using a distributed rendering solution to stream highly dynamic visuals.

Visual subsystem

The main requirements for the AlloSphere visual subsystem are fixed by the constraints of the building and by our desired quality targets. The sphere screen area is 320 square meters. For good performance, we need a minimum of three arc minutes of angular resolution. In terms of light level, we need 50 trolands, although we can limit active stereo to 30 trolands. With these requirements, we have designed a projection system consisting of 14 active stereo projectors that are capable of a maximum 3,000 lumens and SXGA+ resolution.

For the simulations (see Figure 6, next page), we developed our own environment using Oliver Kreylos' Vrui VR Toolkit.[10] This simulator helped us design projector location and coverage in the AlloSphere and measure the effect of the projector characteristics. For on-site tests, we started with a single active stereo projector (see Figure 7) and brought the visual system up to a four-projector configuration. For these tests, we used a range of projectors, moving from 2,000 to 10,000 lumens and including accessories such as fish-eye lenses.

Image brightness

One of the important design goals for the AlloSphere is to make it user-friendly and usable for extended periods. Unacceptably low levels of brightness cause eye fatigue or severely restrict the type of content that we can display. Without considering the stereo requirement, the projected system yields 42,000 lumens and a screen luminance (full white) of 9.26 candela per square meter (cd/m). In comparison, the luminance of a good-quality, multimedia Dome is recommended to be between 0.686 and 5.145 cd/m.[11]

According to our simulations and on-site tests, 42,000 lumens of input flux produce

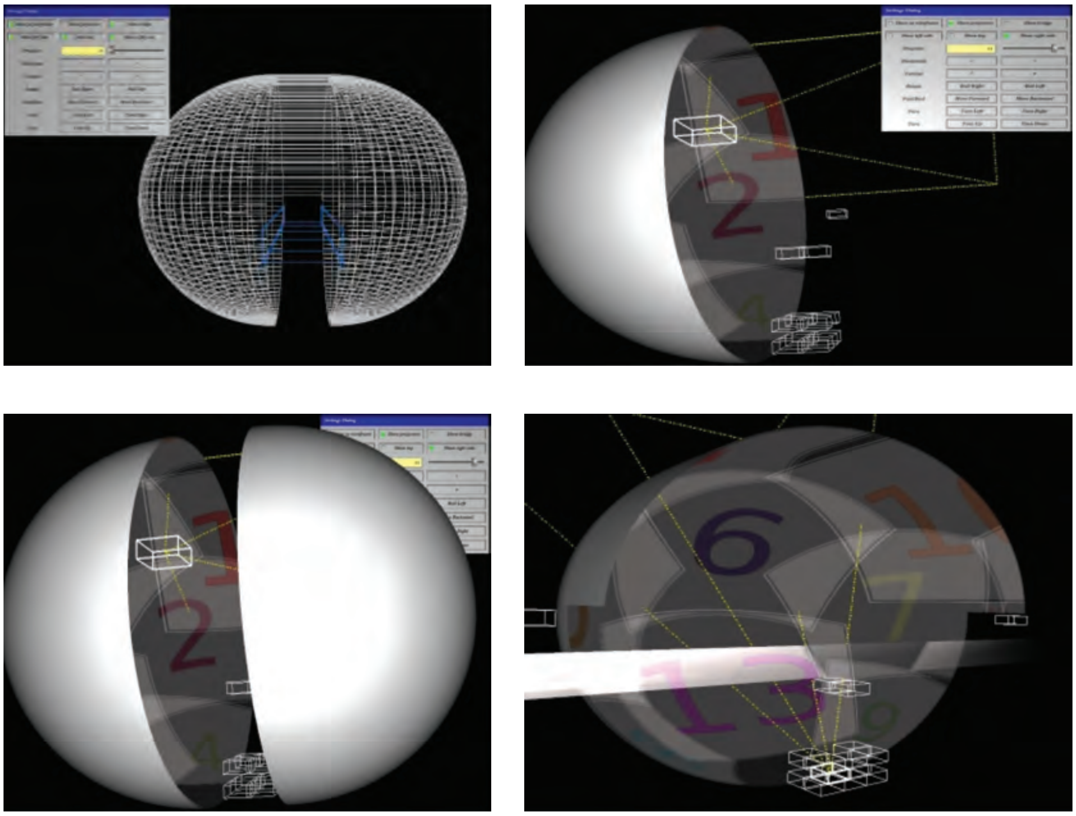

Figure 6. Different views from the simulator for placing projectors and experimenting with coverage models. With this tool, we can simulate projector models with different coverage and experiment with positioning and tiling.

close to optimal results. Besides, augmenting the light flux above this level has several undesired effects, namely cross-reflection and ghosting due to back-reflection.

Stereoscopic display

The performance of the AlloSphere in stereo mode depends on the choice of the stereo display technology. Passive, polarization-based methods are ill suited to the AlloSphere due to the surround nature of the double-concave screen and the nonpolarization-preserving material. While light losses from stereo projection are substantial, the design requirement for stereo-projection brightness is 30 trolands at 50 percent RGB.

Stereo-projection mode falls below the theoretical eyestrain threshold with our projected 42,000 lumens total. Nevertheless, our field studies indicate that this level of stereo brightness is still perceived as high quality. In addition, it allows continuous working times of more than 60 minutes. It's worth noting that the main cause of eyestrain in stereo mode is active shuttering, which is not related to luminance.

Figure 7. Testing the AlloSphere projection with a single stereo projector with backlighting on to show the screen structure. (Image used with permission of Springer Science+Business Media.)

Contrast ratio

Contrast loss due to diffused scattering represented a serious problem for the projection design. Lowering the screen gain reduces the secondary reflections proportionally to the square of the screen-paint gain and translates to a corresponding increase in image contrast. However, doing so has the unwanted effect of requiring more input light flux, which increases back reflections, heat, and noise.

We determined the screen gain after several tests and simulations, taking into account experiences in similar venues (mostly state-of-the-art planetariums such as the Hayden in New York or the Gates Planetarium in the

Denver Museum of Natural History). The screen paint has an field-of-view-averaged gain of 0.12 with a peak value of 0.24, which will, according to the simulation, produce a maximum contrast ratio of about 20:1 for images with 50 percent total light-flux input.

Screen resolution

The AlloSphere's visual resolution is a function of the total number of pixels available and the projector overlap factor, which we calculated to be 1.7. The spatial acuity for 20/20 eyesight is 30 line pairs per degree, which is the average spatial acuity in regular conditions because spatial resolution is a function of both contrast ratio and pupil size. Nevertheless, users perceive resolutions as low as three arc minutes to be high quality unless a better reference point is available.

By using the center of the hemisphere as a common viewpoint, we can infer the number of pixels for a given resolution independently of the screen size or diameter. A three-arc-minute resolution requires 20 million pixels spread over a full sphere. Our target configuration of 14 projectors has 19.2 megapixels, a number that corresponds to our desired resolution.

Image warping and blending

The AlloSphere projection system requires us to warp and blend images from multiple projectors to create the illusion of a seamless image. We warp and blend images on the video projectors and on the graphics cards in the image-generation system.

Most modern simulation-oriented video projectors support some form of warping and blending. Doing the warping and blending on the projectors is convenient, and often results in the best image quality. However, a negative side effect of this technique is that the projector must buffer an entire frame before being able to process it. Another negative aspect of this technique is that projector-based warping and blending is encoded in proprietary software that is hard to access and extend.

In the case of doing warping and blending on the graphics cards, the process happens after rendering the frame buffer. However, doing so consumes resources that otherwise could be used to render polygons. For this reason, we prefer specialized hardware. But such hardware is costly and proprietary, and makes calibration procedures more complex. The benefit of computer-based warping is reduced latency; the video projector doesn't need to buffer an entire frame before displaying it.

In the AlloSphere, we decided to start with projector-side warping and blending. The decision fulfilled many, but not all, of the AlloSphere requirements. Moreover, we are working on extending and adapting existing solutions[12] for a full spherical surface.

Latency and frame rate

For the AlloSphere system design, we had to consider all latencies occurring from after the start of rendering to when the image appears on the screen. Research indicates that unpleasant side effects appear above 120 milliseconds total system latency for VR applications. Below 120 ms, the lower the latency, the more accurate and stress-free the interaction becomes.

In general, a total system delay of 50 ms is considered to be state-of-the-art for systems like the AlloSphere. Furthermore, to deliver flicker-free stereo, we must guarantee a frame rate of at least 100 Hz.

Image generation and rendering

To meet our requirements, we needed an image-generation system capable of producing 20 million pixels through 14 channels and supporting resolutions of at least XVGA+, with active stereo support as well as frame-lock capabilities for synchronizing all channels. To meet these requirements, we designed a rendering cluster consisting of seven Hewlett Packard 9400 workstations, each of which is equipped with an Nvidia FX-5500 graphics card and a G-sync card for frame locking.

To generate large, multitile immersive displays, there are several techniques and tools available.[13] However, the AlloSphere design poses some unique problems that are best addressed through research. For example, the tiles are irregularly shaped and curved and the projection screen is a continuous quasisphere. In addition, because the AlloSphere is not a perfect sphere, conventional warping solutions aren't directly applicable. Moreover, the projection must allow for active stereo to work well in most fields of view, and the system must be flexible enough to adapt to legacy applications. Finally, our design vision requires a middleware layer that can run any OpenGL application, even when source code access isn't available.

We are addressing some of these requirements in our current research. For those applications in which no source code is available or in which viewpoint information isn't relevant for rendering a convincing 3D scene, we are using a distributed-rendering solution based on Chromium. In these cases, a single master runs the application and performs early rendering, offloading the rendering of the specific viewpoints to appropriate slaves.

In those applications for which source code is available or that require complete viewpoint-dependent rendering, we use an approach that is based on distributing the whole application. In these cases, the master manages the application state and processes user input from the interface. The slaves perform the rendering, receiving information about the application state and their particular viewpoint and rendering tile. This approach is similar to that used by VR libraries, such as Syzygy[14] or VRJuggler.[15]

Audio subsystem

One of the unique features in the AlloSphere is that it offers symmetrical immersion through video and audio capabilities. Designing the audio software and hardware subsystems has taken several years because our goal has been to build an immersive interface that provides sense-limited resolution in both the audio and visual domains. This means that the spatial resolution for the audio must allow us to place virtual sound sources at arbitrary points in the AlloSphere. And the system must allow us to simulate the acoustics of measured spaces with a high degree of accuracy.

To provide for ear-limited dynamic, frequency, and spatial extent and resolution, we require the system to be able to reproduce in excess of 100 decibels (dB) near the center of the sphere, to have acceptable low- and high-frequency extension (3 dB below 40 Hz and above 18 kHz), and to provide spatial resolution on the order of three degrees in the horizontal plane and 10 degrees in elevation. To provide high-fidelity playback, we require an effective signal-to-noise ratio that exceeds 80 dB, with a useful dynamic range of more than 90 dB.

To be useful for data sonification[16] and as a music performance space, the decay time of the AlloSphere must be less than 0.75 seconds from 100 Hz to 10 kHz. We have carried out and published detailed measurements of the finished AlloSphere space, its treatment, and the projection screen's acoustical properties.[17] We used several synthetic and explosive sources and careful microphone placement to ascertain the effects of having the aluminum sphere in our anechoic chamber. The space's wideband time of 0.45 seconds means that we can dissipate and absorb the energy we introduce into the sphere, and the mirrored-microphone measurements confirm that the sphere itself is acoustically inert.

Spatial sound processing

There are three techniques for spatial sound reproduction used in current state-of-the-art systems: vector-base amplitude panning (VBAP),[18] ambisonic representation and processing,[19] and wave field synthesis (WFS).[20,21] Each of these spatialization techniques provides a different set of advantages and presents unique scalability and complexity challenges when scaling to a large number of speakers or virtual sources.

VBAP is a signal-processing technique by which a sound source can be located in the space by setting the balance of the audio signal sent to each of several speakers, which are assumed to be equidistant from the listener. The technique's main drawbacks are that it doesn't include a model of a direct distance cue and doesn't support sound sources inside the loudspeaker sphere. Members of our research group implemented a system in which the user can move and direct several independent sound sources using a data glove input device, and play back sound files or streaming sound sources through VBAP using a variable number and layout of loudspeakers specified in a configuration file.[22]

Ambisonics is used to synthesize a spatial sound field by encoding sound sources and their geometry, then decoding them using the ambisonic transform, a multichannel representation of spatial sound fields based on spherical harmonics.[23] One of the advantages of this technique is that it scales well to a large number of moving sources. However, as with VBAP, sound-source positions inside the loudspeaker ring cannot be recreated directly. Graduate researchers from our group implemented higher-order ambisonic processing and decoding.[24,25] To adapt ambisonics to a navigable environment such as the AlloSphere, we implemented multiple distance cues and a source radiation pattern simulation.

Finally, WFS recreates wave fronts with large arrays of loudspeakers by building on the Huygens principle of superposition. Although this technique can produce detailed sound fields, in its current implementations it has two drawbacks. It generally requires offline computation that limits its usefulness in virtual environments, and it doesn't natively allow for speaker configurations in more than two dimensions. Members of our team developed a different method in which WFS filters are calculated in real time with a small computational overhead.[26] In addition, to obtain high-performance 3D effects, we can combine WFS in the horizontal plane with any other technique for the vertical plane using the framework presented later in this article.

The AlloSphere supports the use of any combination of these techniques for sound spatialization. To facilitate these options, we developed a generic software framework that can combine different techniques and speaker layouts with little effort.[27] We based this framework on our own Metamodel for Multimedia Processing Systems[28] and implemented it on top of the Create Signal Library.[29] The framework offers interface layers with increasing levels of complexity and flexibility. In the simplest interface, the user is responsible only for determining the sound position and providing the raw audio material; everything else is determined automatically. However, by using the other layers, the user can determine details such as the spatialization algorithm or the filters.

We have evaluated the scalability of our implementations of the three spatialization techniques according to a multidimensional load model, characterizing performance with several resource metrics, including processing load, memory footprint, bandwidth, and so forth.[30] Each of the algorithms has different load-condition profiles where they can scale quite well, and different modes requiring increasing CPU, RAM, or bandwidth resources.

Speaker system

It was a major project to determine the optimal speaker placement and density because the loudspeaker count and configuration had to support all of the spatial audio techniques. Our design consists of 425 to 512 speakers arranged in several rings around the upper and lower hemispheres, with accommodations at the seams between the desired spacing and the requirements of the support structure.

Our design requires placing densely packed circular rings of speaker drivers running just above and below the equator (on the order of 250 channels, side-by-side), and two smaller and lower-density rings concentrically above and below the equator. The main loudspeakers have limited low-frequency extension in the range of 200 to 300 Hz. To project frequencies below this, we mounted subwoofers on the underside of the bridge. At this moment, because of timing and construction constraints, we installed a prototype system with only 32 speakers installed along the three different rings and two subwoofers under the bridge.

We connected the speakers to the computer via FireWire audio interfaces that support 16 channels. The eventual audio output hardware will consist of several synchronized servers on a switched network, each server supporting multiple 64-channel FireWire or optical interfaces to send audio to the distributed speaker banks.

Test applications

Our goal for the AlloSphere is to have content and demand driving its technological development just as it has driven its design. For this reason, specific application areas are essential in the development of the instrument because they define the functional framework in which the AlloSphere will be used. In the first prototype, we set up an environment consisting of the following elements: four active stereo projectors, two rendering workstations, one application manager, two 16-channel FireWire cards, 32 speakers, one subwoofer, and a Precision Position Tracker system from WorldViz. For user interaction, we used Logitech controllers, Wiimotes, and several custom-developed wireless interfaces.

The research projects described here use this prototype system to test the functionality and prove the validity of the instrument design. All of the projects are being developed by teams of scientists, engineers, and media artists, allowing the scientist to perceive their data in different ways and offering media artists the possibility of converting abstract models and data sets into pieces of art. As a result, the projects offer the option of presenting hard science problems to the general public.

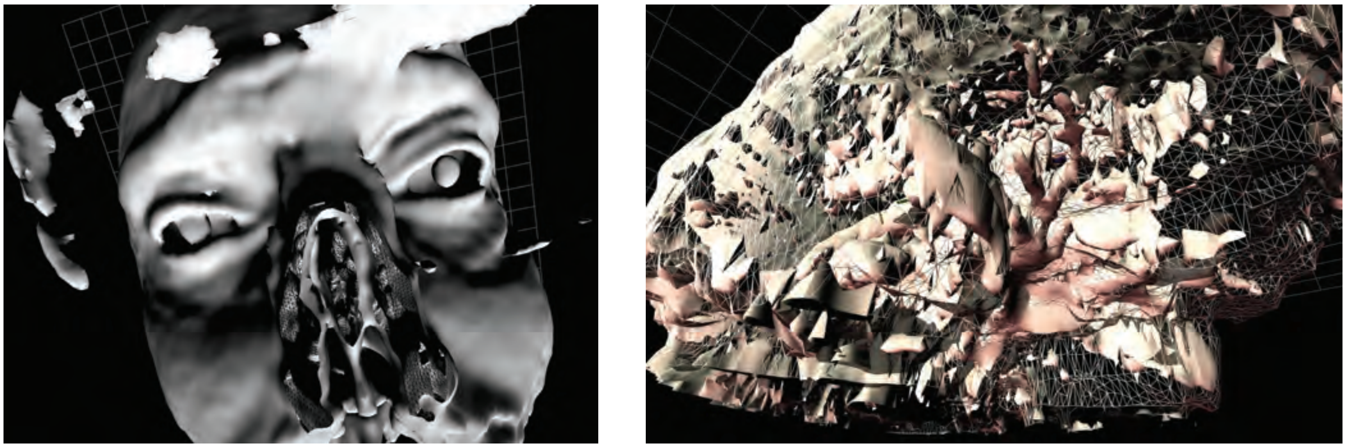

(a) **(b)**

Figure 8. Two screen captures of the AlloBrain interactive recreation of the human brain from functional MRI data. In (a) most tissue layers are activated to allow for visualization of realistic facial expressions. In (b) the outer layers are faded to allow for inner navigation into the brain.

The AlloBrain, a project under the direction of artist Marcos Novak, reconstructs an interactive 3D model of a human brain from macroscopic, organic data sets derived from functional MRI data from Novak's brain (see Figure 8). The current model contains several layers of tissue blood flow and consists of an interactive environment where twelve agents navigate the space and gather information to deliver to the researchers. The systems are stereo-optically displayed and controlled by two wireless input devices that feature custom electronics and several sensor technologies.

The first controller allows the user to navigate the space using six degrees of freedom. The second contains 12 buttons that command the 12 agents (see Figure 9) and allows moving the ambient sounds spatially around the sphere. Its shape is based on the hyperdodecahecron, a four-dimensional geometrical polytope. The final object represents its shadow projected into three dimensions. We developed the shape using procedural modeling techniques and constructed it with a 3D printer capable of building solid objects.

Figure 9. A researcher interacting with the AlloBrain through a custom made wireless controller.

Using these controls and the immersive qualities of the AlloSphere, neuroscientists have explained the structure of the brain to a varied audience. This virtual interactive prototype, currently our most mature project, illustrates some of the key aspects of the AlloSphere and has been featured as an artwork in several exhibitions. Also, Novak's work in the AlloSphere has been featured in several arts and architecture forums and is being studied by digital art researchers. In addition, the AlloBrain has been showcased in the AlloSphere to the general public.

In another project, we are developing an immersive and interactive software simulation of nanoscale devices and structures, with atom-level visualization of those structures implemented on the projection dome (see Figure 10). For this project, we are implementing our scientific partners' algorithms and models, including molecular dynamics on high-end GPUs that allow for enough speed to provide real-time interaction with the simulation.

Another project, the quantum visualization and sonification project, is lead by composer and digital artist JoAnn Kuchera-Morin. This project relies on an audio synthesis model of electronic measurements on a quantum dot.

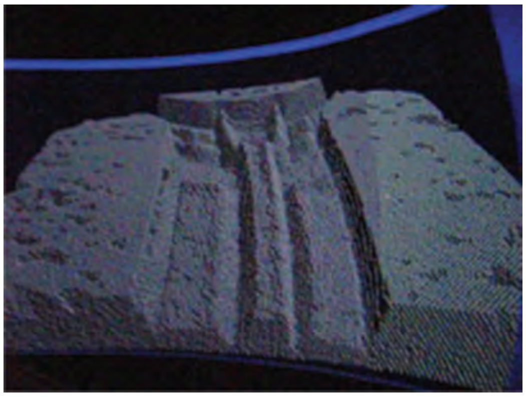

Figure 10. Rendering of a silicon nanostructure in real time as shown in the AlloSphere.

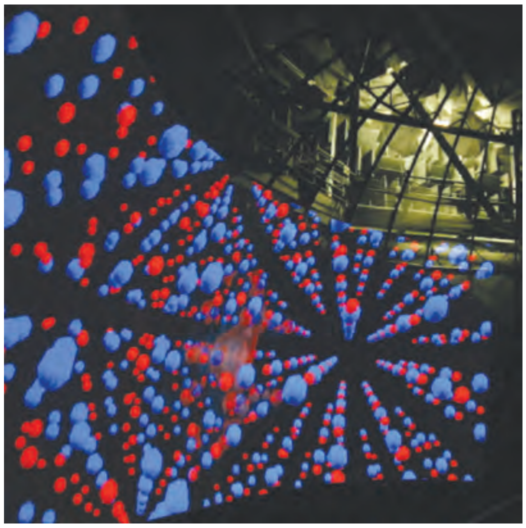

Figure 11. Interactive visualization of atomic bonds for alternative fuel sources as shown in the AlloSphere.

The model is a literal interpretation of experiments undertaken by physicist David Awschalom and his research group in the Center for Spintronics and Quantum Computation (see http://csqc.ucsb.edu). The experiment from which the model is derived is a measurement of coherent electron spin in a quantum dot.

We mapped the mathematical model of the experiment using wavelength as the basis for transposing optical frequencies into audio. We derived the visualizations directly and literally from the audio output and represented it with animation. Conceptually, this project follows in the evolution of sound generation from earlier developments in musical instrumentation that applied electronic pickups on acoustic instruments to analog signal generation and digital synthesis.

In another project, also under the artistic direction of Kuchera-Morin, we are creating an interactive visualization and multimodal representation of unique atomic bonds for alternative fuel sources. The project is a joint venture with Christopher Van De Walle and the Solid State Lighting and Display Center (see http://ssldc.ucsb.edu). The goal of the project is to create an interactive and artistic installation that offers new insights into hydrogen bond formation.

The piece we created allows users to fly through a 2,000-atom lattice and navigate through the sonification of atomic emission spectrums. We derived all the sonic information from transposing the atomic emission spectrums to audio. We created the visualizations (see Figure 11) from mapping the mathematical calculations of the bond through the Schrodinger equation.

These and other tests are helping us develop an open, generic software infrastructure capable of handling multidisciplinary applications that have common goals. In addition, the tests should facilitate the development of an open-ended computational system for data generation, manipulation, analysis, and representation.

Conclusions

We envision the AlloSphere will become an important instrument in the advancement of fields such as nanotechnology and bioimaging, and will help stress the importance of multimedia in science, engineering, and the arts. The results from our initial tests are feeding back into the prototyping process and are demonstrating the validity of our approach.

The prototype work has given us the opportunity to configure one quarter of the sphere so we can test luminance, colorization, pixel mapping, warping, and blending. Aurally, we have tested a 32-channel, 3D-audio system, implementing several sound-spatialization algorithms. And we have experimented with wireless interactive control. With results from this research, we are scaling up to a complete interactive system consisting of 14 projectors and 500-channel audio for true 3D immersion. **MM**

Acknowledgments

The AlloSphere is the result of the work of a large team. Although the project is directed by JoAnn Kuchera-Morin with Xavier Amatriain as the assistant technical director, our colleagues and students are responsible for the bulk of this work. Graham Wakefield, John Thompson, Lance Putman, and Dan Overholt worked with Marcos Novak on the brain simulation and initial prototypes. Alex Kouznetsov, Jorge Castellanos, Graham Wakefield, Will Wolcott, Florian Hollerwerger, Doug McCoy, and Curtis Roads worked on the audio system. Alex Kouznetsov, Lyuba Kavaleva, and Brent Oster worked on the visual system and simulator. Dennis Adderton and Lance Putnam worked on the quantum visualization project. Basak Alper, Lance Putnam, and Wesley Smith worked on the atom bond project.

References

1. C. Cruz Neira et al., ''The Cave: Audio Visual Ex perience Automatic Virtual Environment,'' *Comm. ACM,* vol. 35, no. 6, 1992, pp. 64 72.
2. J. Ihren and K. Frisch, ''The Fully Immersive Cave,'' *Proc. 3rd Int'l Immersive Projection Technol ogy Workshop,* Springer, 1999, pp. 59 63.
3. N. Shibano et al., ''Cyberdome: PC Clustered Hemispherical Immersive Projection Display,'' *Proc. Int'l Conf. Artificial Reality and Telexistence,* IEEE CS Press, 2003, pp. 1 7.
4. R. Kalawsky, *The Science of Virtual Reality and Virtual Environments,* Addison Wesley, 1993.
5. H. McGurk and T. McDonald, ''Hearing Lips and Seeing Voices,'' *Nature,* vol. 264, 1976, pp. 746 748.
6. A. Malkawi and R. Srinivasan, ''Multimodal Human–Computer Interaction for Immersive Visualization: Integrating Speech Gesture Recog nitions and Augmented Reality for Indoor Envi ronments,'' *Proc. 7th IASTED Conf. Computer Graphics and Imaging,* ACTA Press, 2004, pp. 171 175.
7. P.B. Hansen, ''An Evaluation of the Message Passing Interface,'' *ACM Sigplan Notices,* vol. 33, no. 3, 1998, pp. 65 72.
8. M. Wright, A. Freed, and A. Momeni, ''Open soundcontrol: State of the Art 2003,'' *Proc Int'l Conf. New Interfaces for Musical Expression,* 2003, pp. 153 160.
9. G. Humphreys, M. Houston, and R. Ng, ''Chromium: A Stream Processing Framework for Interactive Rendering on Clusters,'' *ACM Trans. Graphics,* vol. 21, no. 3, 2002, pp. 693 702.
10. O. Kreylos, *Environment Independent VR Develop ment,* tech. report, Univ. of California, Davis, 2006.
11. M. Howe, ''A Proposal for Dome Standards,'' *Proc. IPS Fulldome Standards Summit,* 2004.
12. R. Raskar et al. ''Quadric Transfer for Immersive Curved Screen Displays,'' *Proc. Eurographics,* Blackwell Publishing, vol. 23, no. 3, 2004, pp. 451 460.
13. T. Ni et al., ''A Survey of Large High Resolution Display Technologies, Techniques, and Applica tions,'' *Proc. IEEE Virtual Reality Conf.,* IEEE CS Press, 2006, pp. 223 236.
14. B. Schaeffer and C. Goudeseune, ''Syzygy: Native PC Cluster VR,'' *Proc. IEEE Virtual Reality Conf.,* IEEE CS Press, 2003, p. 15.
15. A. Bierbaum et al., ''VR Juggler: A Virtual Platform for Virtual Reality Application Development,'' *Proc. IEEE Virtual Reality Conf.,* IEEE CS Press, 2001, pp. 89 96.
16. J. Ballas, ''Delivery of Information through Sound,'' *Auditory Display: Sonification, Audification and Audi tory Interfaces,* Addison Wesley, 1994, pp. 79 94.
17. D. Conant et al., ''The Acoustics of a Large 3D Immersive Environment: The AlloSphere at UCSB,'' *Proc. ASA EAA Joint Conf. Acoustics,* ASA, 2008.
18. V. Pulkki, ''Virtual Sound Source Positioning Using Vector Base Amplitude Panning,'' *J. Audio Engi neering Soc.,* vol. 45, no. 6, 1997, pp. 456 466.
19. D. Malham and A. Myatt, ''3D Sound Spatializa tion Using Ambisonic Techniques,'' *Computer Music Journal (CMJ),* vol. 19, no. 4, 1995, pp. 58 70.
20. S. Spors, H. Teutsch, and R. Rabenstein, ''High Quality Acoustic Rendering with Wave Field Syn thesis,'' *Proc. Vision, Modeling, and Visualization Workshop,* Aka GmbH, 2002, pp. 101 108.
21. A. Berkhout, ''A Holographic Approach to Acous tic Control,'' *J. Audio Engineering Soc.,* vol. 36, no. 12, 1988, pp. 977 995.
22. D. McCoy, *Ventriloquist: A Performance Interface for Real Time Gesture Controlled Music Spatializa tion,* master's thesis, Univ. of California, Santa Barbara, 2005.
23. M.A. Gerzon, ''Periphony: With Height Sound Reproduction,'' *J. Audio Engineering Soc.,* vol. 21, no. 1, 1973, pp. 2 10.
24. F. Hollerweger, *Periphonic Sound Spatialization in Multiuser Virtual Environments,* master's thesis, Austrian Inst. Electronic Music and Acoustics, 2006.
25. G. Wakefield, ''Third Order Ambisonic Extensions for MaxMSP with Musical Applications,'' *Proc. Int'l Computer Music Conf.,* ICMA, 2006, pp. 123 126.
26. W. Wolcott, *Wave Field Synthesis with Real Time Control,* master's thesis, Univ. of California, Santa Barbara, 2007.
27. J. Castellanos, *Design of a Framework for Adaptive Spatial Audio Rendering,* master's thesis, Univ. of California, Santa Barbara, 2006.

28. X. Amatriain, ''A Domain Specific Metamodel for Multimedia Processing Systems,'' *IEEE Trans. Multimedia,* vol. 9, no. 6, 2007, pp. 1284 1298.
29. S.T. Pope and C. Ramakrishnan, ''The Create Signal Library (''Sizzle''): Design, Issues and Applications,'' *Proc. Int'l Computer Music Conf.,* ICMA, 2003.
30. X. Amatriain et al., ''Immersive Audio and Music in the AlloSphere,'' *Proc. Int'l Computer Music Conf.,* ICMA, 2007.

Xavier Amatriain works at the Telefonica research center in Barcelona, Spain, and is associate professor at the Universitat Pompeu Fabra in Barcelona. His research interests include multimedia systems, information retrieval and personalization, immersive and augmented reality, and audio and music processing. Amatriain has a PhD in computer science from the Universitat Pompeu Fabra. Contact him at xavier@amatriain.net.

JoAnn Kuchera Morin is a professor of media arts and technology and director of the AlloSphere Research Facility at the University of California, Santa Barbara. Her research interests include composition, human–computer interaction, and immersive data representation. Kuchera Morin has a PhD in music from the Eastman school of music. Contact her at jkm@create.ucsb.edu.

Tobias Hollerer is an associate professor of computer science at the University of California, Santa Barbara, where he codirects the ''Four Eyes Laboratory'' on Imaging, Interaction, and Innovative Interfaces. His research interests include augmented reality, wearable and ubiquitous computing, and computer graphics and visualization. Hollerer has a PhD in computer science from Columbia University. Contact him at holl@cs.ucsb.edu.

Stephen Travis Pope is a senior research specialist working in computer music at the University of California, Santa Barbara. His research interests include spatial sound, distributed processing, and music representation languages. Contact him at stephen@heaveneverywhere.com.