

The Allosphere: A Large-Scale Immersive Surround-View Instrument

Tobias Höllerer, holl@cs.ucsb.edu
JoAnn Kuchera-Morin, jkm@create.ucsb.edu
Xavier Amatriain, xavier@create.ucsb.edu
University of California Santa Barbara

Abstract

We present the design of the Allosphere and initial experiences from its ongoing implementation. The UCSB Allosphere is a novel large-scale instrument for immersive visualization and simulation, which in its full realization will be one of the world's largest immersive environments. The three-story high cubical space comprises an anechoic chamber with a spherical display screen, ten meters in diameter, surrounding from one to thirty users standing on a bridge structure. The Allosphere is differentiated from conventional virtual reality environments by its size and focus on collaborative experiences, its seamless surround-view capabilities and its focus on multiple sensory modalities and interaction. The Allosphere is being equipped with high-resolution active stereo projectors, a complete 3D sound system with hundreds of speakers, and interaction technology. In this paper we will give an overview of the purpose of the instrument as well as the systems that are being put into place in order to equip it. We also review the first results and experiences in developing and using the Allosphere in several prototype projects.

1 Introduction

For almost a decade, a team of digital media researchers at UCSB has been fostering a cross-disciplinary collaboration that connects science, engineering, and the arts through the use of new media. The culmination of these efforts is the Allosphere, a scientific instrument that enables manipulation, exploration and analysis of large-scale data sets in a rich, multi-user immersive environment. The Allosphere is situated at one corner of the California Nanosystems Institute building at the University of California Santa Barbara (see virtual model in Figure 1), surrounded by a number of associated labs for visual/audio computing, robotics and distributed systems, interactive visualization, and world modeling. These labs are used to prototype the technologies to be deployed in the Allosphere as well as for content production and media post-production. A main target use of the instrument is scientific visualization/auralization and data exploration, but we also anticipate it being used as a research environment for behavioral/cognitive scientists and artists.

The Allosphere is a novel immersive visualization, simulation, and performance space, forming one of the world's largest high-resolution immersive environments. Up to thirty observers can experience an Allosphere session from a bridge structure.

Figure 2 shows the Allosphere from above. The space consists of




Figure 1: A virtual real-scale model of the Allosphere, situated in the UCSB CNSI Building




Figure 2: Horizontal section of the Allosphere

a 3-story-high trapezoidal chamber that is treated with extensive sound absorption material (4-foot wedges on almost all inner surfaces), making it one of the largest quasi-anechoic rooms in the world. Standing inside this chamber are two 5-meter-radius hemispheres constructed of perforated aluminum that are designed to be optically opaque (for front projection, optimized for low optical scatter) and as acoustically transparent as possible. One of the reasons that the instrument was designed as an ellipsoid instead of a perfect sphere is because of the problem of strong focused sound reflection at the center of a sphere. This spot is now broken up into two less problematic locations above the railing at the lengthwise mid-point of the bridge. The bridge area is unaffected.

There are four primary distinguishing characteristics of the Allosphere, when we compare it to existing immersive environments such as the CAVE [Cruz-Neira et al. 1992], a hemispherical immersive theater [Gaitatzes et al. 2006], or VR spaces facilitated by head-worn displays [Cakmakci and Rolland 2006]. First, the Allosphere makes it possible for a number of users to be active in the environment at once (up to 30 people on the bridge) and collaborate on an analysis task. Second, its seamless surround-view design pro-




Figure 3: *The Allosphere in simulation, as seen from the entrance, looking down the bridge. For visibility purposes, screen segments above the bridge have been omitted*

vides an extreme sense of immersion with little encumbrance and limited distortions away from the center of projection or tracked user. Third, the space was designed as an intrinsically multi-modal interaction environment. The spherical screen is placed in a carefully designed anechoic chamber and is perforated to enable spatialized audio from a planned 120 (currently 30) speaker system behind it. Medium to long-term we have mapped out research efforts on unencumbering tracking and interaction technologies, but these are not addressed in this paper. Fourth, the Allosphere will provide high sensory resolution. The extent of fidelity across the multiple modalities, including graphics and audio in balanced combination, is unparalleled for an interactive immersive environment. At this point in time, only a subset of four out of the planned fourteen projector configuration is operational, providing a sub-optimal but functional prototype, which is currently being evaluated for the best ways to scale the installation up to full immersiveness.

With its unique spherical shape, its high sensory resolution, and its immersive multi-modal capabilities, 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 24 projector C6 CAVE at the Virtual Reality Applications Center at Iowa State University, the Fakespace FLEXT installation at Los Alamos National Laboratory [LANL 2006], which is a 33 projector 5-wall CAVE with 43M pixels in a 15x10x12ft room, the Samuel Oschin Planetarium at the Griffith Observatory in Los Angeles with its E&S Digistar 3 full dome laser projection system with claimed 8kx8k = 64M pixel resolution in its 75 foot diameter dome, the Denver Museum of Nature & Science Gates Planetarium dome [DMN 2006] with approximately 10M pixels from 11 projectors, and the Louisiana Immersive Technologies Enterprise center with its 8.8M pixel 6-sided CAVE and smaller curved display immersive theater and conference room.

Compared with CAVE environments [Cruz-Neira et al. 1992], [Ihren and Frisch 1999], the Allosphere has a different focus: It enables seamless immersive projection, even in non-stereo mode. Room geometry does not distort the projected content. Stereo is possible for a large set of users, since the audio and stereovision "sweet spot" area is much larger, the screen being consistently 5 meters and farther away. Single person head-tracked stereo is still possible, but because of the bridge structure and the exit doorways on either end, CAVE-like "walk-around-the-object" stereo imagery is restricted to railing height or higher. Our prediction is that the most common usage scenario for this display is going to be non-



Figure 4: *Visitors on the Allosphere bridge*

tracked stereo projection for a medium-sized group of people (5-10), who collaborate with each other near the center of the bridge.

2 Towards Immersion for the Masses

The Allosphere is uniquely positioned in between CAVE environments, which give fully immersive experiences to a small number of users, and full-dome planetarium style theaters, which have extremely high outreach potential but a limited sense of immersion. In designing the Allosphere, we targeted the best of both worlds. With it, we hope to develop technologies that immediately benefit recent endeavors at select full-dome planetaria to diversify their science programs and enable more interactive presentations [Neafus and Yu 2007].

The main technical and functional innovations of our instrument can be summarized in the following:

- * It 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 [Shibano et al. 2003] but on a different magnitude scale. It is known that spherical immersive systems enhance subjective feelings of immersion, naturalness, depth and "reality" (see [Kalawsky 1993], for instance).
- * It is fully multimedia as it combines state-of-the-art techniques both on virtual audio and visual data spatialization. There is extensive evidence of how combined audio-visual information can help information understanding [McGurk and McDonald 1976]. Nevertheless most existing immersive environments focus on presenting visual data.
- * It will be a completely interactive multimodal environment [Malkawi and Srinivasan 2004], including camera tracking systems, audio recognition and sensor networks.
- * It is designed to be a pristine scientific instrument – e.g., the containing cube is a large fully anechoic chamber and the details such as room modes or screen reflectivity have been included in the building design discussions.
- * Its size allows for up to 30 people to interact and collaborate on a common research task.
- * Although the instrument will be flexible and general-purpose, its focus is on scientific research in fields such as data mining [Wegman and Symanzik 2002], geographic visualization or bio-imaging.

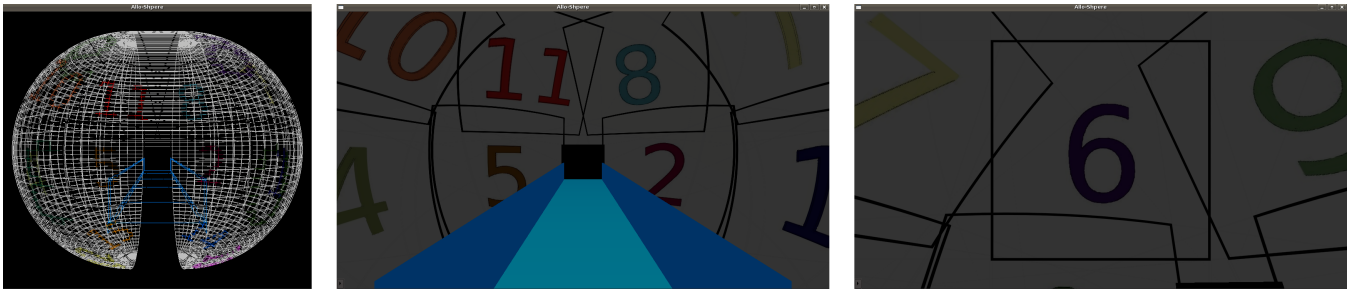


Figure 5: Interactive simulator for placing projectors and experimenting with different coverage models; a) wireframe overview, b) view along bridge, c) view from projector 6

Figures 3 and 4 show a view into the instrument along the bridge, which is 2.13 meters (seven feet) wide. All four of our current projectors are placed underneath the bridge. Audio speakers are mounted behind the display surface. In combination, they provide high-resolution stereo video and spatialized sound.

Although the space is not fully equipped at this point, we have been experimenting and prototyping with different equipment, system configurations, and applications that pose different requirements. This being a research instrument – rather than a static performance space – presenting unique requirements and research questions, the final system will be the result of an evolving prototype. As a matter of fact, we envision the instrument as an open framework that is in constant evolution with major releases signaling major increments in functionality.

Off-the-shelf computing and interface solutions proved to be inadequate in our system design. Allosphere applications not only require a server farm dedicated to video and audio rendering and processing, but also a low-latency interconnection fabric so that data can be processed on multiple computers (in a variety of topologies) in real time, an integration middleware, and an application server that allows clients to interface with the system in a flexible and meaningful way. The Allosphere Network will have to host not only standard/low-bandwidth message passing but also multichannel multimedia streaming. The suitability of Gigabit Ethernet or Myrinet versus other proprietary technologies is still under discussion. In our first prototypes, Gigabit has proven sufficient but our projections show that this will become a bottleneck for the complete system, specially when using a distributed rendering solution is used to stream highly dynamic visual applications.

The remainder of this document will describe basic components and subsystems as well as the way they interact, and the guidelines upon which future design and implementation decisions will be taken. In particular, in Sections 3 and 4 we will give a more detailed overview of the visual and audio subsystems, concentrating on the generative aspect, especially on the graphics side.

3 The Visual Sub-System

This section addresses the visual subsystem of the Allosphere, addressing projection, image generation and rendering components.

3.1 Projection

The main requirements for the Allosphere visual sub-system are fixed both by the building construction and constraints, and by the final quality that is targeted. The sphere screen area is 320.0 m² and its reflective gain, FOV averaged is 0.12. For a good performance we need a minimum of 3 arc minutes of angular resolution,

although ideally 1 arc minute should be targeted. In terms of light level we need 50 trolands although that number can be limited to 30 when projecting active stereo.

With these requirements, we have mapped out a projection system consisting of 14 3-chip DLP active stereo projectors that are capable of a maximum 3000 lumens output and SXGA+ resolution (1400x1050). The projectors are being installed with an effective projector overlap/blending loss coefficient of 1.7. In the following paragraphs we will outline the features that such a system will yield.

3.1.1 Image Brightness

One of the important design goals for the Allosphere is to make it user-friendly and usable for extended periods of time. Unacceptably low levels of brightness will cause eye fatigue or will severely restrict the type of content that Allosphere can display. Without yet considering the stereo requirement, the projected system yields 42,000 lumens and a screen luminance (full white) of 9.26 cd/m². For comparison, the luminance of a typical good quality multimedia Dome is recommended between 0.686 and 5.145 cd/m² [Howe 2004]

According to our simulations and on-site tests, 42,000 lumens of input flux produce very close to optimal results. Augmenting the light flux above this level has several undesired effects, namely cross-reflection and ghosting due to back-reflection.

3.1.2 Stereoscopic Display

The performance of the Allosphere in stereo mode depends on the choice of the stereo display technology. Nevertheless, passive polarization-based methods are ill-suited for the Allosphere due to the curved nature of the screen and the non-polarization-preserving material.

Light losses from stereo projection are substantial. The design requirement for stereo projection brightness level is 30 trolands at 50% RGB average. Stereo projection mode will fall below the theoretical eyestrain threshold with our projected 42K lumens total projector flux. Nevertheless, our field studies have indicated that this level of brightness in stereo mode is still perceived as high quality and allows continuous working times of beyond 60 minutes after some adaptation period. As a matter of fact, the main cause of eyestrain in stereo mode is active shuttering itself and this is not correlated with luminance.

3.1.3 Contrast Ratio

Contrast loss due to diffused scattering has represented a serious problem for the projection system 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. Nevertheless, this would have the unwanted effect of requiring more input light flux and increasing the back reflections (due to the screen perforation), heat, noise, etc. The screen gain has been decided 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 and the Gates Planetarium in the Denver Museum of Natural History).

The screen paint has FOV-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% total light flux input. We are in the process of verifying this data in a set of systematic measurements with projector directions from various angles. These measurements will feed into a simulator we are developing for optimizing projector placement and calibration (cf. Figure 5).

3.1.4 Screen Resolution

The resulting visual resolution for the Allosphere is a function of the total number of pixels available and the projector overlap factor, which is calculated to be 1.7.

The spatial acuity of 20/20 eyesight is 1MAR or 30 line pairs per degree¹. Nevertheless, resolutions down to 3 arc minutes are reported as high quality unless a better reference point is available. As explained in [Howe 2004] by taking the center of the (hemi)sphere as a common view point, we can infer the number of pixels for a given resolution independently of the screen size or diameter. A limit 3 arc minute resolution requires 20 million pixel spread over a full sphere. The projected configuration with 14 projectors has 19.2 Mpixels which comes close to the desired resolution.

3.1.5 Image Warping and Blending

The Allosphere projection system requires image warping and blending to create the illusion of a seamless image from multiple projectors. The two places where image warping and blending can be done are the graphic cards on the image generation system and the video projectors.

Most modern simulation-oriented video projectors support some form of warping and blending. This approach is very convenient, often resulting in the best image quality, and it is typically combined with color correction and whatever other processing the projector has to do for its internal purposes. A negative side-effect of this technique is the fact that the projector has to buffer an entire frame before being able to process it. Another downside is that projector-based warping and blending relies on vendor-specific proprietary software that is hard to access and extend.

In the case of computer-side warping and blending, graphics rendering hardware can perform frame warping and edge-blend mixing after the frame buffer is filled. This consumes additional resources that could otherwise be used to render polygons, and is best done with specialized hardware designed for that sole purpose. Such hardware, however, is costly and proprietary. Calibration procedures also become more complicated. The benefit of computer-side warping is the reduced latency, because the video projector does not need to buffer an entire frame before displaying it. Nevertheless, although some LCD projectors are able to display images with latencies as low as 3ms, DLP projectors require a separate frame buffer in any case, due to the nature of DLP technology.

¹This is the average spatial acuity in “regular” conditions as spatial resolution is a function of both contrast ratio and pupil size



Figure 6: Testing the Allosphere projection with a single 2K active stereo projector and brain imaging data

In the Allosphere, we have decided to start off with a projector-side warping/blending solution that is readily available and solves many but not all of the Allosphere requirements. In parallel, we are working on extending and adapting existing solutions such as [Cani and Slater 2004] to the case of a full spherical surface.

3.1.6 Latency and Frame Rate

All occurring latencies, from the start of rendering until the image appears on the screen, must be considered in the Allosphere system design. The literature indicates that unpleasant side-effects appear above 120ms total system latency for interactive VR applications. Below this value, the lower the latency, the more accurate and stress-free the interaction becomes. In general, a total system delay of 50ms is considered to be state-of-the-art.

3.2 Image Generation and Rendering

The following paragraphs describe the image generation subsystem, which includes the computers and the software that generate video signals, which then are transported by the interconnecting subsystem to the video projectors. In order to meet our display requirements, we needed an image generation system capable of producing 20 million pixels through 14 channels. The system should support resolutions of at least XVGA+, should offer active stereo support, as well as framelock capabilities for synchronizing all channels.

From these requirements, we have designed a rendering cluster made up of 7 HP-9400 workstations, each of which is equipped with an NVIDIA FX-5600 and a G-sync card for frame-locking, which, after initial driver problems, we now run to produce stereo using CUDA.

As far as the software component is concerned, in order to generate large multi-tile immersive displays different techniques and tools are possible (see [Ni et al. 2006]). Nevertheless, the Allosphere poses some unique problems that are still matter of research, namely:

- * Tiles are irregularly shaped and curved
- * The projection screen is a full continuous *quasi-sphere*, which adds a wrap-around problem.
- * Because the Allosphere is strictly speaking not a perfect sphere, but rather two hemispheres separated by the bridge,

warping solutions such as the one presented in [Cani and Slater 2004] are not directly applicable.

- * Projection should allow for active stereo to work reasonably well in most of the field of view for an arbitrary viewer.
- * As flexibility is one of the key design parameters of the instrument, time to adapt a legacy application to the Allosphere should be minimized.
- * A middleware layer should be offered, allowing for the display of any OpenGL application, even when source code access is not available.

Some of these requirements are still to be accomplished and are in fact being addressed by our current research and development. Nevertheless, our first experiments and results point out that a two-way approach is recommended:

For those applications in which *no source is available* or in which *viewpoint information is not “relevant”* for rendering a convincing 3D immersive scene (most scientific visualization packages that render abstract data), we are using a distributed rendering solution based on Chromium [Humphreys et al. 2002]. In this case, a single master runs the application and performs early rendering stages, offloading the rendering of the specific projector viewpoints to appropriate slaves. This is the solution that is being used for two of our “early adopter” applications: molecular dynamics using VMD [Humphrey et al. 1996] and the home-brewn NanoCAD (see section 5).

In those applications for which *source code is available* and which require complete *viewpoint-dependent rendering*, we use an approach based on fully replicating the whole application at each node and synchronizing the execution. In this case, the master manages the application state and processes user input from the interface. Rendering is performed completely on the slaves, which receive information of the application state and have knowledge of their particular viewpoint and rendering tile. This is a similar approach to that used by VR libraries such as Syzygy [Schaeffer and Goudeuseune 2003] or VRJuggler [Bierbaum et al. 2001]. This is the solution used in the AlloBrain project (see section 5).

Other intermediate solutions, e.g. distributing 3D objects using either Open Scene Graph or OpenSG are also possible and are currently being explored.

4 The Audio Sub-System

Our goal for the Allosphere is to build an immersive multimedia interface that provides “sense-limited” resolution in both the audio and visual domains. This means that the spatial resolution for the audio output must allow us to place virtual sound sources at arbitrary points in space with convincing synthesis of the spatial audio cues used in psychoacoustical localization. Complementary to this, the system must allow us to simulate the acoustics of measured or simulated spaces with a high degree of accuracy.

4.1 Acoustical Requirements

In order to provide for “ear-limited” dynamics, frequency, and spatial extent and resolution, we require the system to be able to reproduce in excess of 100 dB sound pressure level near the center of the sphere. We need to have acceptable low- and high-frequency extension (-3 dB points below 80 Hz and above 15 kHz), and to provide spatial resolution on the order of 3 degrees in the horizontal plane (i.e., 120 channels), and 10 degrees in elevation. To provide high-fidelity playback, we require audiophile-grade audio distribution formats and amplification, so that an effective signal-to-noise

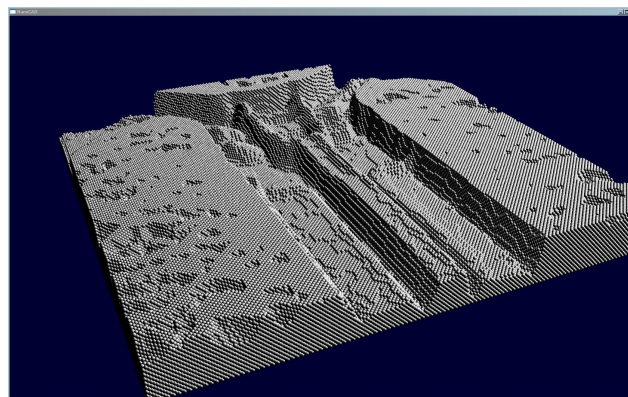


Figure 7: Rendering of a 1M atoms silicon nanostructure on real-time on a single CPU/GPU. One of the first scientific Allosphere demonstrations.

ratio exceeds 80 dB, with a useful dynamic range of more than 90 dB.

To be useful for data sonification [Ballas 1994] and as a music performance space, the decay time (the “T60 time”) of the Allosphere must be less than 0.75 seconds from 100 Hz to 10 kHz. This is primarily an architectural feature related to the properties of the aluminum screen and the sound absorbing treatment in the anechoic chamber.

4.2 Spatial Sound Processing

There are three techniques for spatial sound reproduction used in current state-of-the-art systems: (1) vector-based amplitude panning, (2) ambisonic representations and processing, and (3) wave field synthesis. Out of these, the latter is the most flexible and powerful (see [Rabenstein et al. 2005]). In any case, the Allosphere speaker count and configuration supports the use of any of these for sound spatialization. This implies high speaker density (on the order of one source per square yard of surface, or about 380 channels), and a semi-regular and relatively symmetrical speaker layout.

We have developed a generic software framework in which different techniques and speaker layouts can be combined and interchange with virtually no user effort.

4.3 Speaker System

It has been a major project to derive the optimal speaker placements and speaker density function for use with mixed-technology many-channel spatialization software. Our driver placement design comprises between 425 speakers arranged in several rings around the upper and lower hemispheres, with accommodations at the “seams” between the desired equal and symmetrical spacing and the requirements of the support structure. The loudspeakers will be mounted behind the screen.

Our plans foresee densely packed circular rings of speaker drivers running just above and below the equator (on the order of 100-150 channels side-by-side), and 2-3 smaller and lower-density rings concentrically above and below the equator. The main loudspeakers have limited low-frequency extension, in the range of (down to) 200-300 Hz. To project frequencies below this, four large subwoofer(s) are mounted on the underside of the bridge.

The (passive) speaker elements will be wired to a set of 8-16 networked digital-to-analog converter (DAC) amplifier boxes, each of




Figure 8: Screen capture of the "AlloBrain" interactive tour of the human brain from fMRI data, and a group of middle school girls reaching out at the stereoscopic projection

which supports in the range of 32-128 channels and has a Firewire interface.

5 Testbed Applications

The Allosphere's function is the analysis, synthesis, simulation, and processing of complex multidimensional data in an interactive and immersive environment. Content and demand will drive the technological development just as it has driven its design and conception. Specific application areas are essential in the development of the instrument, as they define the functional framework in which the instrument will be used.

In the first iteration of the prototype, we have set up an environment consisting of the following elements:

- * 4 active stereo projectors (Christie Digital Mirage S+2K), 3000 ANSI lumens, DLP
- * 2 rendering workstations (HP 9400), AMD Opteron 64@2.8Ghz, NVidia Quadro FX-5500 and FX-5600
- * 1 application manager + Audio Renderer (Mac Pro), Intel Xeon Quad Core @3Ghz
- * 2 10-channel firewire audio cards.
- * 16 full-range speakers + 1 subwoofer
- * Several custom-developed wireless user interface devices.

The research projects described below make use of this prototype system to test the functionality and prove the validity of the instrument design.

In the first project, we are developing an immersive and interactive software simulation of nano-scaled devices and structures, with atom-level visualization of those structures implemented on the projection dome of the Allosphere (see figure 7). When completed, this will allow the user to stand in the middle of a simulation of a nano-scaled device and interact with the atoms and physical variables of that device.

Our scientific partners in material science and chemistry are implementing computational science and engineering algorithms in the areas of molecular dynamics and density functional theory using GPU's, transforming a single PC workstation into a 4 Teraflop supercomputer. This allows us to run specific nanoscale simulations

that are 2-3 orders of magnitude faster than current implementations. We will also be able to use this extra computational power to solve for the physical properties of much larger structures and devices than were previously possible, allowing nano-engineers to design and simulate devices composed of millions of atoms.

In the second research project, called *AlloBrain* (see Figures 6 and 4.3), we operate with macroscopic, organic data sets, reconstructing an interactive 3D model of a human brain. This test bed project displays an interactive 3D model of one of our team member's brain, reconstructed from fMRI data. The current model contains two layers of tissue blood flow, and we have created an interactive environment where twelve animated "agents" navigate the space and gather information to deliver back to the researchers. The simulation contains several generative audio-visual systems. These systems are stereo-optically displayed and controlled by two wireless (Bluetooth) input devices that feature custom electronics, integrating several MEMs sensor technologies.

One controller contains twelve buttons that control the twelve agents. The controller also allows you to move the ambient sounds spatially around the sphere. The second controller allows you to navigate the space. The shape of the larger controller is based on the hyperdodecahedron, a 4-dimensional geometrical polytope, the final device shape being represented by its shadow projected into 3 dimensions. It was developed using procedural modeling techniques. The mathematical model was algorithmically sculpted to provide a more organic look and feel while preserving the 600 internal vertices, and it was constructed with a 3-D printer capable of building solid objects.

This virtual interactive prototype illustrates some of the key features in the Allosphere project, such as multimedia/multimodal computing, interactive immersive spaces and scientific data understanding through art.

Finally, we are also starting a project focused on molecular dynamics. We are extending the VMD [Humphrey et al. 1996] package through the use of Chromium in order to have seamless visualization of complex protein molecules and their interactions in the context of the Allosphere.

The development of these test bed applications, among others, is geared towards the development of an open generic software infrastructure capable of handling multi-disciplinary applications with common goals, and it will facilitate the development of an open

ended general computational system for data generation, manipulation, analysis and representation.

6 Conclusions

Once fully equipped and operational, the Allosphere will be one of the largest immersive instruments in existence. But aside from its size, it also offers a number of features that make it unique in several ways: fully immersive spherical projection, multimodal interaction including 3D audio and novel devices, and multi-user capabilities. Such features are only possible because the instrument has been conceptualized and designed with a combination of all these requirements in mind.

We envision the Allosphere as a vital instrument for the future advancement of fields such as nanotechnology or bioimaging. It will stress the importance of multimedia for the advancement of science, engineering, and the arts. In this paper, we have shown first results in the form of projects featuring very diverse requirements. These initial results feed back into the prototyping process but already indicate the validity of our approach.

Although the Allosphere is still at the beginning of its potential, we believe that its uniqueness make these first results already valuable. Besides, in a research instrument such as this, evolution is built into the life cycle in such a way that it will continue to provide new insights and solutions, hopefully for many years to come.

7 Acknowledgments

The Allosphere project is conducted and supported as a UCSB Media Arts and Technology (MAT) Initiative, and additional support and contributions came from many individuals from various departments. We'd like to first and foremost thank MAT professors Marcos Novak and Stephen Pope and the rest of the MAT faculty, and especially the many MAT graduate students who directly contributed to the project with their thesis, project, and class work. Special thanks also to Brent Oster and Liubov Kovaleva among several other graduate students from Computer Science. Other departments such as Chemistry, Psychology, Geography, and Biology have also contributed in different ways.

References

- BALLAS, J. 1994. Delivery of information through sound. In *Auditory Display: Sonification, Audification and Auditory Interfaces*, G. Kramer, Ed., vol. XVIII. Addison Wesley, Reading, MA, 79–94.
- BIERBAUM, A., JUST, C., HARTLING, P., AND MEINERT, K. 2001. Vr juggler: A virtual platform for virtual reality application development. In *Proceedings of the IEEE Virtual Reality 2001 Conference (VR '01)*.
- CAKMAKCI, O., AND ROLLAND, J. 2006. Head-worn displays: a review. *Journal of Display Technology* 2, 3, 199–216.
- CANI, M., AND SLATER, M. 2004. Quadric transfer for immersive curved screen displays. In *Proceedings of the 2004 Eurographics*, vol. 23.
- CRUZ-NEIRA, C., SANDIN, D. J., DEFANTI, T., KENYON, R., AND HART, J. 1992. The cave: Audio visual experience automatic virtual environment. *Communications of the ACM*, 35, 64–72.
2006. The Denver Museum of Nature and Science Gates Planetarium. URL: <http://www.dmns.org/>.
- GAITATZES, A., PAPAIOANNOU, G., CHRISTOPOULOS, D., AND ZYBA, G. 2006. Media productions for a dome display system. In *VRST '06: Proceedings of the ACM symposium on Virtual reality software and technology*, ACM Press, New York, NY, USA, 261–264.
- HOWE, M. 2004. A proposal for dome standards. spatial resolution, luminance and system contrast. In *Proceedings of the IPS 2004 FullDome Standards Summit*.
- HUMPHREY, W., DALKE, A., AND SCHULTEN, K. 1996. Vmd - visual molecular dynamics. *Journal of Molecular Graphics*, 14, 33–38.
- HUMPHREYS, G., HOUSTON, M., NG, R., FRANK, R., AHERN, S., KIRCHNER, P., AND KLOSOWSKI, J. 2002. Chromium: A stream processing framework for interactive rendering on clusters. In *Proc. ACM SIGGRAPH*, 693–702.
- IHREN, J., AND FRISCH, K. 1999. The fully immersive cave. In *Proc. 3rd International Immersive Projection Technology Workshop*, 59–63.
- KALAWSKY, R. 1993. *The science of virtual reality and virtual environments*. Addison-Wesley Pub.
- LANL, 2006. Fakespace Delivers Highest Resolution Visualization Room (Advanced Simulation and Computing (ASC) Program, Los Alamos National Laboratory). URL: http://www.tenlinks.com/NEWS/PR/fakespace/110905_room.htm.
- MALKAWI, A., AND SRINIVASAN, R. 2004. Multimodal human-computer interaction for immersive visualization: Integrating speech-gesture recognitions and augmented reality for indoor environments. In *Proceedings of the Seventh IASTED Conference on Computer Graphics and Imaging*, ACTA Press, Kauai, Hawaii, 171–175.
- MCGURK, H., AND MCDONALD, T. 1976. Hearing lips and seeing voices. *Nature*, 264, 746–748.
- NEAFUS, D., AND YU, K., 2007. The Digital Gates Planetarium at the Four-Year Mark: Innovations and Adventures of a FullDome Theater. (in preparation).
- NI, T., SCHMIDT, G., STAADT, O., LIVINGSTON, M., BALL, R., AND MAY, R. 2006. A survey of large high-resolution display technologies, techniques, and applications. In *Proceedings of the 2006 IEEE Virtual Reality Conference*.
- RABENSTEIN, R., SPORS, S., AND STEFFEN, P. 2005. *Selected methods of Acoustic Echo and Noise Control*. Springer Verlag, ch. Wave Field Synthesis Techniques for Spatial Sound Reproduction.
- SCHAEFFER, B., AND GOUDESEUNE, C. 2003. Syzygy: Native pc cluster vr. In *Proceedings of the IEEE Virtual Reality 2003*, 15.
- SHIBANO, N., HAREESH, P., HOSHINO, H., KAWAMURA, R., YAMAMOTO, A., KASHIWAGI, M., AND SAWADA, K. 2003. Cyberdome: Pc clustered hemi spherical immersive projection display. In *Proc. of the 2003 International Conference on Artificial Reality and Telexistence*, 1–7.
- WEGMAN, E., AND SYMANZIK, J. 2002. Immersive projection technology for visual data mining. *Journal of Computational and Graphical Statistics* (March).