

Sensate abstraction: hybrid strategies for multi-dimensional data in expressive virtual reality contexts

Ruth West*^a, Joachim Gossmann^b, Todd Margolis^b, Jurgen P. Schulze^c,
JP Lewis^d, Ben Hackbarth^b, Iman Mostafavi^c

^aUCLA Center for Embedded Networked Sensing (CENS), BOX 951596 3551G Boelter Hall, Los Angeles CA 90095-1596; ^bUCSD Center for Research in Computing and the Arts, 9500 Gilman Drive, La Jolla CA 92093 – 0037; ^cUCSD/CALIT2 9500 Gilman Drive, La Jolla, CA 92093-0436;

^dWeta Digital Ltd. P.O Box 15208, 9-11 Manuka Street, Miramar, Wellington New Zealand 6022;

^eUCSD Computer Science and Engineering, 9500 Gilman Drive, La Jolla CA 92093-0404

ABSTRACT

ATLAS in silico is an interactive installation/virtual environment that provides an aesthetic encounter with metagenomics data (and contextual metadata) from the Global Ocean Survey (GOS). The installation creates a visceral experience of the abstraction of nature in to vast data collections – a practice that connects expeditionary science of the 19th Century with 21st Century expeditions like the GOS. Participants encounter a dream-like, highly abstract, and data-driven virtual world that combines the aesthetics of fine-lined copper engraving and grid-like layouts of 19th Century scientific representation with 21st Century digital aesthetics including wireframes and particle systems. It is resident at the Calit2 Immersive visualization Laboratory on the campus of UC San Diego, where it continues in active development. The installation utilizes a combination of infrared motion tracking, custom computer vision, multi-channel (10.1) spatialized interactive audio, 3D graphics, data sonification, audio design, networking, and the Varrier™ 60 tile, 100-million pixel barrier strip auto-stereoscopic display. Here we describe the physical and audio display systems for the installation and a hybrid strategy for multi-channel spatialized interactive audio rendering in immersive virtual reality that combines amplitude, delay and physical modeling-based, real-time spatialization approaches for enhanced expressivity in the virtual sound environment that was developed in the context of this artwork. The desire to represent a combination of qualitative and quantitative multidimensional, multi-scale data informs the artistic process and overall system design. We discuss the resulting aesthetic experience in relation to the overall system.

Keywords: immersive multiscale, multiresolution visualization, art-science collaboration, spatialized multi-channel interactive audio, audio rendering, audio spatialization

1. INTRODUCTION

ATLAS in silico is a physically interactive virtual environment/installation and art-science collaboration encompassing the disciplines of virtual reality, auditory display and sonification, visualization, electronic music and composition, new media arts, metagenomics, computer graphics, and high-performance computing. (Web: <http://www.atlasinsilico.net>) The installation was exhibited at SIGGRAPH 2007, the Ingenuity Festival Cleveland 2008, the Los Angeles Municipal Art Gallery 2008 - 2009 and is resident at the UCSD Center for Research in Computing and the Arts (CRCA) and Calit2 (California Institute for Telecommunications and Information Technologies) where it continues in active development. It provides an aesthetically-impelled public-facing encounter with metagenomics data from the Global Ocean Survey (GOS) - a groundbreaking snapshot of microbial biodiversity in the world's oceans [1]. As an indication of the massive scale of the GOS, the initial set of data released in 2007 nearly doubled the number of proteins in publicly accessible genetic databases. The installation explores the potential interplay between artistic and data-driven strategies, based on visual and auditory pattern, in working with massive multidimensional multi-scale, multi-resolution data. Within the virtual environment users explore GOS data in combination with contextual metadata at various levels of scale and resolution through interaction with multiple data-driven abstract visual and auditory patterns at different levels of detail. The notion of “context” frames the entire experience and takes various forms ranging from structuring the virtual environment according to metadata describing the entire GOS data collection, to the use of contextual metadata pertaining to geographical regions nearest sampling sites for individual records from the GOS to drive visual and

auditory pattern generation, to playing a role in both data sonification and in audio design that is responsive to user interaction. Participants experience an environment constructed as an abstract visual and auditory pattern that is at once dynamic and coherently structured, yet which reveals its characteristics as the participant disturbs the pattern through their exploration.

Additionally, to highlight the historical and ongoing interplay between scientific discovery and culture, the installation contextualizes GOS sequence data within global environmental and social data from geographic regions nearest the ocean locations at which the microorganisms were sampled, and combines these to construct the visual and auditory elements of the virtual environment. This historic linkage is reiterated in the highly abstract, and data-driven virtual world by combining the aesthetics of fine-lined copper engraving, lithography and grid-like layouts of 19th Century scientific representation with contemporary digital aesthetics including 3D wireframes, particle systems, and spatialized audio. In parallel to challenges that Darwin's work on natural selection posed to 19th Century representations of nature associated with concepts of species fixity, [2] the new view of nature provided by vast and abstract metagenomics data, such as the Global Ocean Survey, poses a fundamental challenge for our ability in the 21st Century to represent and intuitively comprehend nature. Within ATLAS *in silico*, this challenge becomes a visceral, sensate experience of the abstraction of nature in to vast databases — a practice that reaches back in to the history of expeditionary science of the 19th Century and which culminates in 21st Century expeditions like the GOS.

In the sections below we describe implementation of the installation on the Varrier™ autostereographic virtual reality system and on a C-wall passive stereo rear-projection display system. We present hybrid audio rendering strategies developed to enable the use of audio spatialization as an expressive dimension within this virtual environment to convey contextual information about the data being visualized and sonified as well as to enhance the sensation of abstraction and immersion. The overall design is motivated by a desire to create a visceral experience of vast and abstract data spaces reaching towards a “sensate abstraction.”

2. SYSTEMS AND DISPLAYS

2.1 Global Ocean Survey Data

The Global Ocean Survey Expedition (2003 - 2006) conducted by the J. Craig Venter Institute, studies the genetics of communities of marine microorganisms throughout the worlds oceans. These microorganisms sequester carbon from the atmosphere with potentially significant impacts on global climate, yet the mechanisms by which they do so are poorly understood. Whole genome environmental shotgun sequencing resulted in a vast metagenomics data set that “produced the largest catalogue of genes to date from thousands of new species, with no apparent slowing of the rate of discovery (i.e., attaining saturation of data). These data have potentially far-reaching implications for biological energy production, bioremediation, and creating solutions for reduction/management of greenhouse gas levels in our biosphere [3].” The data contains DNA sequences and their associated predicted amino acid (protein) sequences. These predicted sequences, called “ORFs” (Open Reading Frames), candidates for putative proteins, are subsequently validated by a variety of bioinformatics analyses. It also includes a series of metadata descriptors, such as temperature, salinity, depth of the ocean, and depth of the sample, latitude and longitude of the sample location that describe the entire GOS data collection. Values for each metadata category were recorded at the sampling locations and stored along with the raw data. We use GOS metadata categories as a contextual container and to structure the virtual environment in which the data is explored. Given a collection of records from the GOS, each record is represented visually and sonically in the virtual environment as a graphical/auditory element and simultaneously as a spatio-temporal pattern that represents the sum of the values for all of its metadata attributes. Our visualization and sonification schemas are an active area of research and refinement and as such will be presented in future publications. The goal of our ongoing research is to enable interactive exploration of the entire GOS data set of over 17.4 million records within an immersive virtual reality context. Currently subsets of these records can be viewed and interactively explored within the ATLAS *in silico* installation.

2.2 Display Systems

At Calit2/CRCA the installation utilizes a combination of infrared motion tracking, custom computer vision, IR-reflective head and hand tracking, multi-channel (10.1) spatialized interactive audio, 3D graphics, data sonification, audio design, networking, a 16 node graphics/compute cluster and the Varrier™ 60 tile, 100-million pixel barrier strip auto-stereographic display. For touring, the installation was modified and uses a passive stereo display system. The touring installation is rear-projected in stereo on to a polarization-preserving screen, the multi-channel spatialized interactive audio makes use of 8 channels, and user input/interaction is through an electromagnetic tracking system using both head and hand sensors.

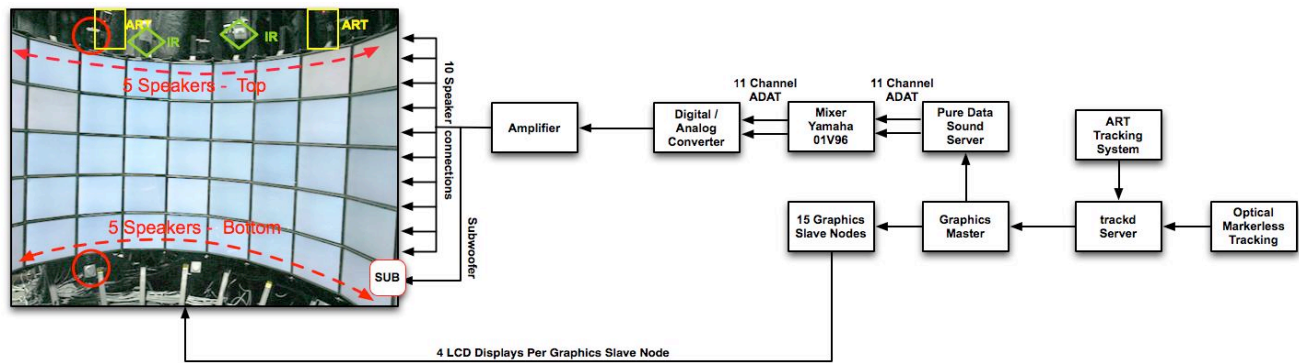


Fig. 1: Varrier system implementation

The Varrier™ autostereoscopic virtual reality display system [4] is comprised of 60 LCD tiles in a semi-circular cylindrical array that surrounds the viewer/participant. The 100-million pixel tiled display utilizes a combination of physical parallax barrier strip and virtual barrier technology to produce a wide field of view, head-tracked virtual reality experience with real-time interactivity yet without the need for users to wear specialized 3D glasses or other hardware. The system produces interleaved left and right eye perspectives by simulating the action of the physical barrier screen in virtual world coordinates in combination with distributing and correlating perspectives through a physical linescreen applied to the surface of each display panel. The correct registration of the physical linescreen and eye positions with the computed virtual linescreen and projection points results in an autostereoscopic image. Computation is performed on a 16-node (1 master, 15 slave) dual-opteron Linux cluster networked via a 1 Gigabit back-plane with output to the 60-panel tiled display. Groups of four display panels are driven by two GeForce 7900 GPUs. Interaction is facilitated by IR-reflective tracking systems from ART [5] utilizing ARTtrack2 IR-cameras to enable wireless head and hand tracking. In combination with the ART tracking system, our installation at SIGGRAPH 2007 piloted a custom optical markerless tracking system [6] for whole-bodied interaction within the Varrier™. Trackd® [7] is used to facilitate communication between the tracking system and the overall system software modules. Viewers position themselves at a location in front of the Varrier™ display, and within the tracking volume, that corresponds with a central viewing area at which immersion and autostereographic viewing are optimal. By positioning the entire data set in its contextual overview mode predominantly behind the implied image plane created by the surface of the LCD tiles, and progressively moving the data-patterns forwards towards the user in the coordinate space (Figure 3 below) while retaining this contextual overview in the background, ATLAS in silico makes use of this optimal viewing and tracking volume artistically and conceptually to enable the blending of quantitative and qualitative representation.

Translating the installation developed on the semi-circular Varrier™ tiled display to a passive stereo rear-projection system with rectilinear geometry required modification of all modules (graphics, audio, tracking). The current version of the installation on the C-wall uses the same positioning and sequencing of visual and auditory patterns as on the Varrier display, yet future research will include exploration of different sequencing and positioning of the data to place the viewer at different positions within the entire data set to make use of the more CAVE-like environment provided by not having to work within the tracking and optimal viewing volume inherent to the Varrier™ display. The audio system components in the C-wall system include 8 Genelec 8040 speakers in a rectangular arrangement. The system does not require a sub-woofer. Trackd® is used to take input from hand and head sensors to an Ascension Flock-of-Birds [8] tracking system. Future work will include continued development of our custom optical markerless computer vision

tracking system to replace the use of the ART IR-tracked hand controller (Varrier™) or head/hand inertial sensors (C-wall) and enable multi-user, multi-hand interaction.




Fig. 2: C-wall system implementation

2.3 Graphics Subsystem

Our implementation utilizes a variety of open-source software libraries. The graphics subsystem is developed in Visenso's COVISE [9] as the virtual reality framework, OpenSceneGraph [10] as the application level graphics interface, and OpenGL [11] for low level graphics routines. It uses Mechdyne/VRCO's trackd[®] for tracking. The audio subsystem uses Pure Data [12] as its interface to all system modules. GOS data and metadata is managed using MySQL.

The graphics subsystem, developed in COVISE's OpenCOVER [13] is the system's main software module and forms the foundation of our application. This module is started/initiated by the user at runtime. On the head node, in cluster-based systems like the Varrier™, it will spawn child processes on each of the rendering nodes. As soon as this application is running, it connects to the sound server using a network socket connection. Whether using the ART, Ascension or our custom tracking system, user input/tracking information is sent to the graphics module via a trackd[®] server. Our COVISE application reads the trackd[®] parameters for location and orientation of tracked head and hand on the cluster's head node and responds to them by updating the graphics/visualization. The rendering nodes on the Varrier™ cluster all run a copy of the COVISE application, receive event data from the head node, and wait for a command from the head node before updating imagery on the display tiles, so that the system stays synchronized. The head node similarly controls both the graphics and the audio components.

The major difference between the Varrier™ and C-Wall system implementation is that on the Varrier™, there is a head node and a set of rendering nodes, whereas on the C-Wall the head node also renders the images. It renders two images, one left and one right eye image for a 3D stereo effect. On the Varrier™, each node renders stereo images on four LCD tiles, which means that every node renders four times as many images as on the C-Wall system. If comparable hardware was used, the C-Wall system would be about 4 times faster than the Varrier™ system, but in reality the difference is not as pronounced because of differences in hardware and computational/networking overhead.

In either the Varrier™ or C-Wall systems, GOS data is accessed via a series of SQL queries that randomly select subsets of up to 100,000 records from a MySQL database containing the entire GOS data collection. This subset of records is used to initialize the virtual environment. Data values within each record include the amino-acid (ORF) sequence, biophysicochemical properties of this sequence such as amino acid composition, molecular weight, hydrophobicity and secondary structure characteristics, as well as the metadata values associated with its physical sampling location, and those associated with the sampling locations in which it was algorithmically identified. Reaping the full benefits of the GOS and metagenomics research overall requires the development of next generation computational and networking technologies to enable the scientific community to work with the scope and massive scale of this type of data. To reflect this broader socio-technological and ultimately human context, a set of contextual socio-economic, technological and environmental metadata descriptors for the regions nearest the geographical locations where the microbial samples were collected was selected and compiled to add this secondary level of context to the GOS data. Contextual metadata

representing CO2 emissions per capita, infant mortality rates and internet users per capita was collected from several internet-based publically available data repositories and combined with biological data from each of the 17.4 million GOS ORFs to drive the generation of visual and auditory patterns.

2.4 Aesthetic Influences And User Experience

Characteristics inherent in both of the above display systems influence formal artistic/aesthetic elements of the installation in the presentation of GOS data, its associated metadata, and contextual metadata within a hybrid strategy that merges quantitative and qualitative representation. This is accomplished by sequencing the presentation of the data, as represented by visual and auditory patterns, in different positions relative to the overall virtual world coordinate system and the user (real-world) coordinates, as well as by utilizing spatialized audio not only as audio objects within the virtual environment, but by spatialization strategies that position and move audio objects relative to the user according to both their interaction with the data patterns, and between and within data objects themselves. While the change in display geometry from the semi-cylindrical Varrier to the rectilinear geometry of the rear-projection system requires adjustment to all system components (graphics, audio and tracking), the installation remains conceptually and aesthetically the same on both displays. The user experience, graphics, and audio progress through three distinct but related modes each representing a different level of scale and resolution of the data. Figure 3 contains representative imagery from the installation for each mode. (Video is available at: <http://www.atlasinsilico.net/gallery.html>)



Fig. 3: (Left) Mode 1; (Center) Mode 2; (Right) Mode 3.

Progressing through each mode changes the scope and scale of data display from an entire subset of up to 100,000 records in mode 1, through a small subset of selected records in mode 2, to exploration of an individual record in mode 3. Each mode offers a distinct treatment of detail-in-context rendering for both visual and auditory representations. These representations themselves transition between modes to maintain coherency as they change in scale and relative amount of detail displayed from one mode to the next. A description of each mode is below.

Mode 1: Small colored “dots” (particles), one particle per GOS ORF, move in large patterns (ORF signatures) within the metadata environment. The entire data set environment is positioned two-thirds behind the plane of the display, and one-third in front of the display. Tracked hand motion brings up text that changes according to the data values in that region of the virtual environment. Text is positioned at the location of the tracked hand as particles (data records) collect in a region around the hand position. Generative audio responds to both hand motion/position and to the number of records collected versus those remaining in the overall pattern. The transition from this mode to the next is marked by the motion of the overall particle-system “freezing” and a color transition from multicolor to white on black indicating the data set pattern is taking its position as overall context for detailed exploration of data in subsequent modes. Simultaneous with this desaturation and back-grounding, each of a small subset (20 to 60) of the collected records (particles) transitions from their initial level of scale as a particle to a 3D geometric shape generated by a custom multi-scale shape grammar. (Note: The visualization schema and algorithms will be presented in future publications.)

Mode 2: Twenty to sixty geometric shapes of different size, number of vertices and color are arranged on a grid and are positioned in the plane of the display. They rotate on an axis centered on this image plane. Tracked hand motion in an area over a shape triggers a sonification of the data used to generate the 3D shape and the shape increases in size to provide a simple form of user feedback. The longer the viewer moves their hand over multiple shapes and “plays” their accompanying data sonification the longer this mode lasts. Hovering over a shape for a few seconds (hovering time depends on how long this mode has been active) selects that shape to progress to the next mode for detailed exploration. The shape increases in size, and moves out from the image plane towards the user, as the others fade out during a transition to the next mode.

Mode 3: The selected individual data object from mode 2 is centrally positioned, fills the majority of the field of view and is positioned three-fourths out in front of the image plane towards the physical location of the participant. A curvilinear line the same color as the shape is visible extending throughout the background. Tracked head and hand motion is used to probe and explore the shape by changing its position on an axis in all directions. Moving the viewer head position changes the point of view. Moving forward/backward allows the viewer to enter inside of the shape, and to exit the shape. The audio merges sonification and speech synthesis articulating the detailed data underlying both graphic and audio renderings. Text appears identifying all the regions of the metadata environment and a sonification tone sounds as each region label appears. The conclusion of this mode is marked by the large shape diminishing in size, scaling back in to a particle, merging with the desaturated/back-grounded data set. At the moment of its dissolution in to this larger pattern, the data set re-saturates with full color and particle motion resumes from the moment at which the pattern was frozen.

Regardless of system implementation (Varrier™ or C-wall), data is loaded from a MySQL database to the head node. There it gets translated into a set of points, which are fed into a simulation engine (Mode 1). The simulation engine updates the positions of all particles every frame. At runtime, the hand position attracts a subset of points, and upon a threshold triggers a selection which transitions the user experience to Mode2. This mode transition happens by first flattening the particle orbits and desaturating the visual field as the points increase in size to reveal their latent detail as full-size 3D objects. Throughout this process the audio system receives updated mode status values and data values to keep graphics, audio and user interaction in sync with each other.

3. AUDIO SYSTEMS: NARRATIVE, IMMERSION AND SONIFICATION

Sound in ATLAS in silico arises from three perspectives that are present simultaneously throughout: representation of data (sonification), support of interaction and immersion, and the embedding of these functionalities in a vessel of composed dynamic narrative that establishes a hybrid approach for data-driven aesthetic coherence facilitating data exploration. These interrelated perspectives are realized via a dynamic and modular audio system architecture and a hybrid approach for audio spatialization and localization that combines both amplitude and delay based panning.

3.1 Audio Subsystem

We utilize Pure Data (PD) to generate, process and adapt in real-time all sound within the immersive and interactive environment. The custom PD patch, running on a dedicated Linux system, receives a stream of control messages from the graphics rendering and tracking systems. A control protocol between all system modules manages the position and progress of the sound for the installation.

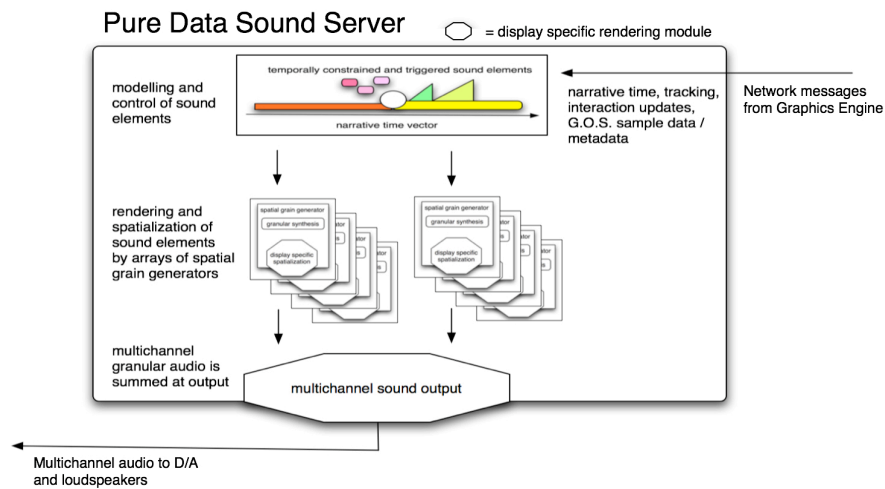


Fig. 4: Pure Data sound server schematic and functionality.

Responding to network messages relaying tracking values from user input, messages from the graphics engine, and all GOS data and contextual metadata values required for parameterization of the sonification, the PD server renders audio output to the multichannel audio display. In a control stage, objects are constrained in activity and processing allocation to certain ranges on a narrative time vector that relates the three interactive modes. Sound objects can be triggered by the progression of time or by specific interaction events and are listening for parameters from interaction/tracking and GOS data and metadata. They in turn trigger and control arrays of granular synthesis [14] generators that have integrated and exchangeable rendering modules to spatialize each grain on the display system with a specific spatialization strategy as defined in our algorithms. This process creates timing behaviors for the various elements, textures, events and layers of sound, as well as the spatialization and distribution of these elements on the multichannel speaker assembly. The resulting multi-channel audio (Figure 4) is forwarded to a display specific output module in which additional loudspeaker oriented calibration takes place. The multi-channel output is then sent to an external digital/analog converter where a final volume bias is applied before the sound is projected by the loudspeakers attached to the physical display. This modular spatialization approach allows the system to accommodate a variety of geometric configurations for the overall system and is thus applicable to both the Varrier and C-wall physical displays. (See Figures 1 and 2 above)

3.2 Narrative, Immersion and Sonification

Within our system, from the perspective of sonification, sound contributes a means of qualitative differentiability. In previous work we have attempted to establish perceptual domains by applying parametric spaces for sound synthesis that provide inexhaustible qualitative detail that we believe will allow for perceptible patterns to emerge from complex data [15]. The most notable ones are additive spectra and impact sounds. To represent multivariate GOS data within the installation, we developed a new conceptual auditory entity, the Scalable Auditory Data Signature (SADS) [16] that encapsulates different depths of data representation through a process of temporal and spatial scaling. On the smallest scale, a SADS represents itself as a microscopic data-determined impact sound from a specific location, displaying a minimum amount of data characteristics. On the largest scale, it can encompass all perceptual potentials of the participant in terms of spectral morphology, spatial resolution, and temporal behavior, and can eventually become a sonic world of its own.

From the perspective of immersion and afforded interaction with the data, sound is used to link visual elements representing data with their sonification. Here, concurrence and synchronization between auditory and visual events both spatially (object location) and temporally (interactive behavior) help to fuse sound, interaction and visual display into a coherent whole. This linkage is contextualized by Chion's theories on "Audio Vision" [17]. Sound that is attached to virtual objects in this way lends them physicality, increasing the sense of plausibility and immersion in the virtual world. These considerations are central to the design of cross-modal SADS. They also play a key role in all aspects of audio-visual design for ATLAS in silico in which sonification and interaction/immersion are embedded in a temporally propelled compositional narrative sequence that progresses through the three different modes (Figure 3), each characterized by a different level of scale and depth of data display, as well as a specific set of afforded interactions, progressing from global overview of the database down to a single data record with its contextual metadata. Criteria for spatial electronic music composition [18] are applied to support and propel this narrative progression, into which interactive affordances and the layer of sonification are interwoven. Modeling of the interactive narrative timeline and temporally constrained and parameterized sound rendering will be presented in a future publication. Below we present our work in hybrid sound spatialization and localization strategies implemented in the context of ATLAS in silico.

3.3 Hybrid Audio Spatialization and Localization

Sound spatialization is more than the placement of a sound source at a specific location. Depending on the strategy used, it changes the subjective quality of the perceived sound and the space the participants within a virtual or physical environment find themselves immersed in. To support the spatialization and localization of data-driven, sonification, user-responsive, and compositional/narrative auditory elements in the open immersive environment of the installation, we developed a hybrid strategy for interactive multichannel audio spatialization and localization. It combines both amplitude and delay strategies simultaneously for different layers of sound to support and highlight their intended purpose within the evolving sonic environment and achieve increased differentiation and perceptual depth between auditory elements as they progress through each of the three modes (Figure 3). Every mode is characterized by a distinct ambient auditory background texture that is layered with elements of interactive sonification.

In the context of virtual reality sound has often been regarded under the paradigm of the representation and simulation of architectural space, with a focus on the localization of sound within the virtual space [19][20], and the creation of simulated acoustic spaces with reflections and reverb, in which the virtual sound sources reside [21][22]. Design of the sound environment for ATLAS in silico poses a different set of questions. First, an architectural “inside” space is not intended as the visual representations hover in an open virtual universe. Second, while the rendered reflections and echoes added by a room simulation may enhance a sense of “realism”, they may also cover up details in the microstructure of the sounds and thereby potentially limit the qualitative differentiability. Although a strategy in which reflections and virtual acoustic room design are used to generate qualitative differences in their own right could be useful for sonification purposes, large scale display systems like the ones currently used for ATLAS in silico, often reside within acoustically problematic physical spaces such as a galleries or labs which add confounding characteristics to this mix thereby making them unsuitable for this type of approach. Consequently, rather than generating an auditory “virtual reality” for ATLAS in silico we utilize sound localization as an expressive modality and conceptual vehicle in a manner analogous to how Electroacoustic Music [18] uses spatialization and localization as a sound property akin to timbre.

Sound events are spatialized conceptually to support immersion by linking of visual and auditory display elements at all levels of scale, and to support events and interactions that emerge and are encountered during the course of all three modes. For sonification, spatialization by itself has shown to be of inferior perceptual resolution and complexity in comparison to other sound properties, such as the enormous perceptual resolution of frequency and spectral perception [15]. Nevertheless, spatialization is used for data display purposes in some instances in ATLAS in silico, for example pertaining to the characteristic of “fuzziness” of the sound localization once audio-visual fusion has been established. In addition, we utilize the spatial resolution of multiple audio elements over a multi-speaker system to allow participants to selectively shift attention between multiple simultaneous streams of sonification, following the idea of auditory streams in Bregman’s “Auditory Scene Analysis” [23]. Based on these considerations, spatialization is central in the design and progression of the composed narrative layer of sound in ATLAS in silico in support of the process of zooming in on detail within the data and unfolding its complexity and context. Future iterations of the installation may include wave field synthesis [24] as an even more versatile vehicle for hybrid spatialization.

3.4 Hybrid sound rendering

To effectively utilize the blended perceptual properties of delay and amplitude based spatial rendering a module was implemented in PD that allows us to shift parameterization continuously between amplitude panning based on the principles of VBAP [19] and to spread out panning with deployable channel delays to enable the precedence effect. This enables the sound for ATLAS in silico to generate a rich and multi-varied spatial sonic environment without the use of virtual acoustic modeling that may cloud the micro-details of the granular synthesis oriented sound. Delay based panning utilizes a psychoacoustic property of our hearing apparatus which collapses multiple correlated copies of the same sound signal arriving at a listeners head from different directions with time offsets into a single sound source dominated by properties of the signal arriving with the first wave front [25]. An interesting property of the resulting psychoacoustic point source illusion is that it deteriorates with a growing decorrelation of the arriving signals. The more decorrelated the signals are, the less they are collapsed into a single perceived point source and instead are heard as having increased width [26][27]. The difficulty of generating correlated signals makes using this technique for multi-channel audio rendering without the addition of amplitude biases difficult. Yet, the added width of the sound source and the resulting qualitative properties can be utilized as design elements to increase subjective properties of the sound such as the richness and immersion.




Fig. 5: Distance based generation of delays (t_n) and amplitude (a_n) values for different display geometries.

A second consideration in the use of delay panning is the “sweet spot” area in which the desired localization can be perceived or the dependency of the localization effect on the position of the participant within the sound field. Both of the display systems for ATLAS in silico present geometric footprints and interactive modalities that include a central interacting participant surrounded by a large group of additional viewers/listeners. Therefore our algorithm for delay based panning differs in this respect from strategies presented in [28] and [29] that focus on a centered listening perspective and also attempt to encode spatial depth in the delay signals – something that we discard in order to increase the control over the sound source location for a larger listening area. This is why we use delays based on the display system geometry and not values that are derived from head-related transfer functions. To minimize delays while maximizing the listening area, our approach moves a virtual sound source along the plane of the loudspeakers. In case of the Varrier™ display this plane is a half-cylinder, whereas on the C-Wall, the plane is rectangular. The delay at which a signal is projected by a loudspeaker is proportional to the distance between its virtual location and the loudspeaker.

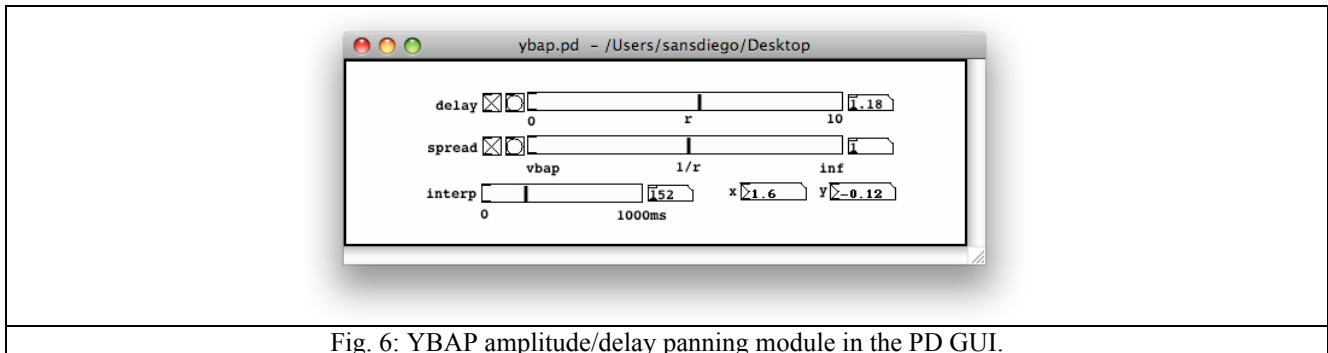


Fig. 6: YBAP amplitude/delay panning module in the PD GUI.

In order to allow a continuous interpolation between pure amplitude and pure delay based panning with all intermediate steps, our algorithms needed to be interpolated. This allows the spread of the sound source to be adjusted from the minimum that VBAP requires to generate a sound localization to $1/r$, which corresponds to a natural drop off, and from there to 0, in which all loudspeakers receive equal amplitudes and therefore only delay affects the localization. Secondly, the delay needs to be adjustable from 0 (pure amplitude panning) to 1 (delay exactly proportionate to distance/speed of sound) and further on to allow an exaggeration of the delay effect. The rules of sum amplitude established by Pulkki [19] are implemented in our schema in order to equalize the perceived amplitude of the virtual sound source with that of the spatialized sound source.

4. CONCLUSIONS AND FUTURE WORK

Our exploration of the combination of artistic and data-driven strategies for working with massive multiscale multidimensional data focuses on a combination of qualitative and quantitative aesthetics and representation. Within ATALS in silico, our graphics subsystem application allows qualitative observations when GOS data is displayed as ORF signatures, and quantitative observations when the user focuses on a single ORF in subsequent modes. The

combination of qualitative and quantitative representation is manifested throughout all visual and auditory elements in our system. A major challenge in this work is creating a system with the capability of presenting the entire set of 17.4 million GOS ORF records simultaneously in a format amenable to real-time interaction. The current prototype system allows subsets of up to 100,000 records to be displayed simultaneously at interactive frame rates. In future work we plan on resolving the existing bottle necks in our simulation engine to enable display of and interaction with the entire GOS data set. The simulation and visualization/rendering of particles in the system each take a certain amount of time and therefore we hope to explore approaches such as parallelization of our algorithms and the use of an additional PC cluster or porting the simulation to CUDA-capable graphics cards. We will explore optimizing the graphics output by various methods including different point rendering algorithms, or by using an adaptive “intelligent” algorithm to display only those points that the user actually sees. Future work in regards to user experience and interaction is required to develop a user interface that would enable this prototype system to function as a scientific research tool. This work will reach towards a unified seamless experience that allows users to move between all levels of scale and representational detail in an open-ended flow rather than progressing through three modes as described above. As part of this we will continue to develop a camera-based optical markerless tracking system to track multiple participants interacting in the environment. While our audio rendering subsystem allows settings that create artifacts potentially undesirable in the context of “pure spatialization” our implementation demonstrates that the side effects of spatialization can be acceptable if they are used consciously to achieve a specific sound design goal, and especially in the context of synthetic sounds for which criteria of “naturalness” do not necessarily exist. We developed hybrid amplitude-and-delay based spatialization as an alternative to Ambisonics and VBAP. It provides greater flexibility in generating spatial perceptions and impressions from monophonic sound sources than these alternate approaches. In the context of ATLAS in silico, this hybrid rendering technique allows elements with very clear localization as well as the rendering of seemingly large sounds with powerful kinetic energy. Our future work on this intersection between physical and mathematical sound rendering in the service of creative sound design will develop additional intersections between data-driven and aesthetic approaches in the context of this hybrid schema. Our ultimate goal is to progress this prototype system to create a more generally applicable tool for exploration of massive multiscale, multiresolution, multidimensional data with no inherent spatial structure.

REFERENCES

- [1] Yooseph S, Sutton G, Rusch DB, Halpern AL, Williamson SJ, et al. “The Sorcerer II Global Ocean Sampling Expedition: Expanding the Universe of Protein Families.” *PLoS Biol* 5(3) (2007).
- [2] Smith J., [Charles Darwin and Victorian Visual Culture], Cambridge University Press, New York (2006).
- [3] Online: “[http:// web.camera. calit2.net/ cameraweb/ gwt/ org.jcvi. camera. web. gwt. download. BrowseProjectsPage/ BrowseProjectsPage.oa#ViewProject.](http://web.camera.calit2.net/cameraweb/gwt/org.jcvi.camera.web.gwt.download.BrowseProjectsPage/BrowseProjectsPage.oa#ViewProject)”
- [4] Sandin D, Margolis R, Ge J, Girado J, Peterka T, and DeFanti T. “The Varrier™ autostereoscopic virtual reality display.” *Proceedings of ACM SIGGRAPH*, 24(3): 894-903, (2005).
- [5] Online: “[http://www.ar-tracking.de.](http://www.ar-tracking.de)”
- [6] Video online: “http://www.atlasinsilico.net/gallery_siggraph2007.html”
- [7] Online: “<http://www.mechdyne.com/integratedSolutions/software/products/trackd/trackd.htm>”
- [8] Online: “<http://www.ascension-tech.com/realtime/RTflockofBIRDS.php>”
- [9] Rantzau D, et al., “Collaborative and Interactive Visualization in a Distributed High Performance Software Environment.” *Proc. International Workshop on High Performance Computing for Graphics and Visualization*, Swansea, Wales, (1995).
- [10] Online: “<http://www.openscenegraph.org/projects/osg>”
- [11] Online: “<http://www.opengl.org/>”
- [12] Online: “<http://www.crca.ucsd.edu/~msp/software.html>”
- [13] Rantzau D, Frank K, Lang U, Rainer D, Woessner U. “COVISE in the CUBE: An Environment for Analyzing Large and Complex Simulation Data.” *Proceedings of the IPTW* (1998).
- [14] Roads C. [Microsound], Cambridge, MIT Press (2001).
- [15] Gossmann J, “Towards an Auditory Representation of Complexity.” *Proc. of ICAD 05-Eleventh Meeting of the International Conference on Auditory Display*, Limerick, Ireland, July 6-9, 2005.
- [16] Gossmann J. et al., “Scalable Auditory Data Signatures for Discovery Oriented Browsing in an Expressive Context.” *Proc. of the 14 International Conference on Auditory Display*, Paris, France, June 24-27 (2008).
- [17] Chion M, Gorbman C, and Murch W. [Audio-Vision], Columbia University Press, (1994).

- [18] Otondo F. "Contemporary trends in the use of space in electroacoustic music." *Organized Sound* 13(1): 77-81 (2008).
- [19] Pulkki V. [Spatial sound generation and perception by amplitude panning techniques] Espoo, Helsinki University of Technology, (2001).
- [20] Eckel G. "Applications of the CyberStage Sound Server." in *Proc. Audio Engineering Society*, March, (1999).
- [21] Jot J-M, "Real-time spatial processing of sounds for music, multimedia and interactive human-computer interfaces." *Multimedia Syst.* 7(1): 55-69 (1999).
- [22] Funkhouser T, Min P, and Carlbom I. "Real-time acoustic modeling for distributed virtual environments." *Proc. of ACM SIGGRAPH*, 365-374 (1999).
- [23] Bregman AS. [Auditory Scene Analysis: The Perceptual Organization of Sound], MIT Press, (1994).
- [24] Theile G and Wittek H. "Wave field synthesis: A promising spatial audio rendering concept." *Acoustical Science and Technology*, 25(6): 393-399 (2004).
- [25] Blauert J. [Spatial Hearing - Revised Edition: The Psychophysics of Human Sound Localization]. The MIT Press, (1996).
- [26] Kendall GS. "The Decorrelation of Audio Signals and Its Impact on Spatial Imagery." *Computer Music Journal* (19)4: 71-87, Winter (1995).
- [27] Potard G. and Burnett I. "Decorrelation Techniques for the Rendering of Apparent Sound Source Width in 3D Audio Displays." in *Proc. of the 7th Int. Conference on Digital Audio Effects (DAFx'04)*, Naples, Italy, October 5-8, (2004).
- [28] Moore FR. "Spatialization of sounds over loudspeakers." in [Current directions in computer music research], MIT Press, 89-103 (1989).
- [29] Braasch J. and Woszczyk W. A " 'Tonmeister' approach to the positioning of sound sources in a multichannel audio system." *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, 130-133 (2005).

ACKNOWLEDGEMENTS

This work was in part supported by National Science Foundation IIS-0841031. Additional support was provided by Calit2, CRCA, CAMERA, Ingenuity Festival, TimeLogic, CENS, NCMIR, EVL and SDSC. Installation hardware and software is provided by Da-Lite Screen Company, VRCO/Mechdyne, Meyer Sound, and mentalimages. Creating ATLAS in silico is a collaborative effort involving: Alex S. Horn, Weizhong Li, Trevor Henthorn, Rajvikram Singh, Javier I. Girardo, Sam Fernald, Toshiro Yamada, and student researchers from both the UCSD ECE 191 design practicum and UCSD Computer Science.