SILVA: AN INTERACTIVE AUDIOVISUAL INSTALLATION

by

Randiel Zoquier

SENIOR THESIS, B.A.

Submitted to the Faculty of the Stevens Institute of Technology
in partial fulfillment of the requirements for the degree of

BACHELOR OF ARTS

Randiel Zoquier, Candidate

ADVISORY COMMITTEE

5/4/22

Lainie Fefferman, Advisor Date

STEVENS INSTITUTE OF TECHNOLOGY
Castle Point on Hudson
Hoboken, NJ 07030
2022

SILVA: AN INTERACTIVE AUDIOVISUAL INSTALLATION

ABSTRACT

Silva is an interactive audiovisual installation that explores the communal process of performing and listening to music, especially for those who have no experience of music performance. Participants connect to a website on their mobile device to control virtual instruments or participate as part of the "audience," and the changes made from the mobile device are reflected in the musical and timbral content of the installation as well as the accompanying light show. Silva is structured as a ring of speakers and vertical light strips, which are all controlled from a central computer. Participants of the installation can freely move around in the center of the ring of speakers to experience the ambisonic soundscape as well as interact with other participants. The musical content of the installation is presented as two scenes, the ocean and the forest, where each scene explores different musical landscapes spanning from ambient electronic music to traditional Latin music. This paper explores the production and functionality of Silva, reflecting on the process of designing and building an installation from start to finish, and the different reactions from participants.

Author: Randiel Zoquier

Advisor: Lainie Fefferman

Date: May 4th, 2022

Department: Music and Technology

Degree: Bachelor of Arts

ACKNOWLEDGEMENTS

Special thanks to the following people who helped in the production and display of the installation:

Mario Diaz de Leon

Lainie Fefferman

Robin Hammerman

Robert Harari

Nancy Nowacek

Kaitlin Osborne

Andrew Stein

Kelland Thomas

Jeff Thompson

TABLE OF CONTENTS

Abstract	**iii.**
Acknowledgements	**iv.**
Introduction	**1**
Inspiration and Production Process	**2**
How the Installation Works	**7**
Displaying the Installation and Reflections	**12**
Conclusion	**14**
References	**16**

Introduction

Silva is a multi-channel interactive audiovisual installation designed to allow participants to explore the communal process of performing music and experiencing sound around them. It is meant to foster curiosity and playfulness, especially for participants who do not have any experience in musical performance. Participants interact with the installation by connecting to a web application, which allows them to control the resultant audio and light show. Participants can choose from four main instruments to control: drums, bass, rhythm, and lead. Participants can also choose to be part of the "audience" which controls the ambient sounds of the installation as more audience members join. The content of the installation is split between two scenes, the ocean and the forest. The ocean scene incorporates timbres from electronic music and inspired by video game soundtracks, while the forest scene incorporates timbres and instrumentation from traditional music from the Dominican Republic.

The speaker system consists of fifteen JBL 705i speakers powered by two Crown DCi 8|300N amplifiers, driven by a MOTU 24Ao interface. The speakers are arranged in a ring, with eight speakers placed high and pointed down, seven speakers placed low and pointed up, all equidistant from the center of the ring. The lighting system consists of twelve custom-built vertical LED light strips, dispersed on light stands in a ring among the speakers. The installation is controlled from a central computer running Max/MSP and QLC+, which receives data coming from the web server to control the audio and lights.

The production process of the installation took about eight months, where the initial planning and research started in late August of 2021. The first few months were spent planning the installation content, researching various technologies, creating small test projects, and finalizing a budget for building lights. The main production and construction of the installation started in February of 2022 and was completed in its first form on the 9th of April for an Accepted Students Weekend event. Additional tweaks were made to the installation and it was then presented again at Innovation Expo on April 29th.

Inspiration and Production Process

This project was primarily inspired by my personal experiences of audiovisual installations. I was amazed by the enveloping feeling of the installations I have visited, and inspired by how a perfect combination of audio and visuals could promote a sense of awe and wonder. The particular installations that inspired me are *Deep Web* by Robert Henke and Christopher Bauder, *Celestial* by ARTECHOUSE, and *Geometric Properties* by Julius Horsthuis. *Deep Web* intrigued me with its use of lasers and illuminated spheres which created nodes and lines of light, underscored with an ambient soundtrack consisting of electronic bleeps and bloops to mimic digital communication. *Celestial* and *Geometric Properties* were both displayed in ARTECHOUSE's gallery space in New York City, where the wall-to-wall projections and immersive audio system transported me to the worlds that the artists were depicting, such as swirling galactical spaces and infinitely folding fractals. After experiencing these installations, I wanted to see if I could

develop something similar, especially since I have always had an interest in the intersection of audio and visual technology.

In deciding the construction and content of the installation, I first considered the basics of what I wanted the participants to experience and the technology I had available to me. I knew I wanted an immersive audio experience and that the Music and Technology program has a multi-channel speaker system that I could access, so I built my initial conceptions around using those speakers for my audio output. I also knew I wanted the experience to be interactive, because I wanted to give participants agency when viewing and participating in the installation. I wanted participants to have this agency so they could have a deeper connection to the experience and so there would be a layer of playfulness to the installation. My first consideration for interaction was having the participants use their mobile devices, since most people have cell phones and it offers many affordances with a touch screen and display. I briefly considered other input modalities such as cameras or game controllers, but I was heavily drawn to the mobile phone since it is a familiar input device for participants and I wanted to explore the use of cell phones in audiovisual installations. I decided that a web interface would be the best choice since it simplifies compatibility between different devices and I have experience in programming websites. For lighting, I first spent time deciding the physical shape I wanted for the lighting, eventually landing on vertical strips of light to mimic trees in a forest. I then researched existing lighting technologies and decided that it would be more economical to build lights from scratch than to use off-the-shelf units.

What took the longest time to completely develop was the content of the installation and the interactive experience. I started off by writing a few word maps to brainstorm ideas that compelled me and sonorities I wanted to explore. The things that stood out to me were themes of nature, curiosity, discovery, ambience, and community. These specific ideas inspired me to structure the installation as a communal performance. I thought about my experiences of performing music with other musicians and the joyous feeling of discovering how my instrument fits in the texture of an ensemble, and I thought that would be a compelling form of interaction within the installation. Curiosity could then be created by offering a minimal yet intuitive interactive experience, so participants would have the opportunity to discover what the controls could do. Going further with the analogy of musical performance, I drafted a plan where four participants act as main instrumentalists and then the rest of the participants act as the "audience," changing the ambience and feeling of space in the installation. Using my word maps I decided on basing the musical and visual content on two distinct scenes, an ocean scene and a forest scene. With these concepts in mind, I was then able to start planning and structuring the technology needed to develop the installation.

I had to do a considerable amount of research to figure out which technologies I was going to use to construct the installation and how to properly implement it. I knew that I needed to learn the basics of dynamic web development, so I followed a basic course in node.js to learn how to build dynamic web applications. I then found that the socket.io library was one of the technologies I could use to create real-time communication between the client and server in the way that I needed for the installation.

I practiced using socket.io by creating a simple tic-tac-toe web app, which really helped me understand how to use the technology, work out any kinks and bugs, and get it running on my personal web server. I then learned how to incorporate socket.io into Max/MSP, which is the software I decided to use to handle all the audio.

In terms of audio, I had to test and learn about various objects in Max/MSP which I would need to use to play the audio in the way I wanted. It took a few tries and tinkering to decide which one gave the best results, but I was then able to refine and scale the process for the entirety of the program. I also had to research how ambisonic technology worked, and how I would translate my audio channels to a multi-channel speaker system. I was glad to find a comprehensive suite of plugins by the Institute of Electronic Music and Acoustics that really helped me apply what I learned about ambisonic audio and easily incorporate it into the installation.

Before I composed any of the music I wanted to gather inspiration from the music I listen to and research new genres so I could have a clear objective in how I wanted everything to sound. I first decided the basic instrumentation for each scene, since the instrumentation would be the basis of the user interaction. I wanted there to be distinct instruments that could meld together and create a full sound experience, so I decided on the familiar ensemble of drums, bass, a rhythm instrument like keys or guitar, and a lead instrument like a lead synth. After deciding the instruments, I decided on the genres and sonorities I wanted to explore with each scene. For the ocean scene, I decided on a dark and washy electronic sound, and I looked to some of my favorite video games and the music they use in their ocean levels to inspire the musical content. For the forest scene, I

decided on exploring traditional music from the Dominican Republic, to dive deeper into my family's culture and incorporate more aspects of my identity into the installation. The woody and dancy textures contrast the calm and ambient atmosphere of the ocean scene, so I felt it would be a great musical genre to explore. I did not have much experience in creating Dominican music, so I researched the primary genres of merengue and bachata, learning about the instruments traditionally used in the ensembles and the types of rhythms and melodies played. I borrowed my uncle's *tambora* and my cousin's *guirra*, which are traditional Dominican percussion instruments, to record the music for this installation. This musical research helped to properly define and focus my compositional process and create an effective musical landscape.

The process of developing the lighting system required a significant amount of research as well. I decided I wanted to use LED light strips to construct tubes of light, and I wanted to control the lights with traditional lighting control technology such as DMX. I did a lot of searching for particular controllers and light strips, but what really advanced me in my research was speaking to Stephen Cowan, a lighting director I met while seeing Slow Magic at Elsewhere in Brooklyn. I was interested in the tube lighting system at the venue, which was very similar to what I was trying to achieve with my lighting setup, and he explained to me how those kinds of lights are typically constructed. He pointed me towards using NeoPixel LEDs for individual pixel control as well as the Advatek lighting controller which I ended up using in the final installation. I then researched power requirements for all the LEDs I would use, the materials I would need to construct the housings, and the best way to organize and wire everything together. This

allowed me to craft a budget proposal for ordering all the parts, and effectively construct the lights when the time came.

Iteration was an essential part of the project process. I constantly consulted my advisor about how the installation experience sounds and feels, and used those notes to improve upon the experience and fine tune my ideas. For the user interface, I consulted my graphic design professor Nancy Nowacek for notes on improving the interface and giving the user just the right amount of information to foster curiosity and discovery. At many points during the process I stopped and re-evaluated the experience, asking myself questions about the simplicity of the experience, how audible the changes were from the user perspective, and how polished the experience feels. Many little changes added up to the final installation as it stands, and many initial ideas changed over the course of the production process to address new problems and new desires that arose. For example, I wanted to create field recordings to act as part of the ambience in the installation, and initially I tried recording some ambisonic audio using an ambisonic microphone. For a multitude of reasons, I was unable to get the ambisonic recording to properly translate into the installation space, so instead I created field recordings using a handheld stereo recorder and placed those recordings in the ambisonic space like the rest of the instrument tracks. This could be seen as a setback, but it allowed me to move the ambience around the space and give the audience members a tangible layer of control.

The production and development process of this installation was a great exercise in concept development, research and planning, and most importantly, problem solving. Being flexible with available resources and diligent with time management was the only

way I was able to effectively create the installation and develop it into the state where it is now. The skills and the lessons learned from this production process will translate effectively into any long-term and high-concept project.

How the Installation Works

This installation has many moving parts that create the final immersive experience, all communicating actively to allow for seamless interaction. The main parts of the installation are the client website which handles user interaction, the web server which handles data routing, the Max/MSP patch which handles audio playback and mixing, and the QLC+ workspace which handles the light display. The main technology that serves as the backbone of the experience is socket.io, which allows for near real-time communication between the client, server, and the Max/MSP patch. Here is a brief overview of how the installation functions.

The client website is the main venue of participation for the installation, and there are a few different coding technologies to make it function. The most important is socket.io, which is how the client connects to the main server and sends message data to the server. When the user first opens the website, the server sends preliminary data about the status of the installation to the client, such as which instruments are currently chosen, whether the installation is active or not, which scene is currently running, and time data for rendering the loop progress bar. The graphical interface is then rendered with p5.js, which also handles the touch input. Each user interface element is custom-coded to render in a certain way and format the data so whenever a button is pressed, a slider is moved, or a state is changed, data is sent from the client to the web server using socket.io

messages. The interface is mainly rendered as simple shapes and lines, although some parts of the interface like the instrument icons and sequence icons are images that were created beforehand in Adobe Illustrator.

The web server program is like the traffic cop of the whole operation, handling incoming requests, routing data, and keeping track of general information. The server stores which instruments are currently selected, which specific socket.io client is the one that represents the Max/MSP patch, which scene is currently playing, and which clients are currently acting as audience members. Whenever a new client connects, the server checks to see if the client is the Max/MSP client, and if it is the Max/MSP client the server stores its specific client ID for later use. If the new client is just a normal user, the server sends all the important state data to the client so the new client is up to date. The server also handles disconnects, so when any client disconnects from the server the appropriate information is reset or updated. For example, if someone closes their web browser tab after they are playing an instrument, that instrument is freed up and the whole installation is updated. Finally, the server receives input data from the instrument clients and sends it to the Max/MSP client so it can update the audio experience. Since the server keeps track of which client is the one controlled by Max/MSP, the information can be sent to the right client. Once in a while, the Max/MSP client will also send to the server important information such as scene changes and downbeat messages, so the appropriate UI elements in the client can be updated.

The Max/MSP patch handles the routing of messages coming from the server to the instruments, audio playback, audio mixing, and the outputting of MIDI data to QLC+

to control the lights. The routing subpatch takes the incoming message data from the web server and parses it so the proper instrument information can be updated. For example, a message from the server can look like “0 sl 0.6,” and the patch then unpacks the message so the value “0.6” can be sent to the slider input of the drum instrument. The messages coming from the server that get routed are input messages, instrument choice messages, and audience messages.

Another essential part of the Max/MSP patch is the clocking subpatch. It handles the transport object, which controls the timing of the audio files so they stay in sync, and then resets the transport to the beginning after eight bars to create a continuous loop. Another part of the timing subpatch is the loop counter, which counts how many loops have passed to switch the scenes after fifteen loops. The clocking subpatch also sends downbeat data to the web server so every client can have a progress bar that matches the loop progress.

Each instrument then has its own subpatch for each scene, which each have different functions specific to each instrument, but in general work very similarly. The ambience, drums, bass, rhythm, and audience instruments all are based around audio loops. Each audio loop is played through a [groove~] object, which is synced to the main transport to keep all the loops in time. To prevent clicking at a zero crossing point when a loop starts and stops, a [trapezoid~] object ducks the audio volume at the loop endpoints, smoothing out any rough edges. A [matrix~] object is used to smoothly switch between a choice of four different loops coming from their own [groove~] objects, and then the audio is processed by filters and effects. From those effects the audio from each

instrument is routed to the main mixer subpatch. The effect parameters and loop switching is handled by incoming messages from the routing subpatch. The only difference for the lead instruments is that instead of loop switching messages they receive note on and note off messages to play software synthesizers.

After all the audio is processed, it is sent to the main mixing subpatcher. The signal flow in the mixer is quite complex, since it has controls for smoothly switching between the two scenes (which means the audio for both scenes are technically running simultaneously), turning on and off instruments, automatic panning, and then handling the various ambisonic aspects of the experience. To simplify, there are two halves of the mixer: the instrument mixing and the reverb/delay mixing. For the instrument mixing, each instrument audio track is patched into a MultiEncoder plugin, which allows each track to be placed in 360 degrees in the ambisonic domain. For the reverb/delay mixing, some instruments are run through separate delays and reverbs, which are either placed ambisonically with StereoEncoder, or in the case of the delays, processed directly into the ambisonic domain. When the signals are encoded into the ambisonic domain, they are output into 16 channels, which is the amount of channels needed to encode and decode 3rd order ambisonics. With a higher order of ambisonics, the more specific and acute the directionality of the experience becomes, at the cost of higher channel count. I found that 3rd order ambisonics was a good balance of directionality and channel count. Once the variety of 16-channels signals from the instrument mixer and the reverb/delay mixers are combined, they are sent to the decoders. When I was developing the installation from home, I used the BinauralDecoder plugin, which allowed me to experience the ambisonic

audio binaurally with headphones. For the final installation, I decoded the ambisonic signal using the AIIRADecoder, which allowed me to map the specific speaker setup I was using and play ambisonic audio through the 15 speakers. The flexibility of the AIIRADecoder means I can adapt the installation to any speaker setup I could ever want.

Finally, as a supplement to all of the complex signal flow, data from the installation is translated into MIDI information to send to QLC+ for lighting control. To keep things simple, the only data sent to QLC+ are the scene changes, the instrument states (on/off) and some audio levels for each instrument so the lights could pulse along. Panning data has been ignored for these early iterations of this installation, but I would love to explore it in future iterations for an even more immersive experience.

The custom-built lighting system is controlled by QLC+, which is an open source software lighting controller. The lights are constructed as 12 tubes, each with an LED light strip of 30 pixels each, for a total of 360 individual pixels. The LED strips are mounted in an aluminum channel and surrounded by a plastic tube. Custom designed 3D-printed end caps hold onto the aluminum channel and plastic tube, while also allowing for mounting onto light stands. The way the QLC+ communicates with the LED light strips is with an Advatek pixel controller, a special board that treats LED light strips as a networked lighting system, communicating using the sACN protocol. The Advatek has four outputs, which I sent into four custom-built breakout boxes. The breakout boxes handle three lights at a time, and are wired so each light is wired sequentially. QLC+ has special software functions that map simple patterns to the pixel array, which I programmed to create an ambience and react to the incoming MIDI data from Max.

Though some of the built-in pixel functions were somewhat lackluster, I was able to get a result I was happy with after some tweaking.

All of the code for this project is available on my GitHub, at this link: https://github.com/Adv3ntur3rz/Silva/.

Displaying the Installation and Reflections

As much as I could work on each individual part of the installation in isolation and enjoy each part of it, there is nothing like the experience of putting it all together and sharing it with people in its entirety. The setup process is quite simple: arrange the speakers on stands, plug the speakers into the amplifiers and audio interface, arrange the lights, plug them into the breakout boxes and lighting controller, and finally plug in the audio interface and lighting controller into my laptop. The first time I had the installation displayed was for Accepted Students Weekend on April 9th and 10th, in Babbio 310. It was an exhilarating experience to see my vision realized in its full form. However, the first full run-through was not without its hiccups, and an issue with the system running for a long period of time causes some nasty audio glitches that incapacited the installation for about a half hour. After some changes to the MAX/MSP patch to make the CPU load lighter, and some adjustments to the rack spacing to allow for adequate cooling of the amplifiers and audio interface, the installation ran smoothly on the second day of the weekend.

It was interesting to see the different reactions from people interacting with the installation. I intentionally gave people as little information about the experience as possible so that they could really explore the experience and discover what the controls

do. I had a poster set up explaining the basic premise of the installation and a QR code so people can scan and join the website. Some people were initially confused about what was going on, but after I gave them a short demonstration with my phone, they caught on quickly and enjoyed tinkering with the experience. Some people were really amazed by the interactive experience and the immersive sound, while others silently took everything in. People were also interested in understanding how the whole installation functioned, and since the Max/MSP patch was on my laptop near the installation, I brought them over to see the brains of the operation. Musicians really seemed to enjoy the experience, especially with playing the lead instrument, where they could really explore a familiar musical interface. Non-musicians seemed to enjoy the installation as well, which was comforting because that is the main target audience for the installation. It was great to see the many reactions and interactions that people had with the installation, as well as the unique perspectives everyone brought to the experience. The goal to foster joy and curiosity was certainly achieved, and I hope to be able to continue to display this installation for more people to experience.

Conclusion

Working on Silva was a tremendous learning experience and an opportunity of growth for me as an artist. The production process was among the longest I have done, and it really gave me perspective on the intricacies of taking a project from initial ideation and developing it to a final project. I gained many useful skills that I might not have learned otherwise, such as CAD design for the 3D-printed parts, soldering and electrical design, advanced website networking and creating dynamic web applications,

and more. There are definitely many things I would have done differently, especially with constructing and designing the lighting system, but I am still happy with the end result. I learned how certain things take a lot more time than anticipated, and so planning a lot of extra time is essential. I also learned how to be flexible and compromise with certain aspects of the experience, while still sticking to the initial purpose of the project.

In the future I would like to present this installation in a variety of venues, so more people can enjoy the experience. I would also like to expand on this concept of shared interaction and performance, evolving the current installation to more high-definition audio systems or developing the lighting system even further with more precise control. In the future I would also like to design entirely new installations based on the same concept but exploring it from different angles: How would the installation feel if everyone has the same instrument? How would the experience change if I used projections instead of lights? What if I created a networked virtual reality experience where people could connect all over the world? The sky is truly the limit in evolving this project, and I cannot wait to see where it goes.

REFERENCES

Socket.IO, https://socket.io/.

ARTECHOUSE. “Celestial.” *Artechouse*, 21 February 2021, https://www.artechouse.com/program/celestial/.

Bauder, Christopher, and Robert Henke. “Deep Web.” *Robert Henke*, https://www.roberthenke.com/installations/deepweb.html.

Burgess, Phillip. “The Magic of NeoPixels | Adafruit NeoPixel Überguide | Adafruit Learning System.” *Adafruit Learning System*, 30 August 2013, https://learn.adafruit.com/adafruit-neopixel-uberguide.

HORSTHUIS, JULIUS. “Geometric Properties | Experience Now at ARTECHOUSE.” *Artechouse*, 1 March 2021, https://www.artechouse.com/program/geometric-properties/.

Institute of Electronic Music and Acoustics. “IEM Plug-in Suite.” *IEM Plug-in Suite*, https://plugins.iem.at/.