California Institute of the Arts

Next Music

Exploring Music in the Augmented Space

by

Colin Honigman

A thesis submitted in partial fulfillment for
the degree of Master of Fine Arts

in the
Herb Alpert School of Music
Music Technology: Interaction, Intelligence & Design

Supervisory Committee

Dr. Ajay Kapur

Mentor

Dr. Jordan Hochenbaum

Committee

Dr. Owen Vallis

Committee

Kenneth Cameron

Committee

Randall Packer

Committee

Abstract

In this document the idea of Next Music is introduced and explored by way of three projects created by Colin Honigman during the course of earning his Master of Fine Arts degree in Music Technology at the California Institute of the Arts. Next Music is the author's own term for the emerging musical practices that are tied to the augmented space. An augmented space is a physical space that has been enhanced through computing and sensing. By combining these elements a new kind of space emerges, one that has immense potential for artistic endeavors. In music, this kind of space can be utilized in many contexts, including performance, installation, studio, composition, and even pedagogical. This document explores Next Music by approaching the augmented space with different combinations of scale and context. Presented in order from smallest to largest these projects are Cultivating Frequencies, The Third Room, and Hear in LA. Cultivating Frequencies is an interactive garden that collects data on plant health and the environment and transforms them into sound through a generative musical system, extending the interaction to the observer by way of touch sensitive plants. The Third Room is an interactive performance space, utilizing multi-user-tracking and depth camera technologies to place a user/users in a virtual reflection of their environment, populated with its own unique set of instruments, allowing them freedom to explore and create, alone or in groups. Hear in LA is a mobile device application that reconnects Los Angeles drivers with the cultural landscapes they drive through each day by allowing them to hear the sounds of the neighborhoods they travel through and even record and upload their own sounds as well. Together, these projects only represent a small portion of what is possible when working in the augmented space. As technology brings us closer to ubiquitous computing and sensing, many of our spaces are already becoming augmented. Artists have a unique chance to participate and influence what the augmented space becomes by showing the world what is possible. Next Music, and this document, is an introduction to the musical possibilities within the augmented space and an open invitation to take part in the making this new space.

Acknowledgments

This thesis would not have been possible without the help, encouragement, and guidance of a few people. First of all, to my parents, thank you for your support, which knows no limit, your love, which has always encouraged me, and your continued belief that I could do whatever I set my mind to. To my brother, who is a more selfless and giving individual than I, thank you for always telling me that my work, however selfish I feared it might be, was also important. You are an inspiration and I am continuously proud to call you my big brother. To Jess, thank you for your love, encouragement, and most of all your patience. The past four years have been a struggle for us both, but I do not believe I could have made it through them, or even would have tried, if it had not been for you. To my friend Andrew Capeluto, thank you for all your help. Much of the work presented herein was influenced by and a direct result of our continued collaboration, it's been a great journey Tesla-bro. To my fellow classmates and friends, it has been a real honor working and growing alongside you and I am forever changed from knowing you all. To Randall Packer, with whom a chance encounter and an experimental class introduced us, your direct and keen observations pushed me in the direction of what would become Next Music, and for that I thank you. To Kenneth Cameron, another chance encounter brought us together to work on what would become Cultivating Frequencies, through that project I have been privileged enough to come to call you my friend, but your mentorship and insight brought my work to a new level, and for that I am so very grateful. To my teachers Owen Vallis and Jordan Hochenbaum, you both have helped me on so many occasions I could not even recount them all. Always patient and insightful, your intelligence and talent has touched all aspects of my work in the past three years that I've known you. Thanks for everything. Lastly, to my mentor Ajay Kapur, without you I would not be writing any of this. Thank you for convincing me to come to CalArts and pursue a BFA when what I was looking for was a graduate program. And thank you for creating a MFA program so I could achieve what I had set out to do in the first place. Your guidance and your mentorship have been transformative. You told me, that if I came to CalArts, that I would be a "rock star" when I left. Looking back, I thought you were

just telling me what I wanted to hear, but I was convinced and I have not regretted the decision to come back to school since. You gave me the opportunity to prove my worth to myself and to the artistic community at large, and in doing so I have emerged a better person and a stronger artist. So for all you have done and continue to do for me, thank you.

Contents

Abstract i

Acknowledgments iii

Contents v

List of Figures vii

List of Tables viii

Chapter 1 Introduction 1

1.1 What is Next Music? 1

1.2 Next Music and Next Nature 2

1.3 Sound, Space, and Music 4

1.4 Technology and Music 5

1.5 Computing, Sensing, and Music 6

1.6 Thesis Outline 8

Chapter 2 Cultivating Frequencies: The Interactive Garden 11

2.1 Background 13

2.2 System Overview 15

2.3 Designing a Musical System 17

2.3.1 Motivation 19

2.4 Generative Musical System: A Detailed Explanation 21

2.4.1 Input a Dataset, Output a Song 21

2.4.2 Data to Sound 22

2.4.3 Time and Tonality 23

2.4.4 Weather Station 24

2.4.5 Soil Moisture and Probability 25

2.4.6 Temperature and Tempo 25

2.4.7 Humidity and Wetness 26

2.4.8 The Importance of PH 27

2.4.9 Touch 27

2.5 Capacitive Sensing 28

2.5.1 First, A Note on the Importance of Touch ... 29
2.5.2 Capacitive Sensing: Methods Old and New ... 30
2.5.3 Evaluation: Methods Old and New ... 33
2.6 Future Work ... 36
2.7 Closing Thoughts ... 37

Chapter 3 The Third Room and the Augmented Musical Space ... 39
3.1 Background ... 41
3.2 System Overview ... 43
3.3 Natural Interaction and Performance ... 45
3.4 Accounting for the Physical in Virtual Performance ... 49
3.4.1 Ember: Digital Mbira Interface ... 49
3.4.2 The Fluid: Wireless Sonic Hand Sculpture ... 53
3.5 Extending virtual interaction with mechatronic sound-objects ... 56
3.6 Future Work ... 57
3.7 Closing Thoughts ... 58

Chapter 4 Hear In LA: Augmenting the Acoustic Landscape ... 61
4.1 Overview ... 62
4.2 Background and Motivations ... 64
4.3 Field Recordings versus Compositions ... 66
4.4 Crowd Sourced Collaboration and Composition ... 68
4.5 Implementation Strategies ... 69
4.5.1 JUCE and Developing a Cross Platform Application ... 70
4.5.2 Interfacing with the SoundCloud API ... 71
4.6 Future Work and Expansion ... 71
4.7 Closing Thoughts ... 72

Chapter 5 Conclusion ... 75
5.1 Summary ... 75
5.2 Primary Contributions ... 76
5.3 Final Thoughts ... 77

Bibliography ... 79

List of Figures

Figure 1: Next Music Conceptual Diagram 3
Figure 2: Next Music defined as music for the augmented space 4
Figure 3: Cultivating Frequencies Garden - Aiguablava, Spain 11
Figure 4: Cultivating Frequencies System Diagram 16
Figure 5: Temperature in Barcelona over 24 hours (left) compared to audio waveform magnified (right) 18
Figure 6: Generative Musical System Flow Chart 22
Figure 7: Tonic and Month Diagram 24
Figure 8: CapSense Circuit Diagram (* Resistor Value of 100K-50M) 30
Figure 9: SFCS Circuit for Arduino Mega - Send pins change depending on Arduino model and Timer used 31
Figure 10: SweepingCapSense Test - Electrode - All Gestures - Position 1 35
Figure 11: SweepingCapSense Test - Plant - All Gestures - Position 1 35
Figure 12: SweepingCapSense Test - Plant - Two-Finger Pinch - All Positions 35
Figure 13: Third Room on display at the CalArts Digital Arts Expo 2013 (photo credit: Scott Groller) 39
Figure 14: Third Room System Overview Diagram 44
Figure 15: Third Room Virtual Rendering in Cinder and OpenGL 46
Figure 16: Traditional Mbira (left) and First Ember Protoype (right) 50
Figure 17: Second Ember Prototype 51
Figure 18: Ember final version 53
Figure 19: Fluid - Version 2, with copper plated touch points 54
Figure 20: Custom Arduino Fio PCB Shield for Fluid 55
Figure 21: Modulets - a mechatronic grid sound object 57
Figure 22: Hear In LA - Application Flow Chart 63
Figure 23: Geographic Scale of each project 76

List of Tables

Table 1: Time and Scale relationship 23
Table 2: CapSense Test Results 34

Chapter 1
Introduction

As technology and culture continue to merge a new landscape is emerging. The combined forces of networked and mobile devices, advanced sensing technologies, and a continued push towards ubiquitous computing allow humans to communicate, not just with each other, but with their environments as well. These manufactured environments, where a physical space is overlaid with technologies that provide "dynamic and rich multimedia information," are what new media theorist Lev Manovich calls *augmented space* (Manovich 2010). The emergence of such a space is an exciting occurrence. While there are some caveats (surveillance, monitoring, privacy, etc.) that stem from the same technologies that help to create the augmented space, it still remains a new space to be explored. So far, many of the applications of this space are largely commercial though some applications verge on the edge of ubiquitous computing. Manovich sites the shopping and entertainment centers of Seoul, Tokyo, and Hong Kong, which are completely covered with dynamic displays, as commercial examples of the augmented space. However, the other side of the spectrum is "*any human-constructed space* where subjects can access various information wirelessly on their cell phones . . . or laptops" (Manovich 2010). Furthering Manovich's idea, the ubiquity of wireless connections and the increasing sophistication of our mobile devices have our society on the path to complete augmentation. However, if access to digital information becomes truly ubiquitous, how might this affect the augmented space? If everywhere is augmented, what makes this space unique? Next Music is not an answer to these questions; however, it does present a unique approach to Manovich's idea.

1.1 What is Next Music?

Next Music, put simply, is music created inside the augmented space. Read "inside" as for, with, about, etc. This means, that any kind of music that intentionally uses the augmented space for its creation, performance, and/or listening, falls within the realm of Next Music. Next Music is created for and with the augmented space, and it is this intention that sets it apart from other music that is meant to be distributed through standard channels, some of which may fall within Manovich's definition of the augmented space, such as internet radio and other streaming services. For the definition of Next Music, music shall be understood both in the traditional sense of melody and song, but also in the sense of deliberately curated sounds, including, but not limited to, field recordings, ambient sounds, and noise. Coining the term Next Music is an attempt to categorize an emerging practice that is signified by its unique manipulation of space through technology, and its use as a key element within the many practices of music.

1.2 Next Music and Next Nature

The idea for Next Music is closely related to, and inspired by, the ideas of theorists working in the field of Next Nature. Next Nature is a school of thought that attempts to re-contextualize beliefs about nature and technology, and specifically how culture plays a large part in defining this context. Humans evolved out of a natural world and our subsequent success as a species has altered that world forever. Nature is no longer the place humans find themselves in, rather it is more often thought of as something one goes to visit or see in a designated "natural" setting. In the extreme, nature is the unstoppable force that can cause disaster and devastation, despite all human efforts to prevent this. As a species, humans have conceptually and culturally removed themselves from nature to create civilizations, and so, nature becomes viewed as something "other" and outside of human civilization. The primary method of this removal is the creation and development of technology. From stone tools, controlled fire, and agriculture, to antibiotics, nuclear weapons, and the internet, human activity has primarily revolved around the ability to create new technologies which have, directly and indirectly, allowed humans to withstand all but the greatest of natural threats (Mensvoort and Grievink 2012). It is because of this technology, and the active control humans enact upon nature through technology that has transformed nature into a "cultural category" and "a simulation" (Mensvoort and Grievink 2012). Simultaneously, "our technological environment becomes so complex and uncontrollable, that we start to relate to it as a nature of its own" (Mensvoort and Grievink 2012). Within the ideas of Next Nature these are facts that are accepted without positive or negative connotation. Instead it is more important to guide people to this new understanding so that, as a species, humans can move forward with the knowledge that their activities and technologies will play a substantial role in the shaping of the future. According to the author Bruce Sterling, this means that the human race is the new catalyst of evolution and "with our urge to design our environment, we cause the rising of a next nature that is wild and unpredictable as ever."[1]

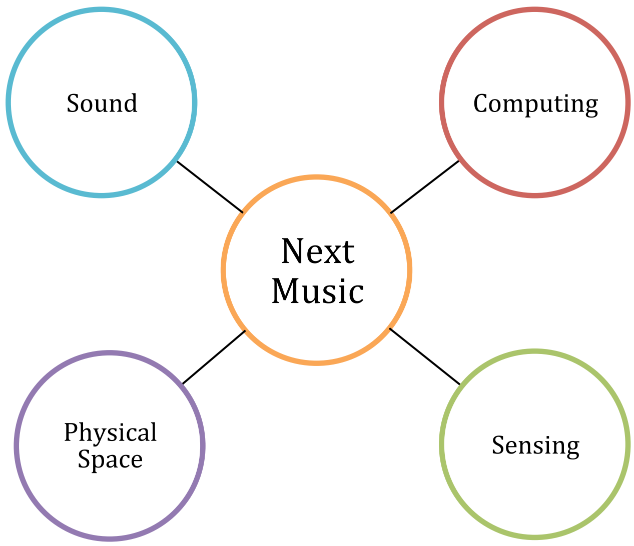

Figure 1: Next Music Conceptual Diagram

It is this notion of evolution through activity, design, and technology that first inspired the idea of Next Music. In trying to connect the following works with a larger theme, it was obvious that there were definite similarities between them, but how they all came together to form a larger idea was, at first, unclear. The common elements of these works broke down into the following categories: Sound, Computing, Sensing, and Physical Space (Figure 1). It was not until the ideas of computing, sensing, and physical space, conceptually converged into a singular idea, that of the augmented space, that Next Music clearly emerged as the sound and music of the augmented space (Figure 2). However, each of the ideas that merged to form the augmented space are still important ideas that lead to this new practice, and have their own histories and implications within the realm of music.

[1] http://www.nextnature.net/themes/nextnature/

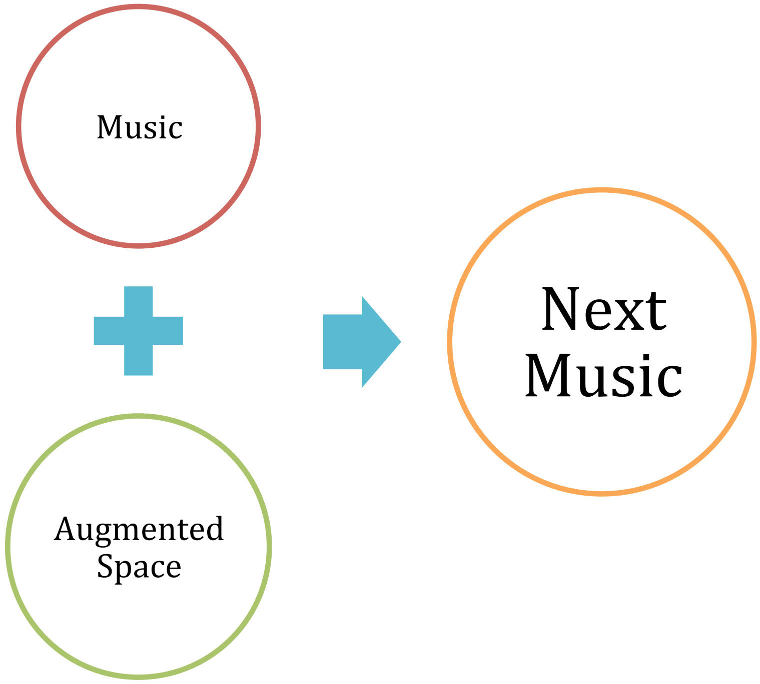

Figure 2: Next Music defined as music for the augmented space

1.3 Sound, Space, and Music

Music often exists within its own space. This is a physical space that is created, or perhaps altered, by the mere presence of music, whether it is being performed, played over loud speakers, or listened to on headphones. The power of music is inherently transformative, as all music is created through some transformation or another. Breath is transformed into song; a tightened skin over a hollow log is transformed into rhythm; a vibrating string is transformed into a note; a computer bit is transformed into physical sound. The acoustic space itself can also be transformative. A lone flute played in a large field can be perceived as fundamentally different than that same flute played inside a cathedral. The musician David Byrne, in his book How Music Works, claims that, historically, the acoustic space where music was written and performed is directly linked with the kind of music that was created (Byrne 2012). According to Byrne, this explains why African music is largely comprised of loud and rhythmic drumming and medieval western music was primarily slow moving and modal. The African drums were played outside where the acoustic space readily absorbed the sound and reverberation was not an issue. Whereas, inside a gothic cathedral, the reverberation lasts several seconds and each note lingers, altering the soundscape long after its initial ending. While mostly based on observational evidence, Byrne's theory is a compelling argument for how and why musical styles differ from culture to culture, and within a given culture, from space to space. Today, it could be argued that one of the most common musical spaces is the personal space provided by headphones. Within the artificial space provided by modern headphones the effect of the physical acoustic space surrounding the listener is made obsolete. This allows for the recreation of the ideal listening space within almost any environment. In the scope of Next Music, physical space is a requirement in the creation of augmented space, and therefore plays an incredibly important role. It is the physical acoustic space that is augmented by way of technology that allows for the creation of new musical experiences driven by the user's movements and activities within the augmented space.

1.4 Technology and Music

Headphones are a small example of how modern technology has altered the relationship between sound, space, and music. However, this trend is not a new development. Music, as it is understood today, is a product of technological developments throughout human history. As stated above, even the engineering and architectural achievements of the great European cathedrals, products of technology themselves (Gies and Gies 1994), can be credited with influencing the evolution of western music. Even before then, the math used to understand simple relationships between harmonic tones (Barker 2004) helped to define the future of tonal western music. Just as the metallurgical methods of evolving a single conical horn into a valve trumpet created a new world of possibilities for instrument makers, musicians, and composers alike (Holmes 2013). Technology, defined as the practical application of knowledge especially in a particular area (Merriam-Webster 2005), is not, however, limited to math and engineering. Even the symbolic codification of western music in the ninth century is a technology that has influenced music for hundreds of years (Taylor 2001).

With the introduction of electronics into mainstream life and culture came a paradigm shift in music. As technology, electronics, and music converged, new sounds, instruments, practices, and spaces have emerged. Many of these new musical artifacts challenged the popularly held notions of music at the time, but practically none of them can be said to have remained completely removed from the evolving and inclusive definition of music. For centuries, developments in music technology have pertained mostly to the creation of new instruments, tuning systems, and even methodologies for composition and performance. Even the rise of electronics is marked by these same trends. However, there were many new and unprecedented developments that would change music forever. Most markedly, is the development of audio recording technology, which changed, not only the act of listening to music, but its composition and eventually its performance. These developments started gradually and then began to increase in speed as the science behind audio recording technology developed into its own field. Many of the effects recording technology has had on music were not immediately foreseeable but grew out of the continued practice of recording. Use of a technology, in itself, seems to spur the evolution and improvement of that technology and experimentation leads to the development of new techniques, especially as the idea takes on its own life as part of a cultural consciousness.

Technology has also transformed the spaces music occupies on many different levels. With recorded music people could listen to their favorite artists at the leisure. With radio, these songs could be broadcast over vast distances to millions of people simultaneously. Today, people have the ability to curate and consume their own mass collection of music practically everywhere they go. When designing musical spaces, acousticians are often consulted to help design a space that can enhance the sound, in the case of performance venues, or deaden the sound in the case of recording studios. From a musician's point of view, modern computing allows them to record their own songs almost anywhere, edit those recordings, and even make them sound as if they were played and recorded in famous venues around the world,[2] or perhaps a space of their own design.[3] Truly, modern music hardware and software continues to change the way in which music is written, produced, performed, and consumed.

1.5 Computing, Sensing, and Music

In the recent past the push towards ubiquitous computing has changed the way a large portion of the world's population interacts and continues to change the dynamics of social movement and thought. More than ever, people are earning a livelihood using a computer and simultaneously many are delegating their social interactions to their personal computers as well. The device at the forefront of this movement is the smartphone, which has put a highly advanced computer in the pocket of over a billion people throughout the world ("Smartphone

---

[2] http://www.audioease.com/Pages/Altiverb/
[3] http://www.uaudio.com/store/reverbs/realverb-pro.html

Users Around the World - Statistics and Facts [Infographic]" 2013). This device, which began as a cellular phone, has become a mobile beacon of data transmission and information retrieval. The accompanying evolution of design and use has rendered a "phone" which is more likely to recommend a place for lunch, provide directions to its location, take a photo of that lunch, and upload that photo to a social media platform, than it is to make an actual phone call. Still, it is marketed as a phone, so that is how it is perceived, even though it is truly a mobile sensing and computing device. However, despite the obvious ubiquity of computers and mobile computing devices, the ideal of "ubiquitous computing" has not yet been realized. The man who coined the term, Mark Weiser, believed "that only when [computers] disappear . . . are we freed to use them without thinking and so to focus beyond them on new goals" (Weiser 1991). In his mind, only when computing reaches the level of ubiquity that written language has, will his dream of ubiquitous computing be achieved and the full potential of computing unlocked.

As computing reaches true ubiquity, our perception of the computer, as well as the computer itself, will likely evolve as it ceases to be a static collection of hardware and becomes an extension of human life and even the human body. While still several, if not many, years away, the related concept of ubiquitous sensing is helping to pave the way. The current trend is that sensors and sensing technology is getting better, cheaper, and smaller, allowing large companies and individuals alike to create some truly remarkable things. At the International Conference on New Interfaces for Musical Expression (NIME), they are dedicated to researching and developing new methods and techniques for working with sensors, specifically towards the creation of new musical interfaces and instruments, as well as strategies to help make effective musical use of these new instruments and technologies. Everything from augmented and hyper-instruments (Overholt 2005; Heinrichs and McPherson 2012; Kapur et al. 2003), to studies on new sensor techniques (Tindale et al. 2005; J. Murphy, Kapur, and Burgin 2010; McPherson 2012; Wiley and Kapur 2009), to mechatronic instruments (Kapur, Hochenbaum, et al. 2011; D. J. Murphy et al. 2013), and even working with bio-data and sensors to create musically expressive instruments (Angel 2011; Dubost and Tanaka 2002; Nagashima 2003). The work done by this growing community, which is documented and shared freely, promotes the growth and further development of old and new technologies alike towards the goal of musical expression.

Sensors and computing are at the heart of both the augmented space and Next Music. It is the combination of these elements with a physical space that allows for this new space and musical practice to emerge. If and when ubiquitous computing is achieved as Weiser envisioned it, most advanced industrial societies will be filled with augmented spaces. This will no doubt affect many aspects of daily life for better or for worse. Next Music is not, by definition, a commentary on these effects, it is merely an attempt to envision how musical practices will change and explore what the creative possibilities are when working within this hybridized environment. Just as early western composers explored the reverberations of the great cathedrals by experimenting with harmonies and counter-melodies, Next Music will explore the augmented space. The biggest difference being that the augmented space can be altered and changed dynamically. As discussed below, the method of augmentation depends on the space itself, and both the space and the means of augmentation contribute to the style of interaction and the kind of music produced. These projects are by no means comprehensive as representations of what Next Music can be. Instead they offer a small glimpse at a few different kinds of spaces that occur on different geographic scales.

1.6 Thesis Outline

The following projects are the very beginnings of what could become a much larger practice. Each project represents a different kind of space and features different methods for augmenting them. Furthermore, each project offers a unique approach to musical creation and interaction that stem from the context of the space itself. In Chapter 2 (pg. 11), Cultivating Frequencies presents an augmented garden that not only senses the conditions of its environment but also can sense the presence and touch of an observer, culminating in an interactive and generative musical experience. This project also presents a new Arduino library, *SweepingCapSense,* which helps create more robust touch sensors out of conductive materials, even biological materials, like plants. In Chapter 3 (pg. 39), the Third Room imagines an augmented performance space that features both virtual and physical interfaces as well as a new style of mechatronic sound-object that is both modular and highly spatial. The actual project is more of a case study to explore how one might augment a more traditional performance space or even a studio/compositional space, where hardware and software, physical and virtual, elements combine to extend the creative control and expression of the performer and/or composer. Chapter 4 presents Hear In LA (pg. 61), a location-based interactive sound installation that uses crowd sourcing to create a citywide soundscape based on field recordings created and uploaded by the users. Aimed at reconnecting the commuting population of Los Angeles with the cultural landscape they often overlook, this project augments space on a global level by utilizing the location sensors built into Android and iOS smartphones and mobile devices. Lastly, the future of Next Music will be discussed in the context of both the author's own work and within the music and art worlds as a whole.

Chapter 2
Cultivating Frequencies: The Interactive Garden

Figure 3: Cultivating Frequencies Garden - Aiguablava, Spain

At its core, Cultivating Frequencies transforms a garden into an interactive sound installation. To be more explicit, Cultivating Frequencies is an interactive biotechnical sonic sculpture that challenges the duality of nature and technology by exploring the artificiality in nature and the natural in technology through the medium of an interactive and musical garden. Nature has largely become governed by culture. Not completely under human control, but much of what was once purely natural has become a product of human activity and culture. Genetically modified organisms are an extreme example of the kind of control human culture can have over the natural world, but even the act of landscaping and gardening brings nature into the realm of culture. At the same time, culture, which can be defined as the cumulative products of human civilization, can now create things so intricate and complex that they often behave like organisms themselves.[4] An economy, for instance, is a man-made entity that is so complex that humans cannot control it, despite their desire to do so. Next Nature is a movement to acknowledge these emerging truths as a means of creating a new foundation for understanding both culture and nature. Technology plays a large part in both culture and nature, influencing both and in turn being influenced by them in a complicated relationship that is anything but straightforward. Technology, however, specifically descends from studies of nature, and in many ways exploits the laws and material properties of nature to create useful and amazing tools. The study of electricity lead to the electronic and then digital revolutions which both lead into the information age. No matter how far removed technology now seems from nature, it will always have its roots in natural phenomenon and innovators will continue to look at nature for inspiration.

Cultivating Frequencies plays on these ideas by taking a culturally designed slice of nature, a garden, and infusing it with sensor and computing technology to create an interactive musical experience that embodies the relationship between nature, humans, and technology. By removing plants from nature and installing them in a garden humans sever the plant's innate ability to thrive. This removal creates a dependent relationship whereby the plants need human attention to survive. Using a system of environmental sensors, data is collected on elements of plant health and growth, allowing the garden to have its own memory. This memory is used as the basis for a generative composition that the garden can "perform" when in the presence of viewers. The music performed is designed to communicate the data collected by the various sensors and give the listener an indication of the plant's overall health. What's more, if the observer approaches the garden and touches the plants, they will find that the plants have been

[4] http://www.nextnature.net/about/

transformed into touch sensors and can be used to play along with the generative piece. Furthering the notion of nature's connection with technology, it is by the laws of nature that this feat is even possible to begin with. Due to the moisture and water present within the soil, the roots, and the plants themselves, a touch capacitance circuit embedded in the plant's soil extends through the plant, like a natural electronic component (discussed in detail below in section 2.5). The electronics are a part of the sculptural structure, adding a level of ubiquity that creates a unique experience for the observer.

In the context of Next Music, the augmented space of Cultivating Frequencies occupies the smallest and most vulnerable of spaces. The garden space is augmented through a collection of sensors and computing to create a sense of awareness that allows a simple form of communication between two completely different kinds of organisms. While this communication is inherently abstracted and conceptual in form and presentation, it nonetheless attempts to allow the transmission of information between a plant and a human being. By augmenting the physical space of a garden with an interactive musical system derived from the data collected about the elements most significant to the health and status of that garden, this piece creates a new kind of biologically driven and influenced musical space. In this musical space, the intimate interaction that occurs between plant and observer is transformed into a musical one. The music itself is influenced both by the plants' health as well as the observer's physical interaction with the plants. The symbiotic relationship is simultaneously exploited and encouraged through an interactive musical medium.

In this chapter the project Cultivating Frequencies will be presented and discussed at length. First, a brief background of Cybernetic art will be presented (pg. 13). Second, a system overview will briefly list the complex collection of technology and musical concepts used to create the piece (pg. 15). Third, the motivations behind the musical system will be listed and discussed (pg. 19). Fourth, the generative musical system is explained in depth (pg. 21). And fifth, a discussion on capacitive sensing, showcasing its role within the piece, and a new open-source method and library will be presented (pg. 28).

2.1 Background

In a larger artistic context, Cultivating Frequencies embodies a lot of what Roy Ascott calls the cybernetic vision. Cybernetics was popularized by Norbert Wiener in the mid 20th

century and was defined by him as "the science of control and communication, in the animal and the machine" (Ashby 1957). In short, cybernetics is interested in the behavior of things. Ascott saw the potential for artistic exploration of the cybernetic ideas on systems, feedback and information (Shanken, Clarke, and Henderson 2002). So much so, that he viewed it as an opportunity for a new a kind of art that was fundamentally different from art of the past. He wrote:

> *"The dominant feature of art of the past was the wish to transmit a clearly defined message to the spectator as a more or less passive receptor . . . Modern art, by contrast is concerned to initiate events and with the forming of concepts of existence. The vision of art has shifted from the field of objects to the field of behavior and its function has become less descriptive and more purposive." (Ascott and Shanken 2003)*

Ascott's "Cybernetic Vision" focused on behaviorist art, which is inherently interactive and seeks to break the passive role of the observer, transforming them into participants. In his view, art has historically been a matter of one person, the artist, strictly controlling their medium as a means of communicating a message to the audience, whereas modern art now allows for "a two-way exchange between the artwork and the audience" (Ascott and Shanken 2003). The modern computer has made this kind of interactive art more accessible than ever, allowing for audiences to view and participate with artworks in ways previously unimaginable. However, these ideas occurred years before the personal computer even existed (Jordan and Packer 2002); even so, artists were inspired to create artworks based on the ideas of Wiener, Ascott, and others.

One of the earliest example of a cybernetic artwork is Nicolas Schöffer's "CYSP-1," which was a robotic sculpture "that was not only active and re-active . . . but also autonomous and pro-active" (Brown 2008). Powered by an "electronic brain," the mobile sculpture housed light sensors and a microphone, sensing and reacting to the world around it, but ultimately seeking equilibrium, as it was essentially a "homeostat on wheels."[5] The homeostat, originally conceived of and designed by William Ross Ashby, is an ultra-stable system "that does nothing more than run to a state of equilibrium" (Ashby 1957). His experiments would show that this dash towards equilibrium is achieved through complex behavior, but the end goal is always stability (Franchi and Bianchini 2011). Schöffer's CYSP-1 was also able to achieve complex

[5] http://cyberneticzoo.com/cyberneticanimals/1956-cysp-1-nicolas-schoffer-hungarianfrench/

behaviors, so much so that he considered it an actual "being" saying "for the first time . . . it [is] possible to replace man with a work of abstract art, acting on its own initiative," which he proved by allowing it to perform alongside various dance troupes (Brown 2008).

Cybernetics influenced many artists through the second half of the twentieth century, both directly and indirectly. As the ideas of information, feedback, and systems became part of the cultural consciousness and everyday life, more and more artists would strive to create works that represented, questioned, and used these ideas as the conceptual foundations of their art. In the development of Cultivating Frequencies these ideas were paramount as the entire project was centered on these very principles. In the following sections, the system of Cultivating Frequencies will be explained in detail, including the motivating and reasoning behind many of the decisions made during the process of development.

2.2 System Overview

Cultivating Frequencies is a sonic sculpture and bio-technical interface, meaning it is both musical and interactive. It integrates a mini weather station, a Mac Mini running Max/Msp, four ultrasonic rangefinders, and an Arduino with a custom PCB shield. This can be thought of as integrating data collection, a generative musical system, user detection, and touch sensing plants, respectively.

At all times the weather station is keeping track of the Temperature, Relative Humidity, and Soil Moisture, which is both logged into its own memory and is stored on the Mac Mini as well. The Mac Mini communicates with the weather station over serial protocol from within Max/Msp by waking up the weather station and requesting the current conditions. These current conditions are returned as a series of characters and are input into the Max/Msp program where they are stored and then uploaded to an online data hosting service called ThingSpeak.com. This data is used to set parameters within the generative musical system that controls things like rate/tempo, reverb, and movement. However, the musical system does not output sound until a user is detected.

Using the weather resistant ultrasonic rangefinders that point outwards from the installation's frame, it is possible to detect when a person is approaching, walking past, around, or standing in front of the installation. When a person is detected, the music fades in and outputs a generative composition based on the current conditions. The composition will continue to play as long as a person is still detected and will fade out shortly after the person leaves. However, if the person chooses to approach the garden and interact with it, another level of the composition unfolds.

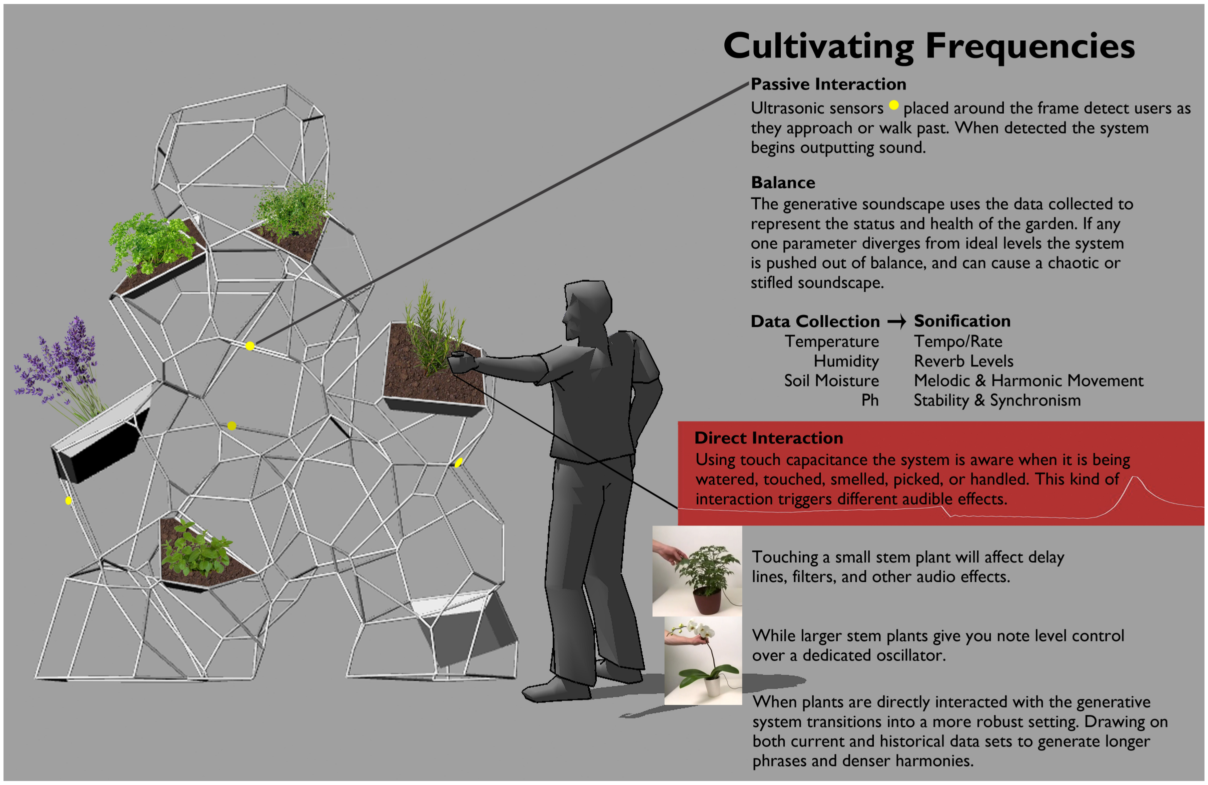

Figure 4: Cultivating Frequencies System Diagram

Each plant has been turned into a touch sensor by a new technique called Swept Frequency Capacitive Sensing. Essentially, this technique adds a second dimension to the data received from a touch sensor. Normally, a single value is returned which can show proximity and amount of touch, but is less accurate in these regards and is more reliably interpreted as a binary value, displaying a touch versus no touch. This new technique returns a whole dataset that can be interpreted more precisely as different gestures when paired with machine learning techniques. The details of this new style of capacitive sensing are presented and compared with the older method below in section 2.5. When a person (technically any conductive object) touches a plant in the system, there is a musical response. Depending on the plant, this musical response can vary from note-level control, where one can play notes by touching the plant in different places, or effect-level control, whereby touching changes values in delay lines or reverb settings.

Touching and interacting with the plants is encouraged and rewarded by a secondary generative composition when enough time has been spent engaging the plants. This secondary composition is different from the primary composition in two important ways. One, the primary composition uses an event and pulse-based triggering system, where a single event/pulse occurs and then cascades down through the system. In the secondary composition, this pulse system is replaced by a continuous flow of sounds and events that crest and trough like a wave, giving way to long notes and phrases. And two, the primary composition uses only the current condition data as parameters for the sound design, but the secondary composition pulls historic data as well, looking back to different times and using the data to create counterpoint melodies and harmonies based on the garden's memories. To discourage users from over-interacting with the plants there is a point where the interaction triggers the end of the secondary composition and returns to the primary. Or, if the user stops interacting with the plants it will also trigger a return to the primary composition. Then, when the user walks away, the system fades out to silence. In the sections below (2.3 and 2.4) the conceptual motivations and implementation of the generative musical system is explained in depth.

2.3 Designing a Musical System

"The complexity we see around us is not the result of complex initial conditions, but the result of interaction of simple initial conditions." Brian Eno

Cultivating Frequencies, the title, was a point of departure and conceptual foundation for the design of the generative musical system. When dealing with data, the tendency is to graph the data to make it easier to understand and define patterns. The timeframe of the data inevitably changes the nature of those patterns. Data collection in this project deals with the health of the plants (PH, Soil Moisture) and the environmental conditions (Temperature, Humidity). Graphing this data over varying lengths of time creates many shapes that resemble audio waveforms, as can be seen below (Figure 5) where a graph of the hourly temperature in Barcelona in May 2013[6] is juxtaposed with a sample of visually magnified audio. This observation lead to the idea that when data is collected the dataset is really just the waveform

[6] Data taken from http://www.eurometeo.com

representation of that collected data, and, when approached as a waveform, this data can be said to have a certain frequency.

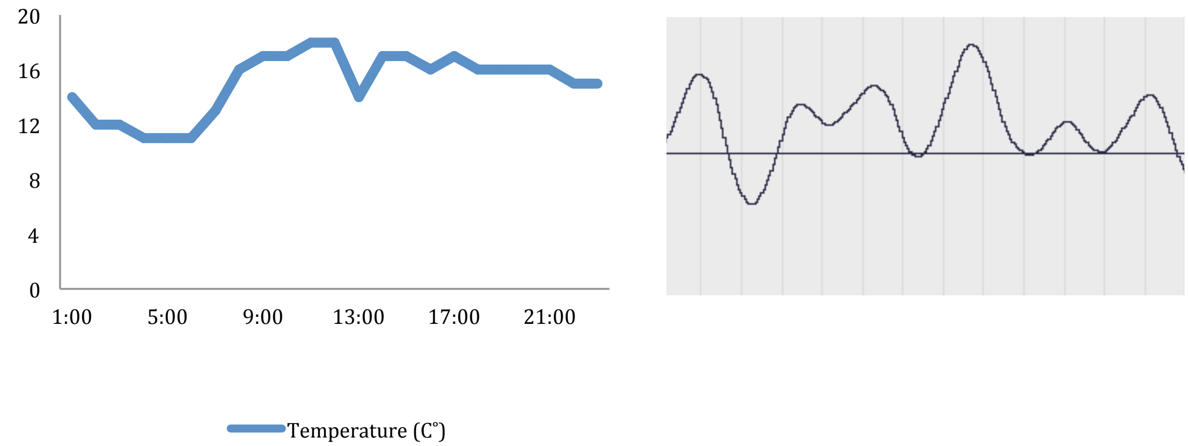

Figure 5: Temperature in Barcelona over 24 hours (left) compared to audio waveform magnified (right)

Frequency, being defined as the rate that something occurs, is not conceptually limited to cycles per second here, as in the case of sound waves. In this context, frequency exists on many time scales simultaneously. Beyond the data collection, it could be said that the garden itself has its own frequency. The life cycle of the plants; the seasonal cycle of growth and dormancy; the cycle of watering and of drying; the cycle of tending and harvesting. All of these cycles can be said to have frequency. At the same time, the act of collecting data at a predetermined rate, conceptually relates more to an impulse than a waveform with a frequency. This impulse is the moment of collection that contributes to the waveform as a collected whole, but in itself, is singular. This idea furthers the conceptual comparison of data collection and audio waveforms because analog to digital audio conversion occurs by taking discrete samples at a rate so fast that they seem continuous.

A duality occurs, between nature, the undulating continuous frequency, and technology, with its discrete precision and binary logic. It is these two insights, continuous frequencies versus discrete impulses, these two states of being gleaned from considering data as a medium for artistic inspiration, that best define the musical system designed for Cultivating Frequencies. It is important to keep these ideas in mind as the other major elements and finer details of the generative musical system are described below.

2.3.1 Motivation

Each aspect of the musical system was thought about and approached very deliberately. The motivations were both conceptual and practical as the piece needed to meet certain design requirements and possess certain qualities, both technically and artistically. The decisions made had to be thoughtfully addressed because in a system, everything is inherently interconnected, which means even little changes could have drastic effects elsewhere. In this section the motivation and approach taken to develop the compositional system used in Cultivating Frequencies is presented.

2.3.1.1 *Time*

The biggest influence on all compositional and design choices were influenced by the fact that the installation would be permanently installed at a private residence. This meant that the owner will experience the piece over an extended period of time, and therefore the musical system would need to be able to grow over the years, just as the garden would. However, the owner would not be present year-round, so the system could not be too slow in its development and output, and should have noticeable differences in as many timeframes as possible. The idea of temporality is present in each layer developed for Cultivating Frequencies. At any time one can interact with the garden and hear a song. This song is defined by many temporal parameters; each day the system cycles through a set of predetermined scales; each month the tonic of that scale changes; seasonal temperatures will change the overall feel of the music. And as the years go by the memory of the garden increases and the archived data is incorporated into the generative musical system, creating another layer of complexity and sonic possibility.

2.3.1.2 *Homeostasis: Biology and Cybernetics*

Homeostasis is the tendency toward a relatively stable equilibrium between interdependent elements. In botany and biology, homeostasis is a generally desirable state of being. In Cybernetics, homeostasis is a major conceptual concern as it is the study of regulatory systems. The interesting part of homeostasis is that it is always in a state of flux. As a system strives for balance, it is also continuously being affected by a maelstrom of other influences, all of which affect the point of equilibrium as it oscillates above and below what might be considered the ideal level. A healthy ecosystem is made up of a lot moving parts, all of which contribute to the overall stasis. A garden could be described as a false ecosystem living inside a larger true ecosystem. This false ecosystem is tended to directly, largely in reaction to its stasis as well as the conditions found in the larger ecosystem surrounding it. Hydroponic gardening is the practice of removing all other influences and directly administering all the important nutrients, normally found in the soil and environment, by mixing it carefully with water and using "grow" lamps. Even under these strict conditions, equilibrium is not a fixed point, one can find the same sort of oscillation above and below ideal levels, though on a smaller scale, as would be found in an ecosystem. As exemplified by the Ashby's homeostat, the move towards equilibrium creates complex behavior, and in the case of Cultivating Frequencies, this complex behavior is what drives the generative composition. This notion of a system, one made up of important and separate parts, being in constant flux, but striving for equilibrium, is not only represented by the captured data from the system's input, but has been incorporated into the design of the generative musical system itself.

2.3.1.3 *Interactivity*

From its inception this piece was meant to be interactive. Currently there is a large movement of interactivity in the arts, which fits into Ascott's ideas about the Cybernetic Vision where "the spectator [is drawn] into active participation in the act of creation . . . [and given] the opportunity to become involved in creative behavior on all levels of experience" (Ascott 2002). Cultivating Frequencies embraces this idea and there are multiple layers of interaction built into the design, namely user detection, touch sensitivity, and, eventually, remote access. When a user approaches or walks past the garden, ultrasonic range finders detect their presence and fade-in the music. The user can choose to just listen, walk away, or approach the garden. If they approach they will find that each plant is also a touch sensitive sensor. The novel idea of turning the plants into touch sensitive sensors was inspired by the work done by Disney Research on their *Touché* (Sato, Poupyrev, and Harrison 2012) and *Bontanicus Interacticus* (Poupyrev et al. 2012) projects. Details and information about this work and the implementation used here can be found below in the section on capacitive sensing (2.5). In short, each plant has been transformed into a continuous touch sensor. Interaction is encouraged and rewarded by direct musical control as well as an unveiling of an additional compositional layer. However, too much of a good thing can be a bad thing, therefore superfluous or aggressive interaction is discouraged by removing the sensitivity for a period of time. The idea is to promote a healthy garden and protect the plants. When the user walks away, their presence is no longer detected by the range finders and so the music fades out. The other, more remote, forms of interaction include an iPad app for the owner, a website that hosts collected data, and a standalone application for the general public. The iPad app allows for certain master-level control over the installation, including an On/Off control, master volume, info-graphics on the current conditions, and some customizable aspects of the sound design. The website will display all collected data, general information about the project, photos, videos, and recordings of generative compositions (www.cultivatingfrequencies.com). Lastly, there will be a standalone version of the musical system that can be downloaded, which allows anyone to hear what the garden might sound like at any time from anywhere with an internet connection.

2.4 Generative Musical System: A Detailed Explanation

Generative music is a term coined by Brian Eno "to describe music that is ever-different [and] changing, created by a system."[7] The author defines a generative musical system as a combination of stochastic and algorithmic composition that is achieved programmatically. This means that there is both a random element and a set of rules that work together to generate some kind of musical composition. The generative musical system for Cultivating Frequencies was developed using the programming environment Max/Msp that interfaces with all the hardware, collects the data, pushes the data online and transforms it into music. In this section the details of the generative musical system developed for Cultivating Frequencies are covered in depth.

2.4.1 Input a Dataset, Output a Song

In the simplest terms, generative music is like automatic composition. It combines probability, randomness, and algorithms to create music. This differs from a purely random piece, or stochastic composition, because the randomness that is used is sampling from a predetermined set of rules. Brian Eno first popularized the idea of generative music, and gave the simple example of a wind chime as a system where the only influence one has over the music it makes is the initial choice of tones for each chime (*Robert Sandall and Mark Russell Interview Brian Eno about Generative Music in 1996* 2013). The essence of a generative piece is that it is always changing and it is always different. The initial choices made, the routing of the data, and the random probability allows that the same system can create thousands of variations of itself.

[7] http://intermorphic.com/sseyo/koan/generativemusic1/

Which is to say, each time one experiences Cultivating Frequencies they will hear something that sounds familiar but is almost completely new.

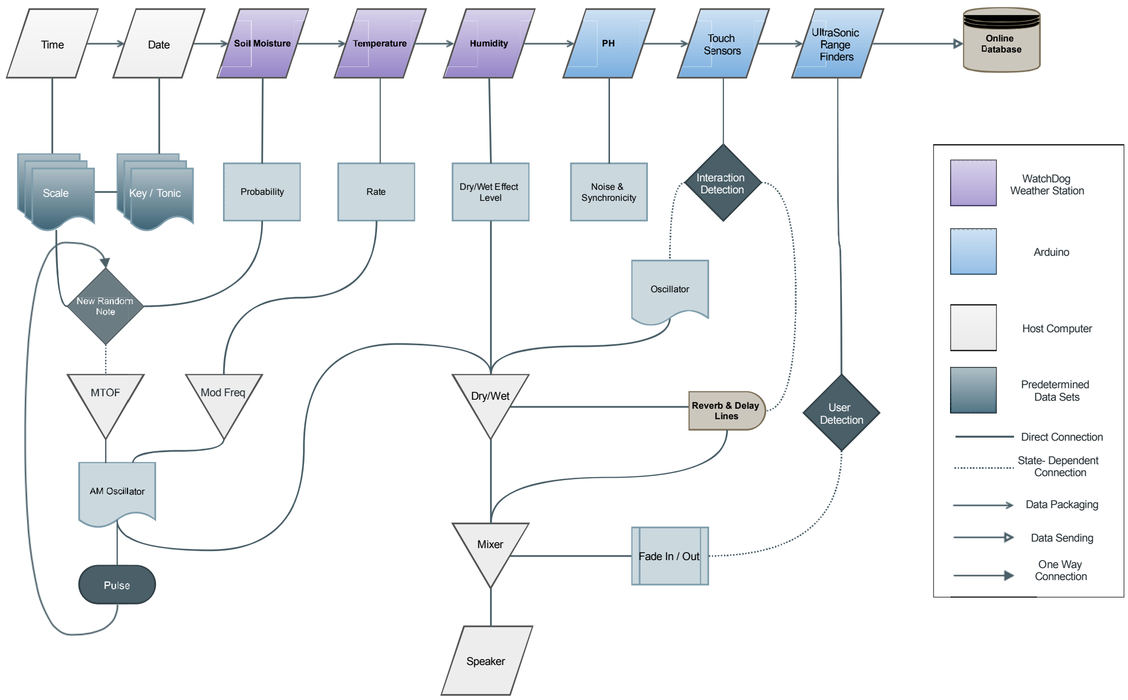

Figure 6: Generative Musical System Flow Chart

2.4.2 Data to Sound

Above, in Figure 6, the majority of the Cultivating Frequencies generative musical system is laid out in the simplest terms possible. While this figure does not capture everything happening in the system, it covers all the important parts. At the top is a list of input data that drives the rest of the system. Each piece of data is color coded to represent where the data is coming from. By following the vertical lines, one can see where each piece of data is used within the musical system, and how some of the data works together to determine the outcomes. The horizontal lines connecting the data represent the periodic collection and storage of the data on the host computer and online. The dotted lines denote a state-dependent connection, which means that the connection only occurs when a predefined requirement has been met. As the system is discussed in detail please refer to Figure 6 as a point of reference and a means of understanding the flow of data as it is transformed into sound.

2.4.3 Time and Tonality

The first piece of data that is collected is the time and the date, also referred to as the timestamp. The timestamp is important because it makes macro level choices concerning the tonality of the whole composition. The time is used to select a pitch collection, or scale, which determines the notes that can be generated later in the system. Pitch collection is used here as a slightly more accurate description than scale because the collection has been slightly weighted, giving certain pitches a higher probability than others. However, the collection of pitches is derived from a certain scale, so the two terms will be used interchangeably. The idea of tying scale to time has a long history in Indian classical music, where the raga drives the improvisation driven melodies (Bagchee 1998). A raga is kind of like a scale as known in Western music, however, with each raga there are certain rules on how the notes are approached and used within a melody. Furthermore, ragas have an explicit relationship to a time of day or season. The implication is that certain sounds more accurately represent different times of day and year. This idea, of connecting time and tonality, was appropriated into the generative system of Cultivating Frequencies.

Time	00:00 - 06:00	06:00 - 10:00	10:00 - 14:00	14:00 - 18:00	18:00 - 0:00
Scale	Minor	Mixolydian	Major	Major Pentatonic	Minor Pentatonic

Table 1: Time and Scale relationship

Together the time and the date set the key and the tonality, however, each changes at a different rate. Every day, the tonality cycles through the pitch collections, essentially splitting the day into five sections: Morning, Midday, Afternoon, Evening, and Night. Each segment of the day is assigned its own scale as shown above in Table 1. The decision to have the pitch collections change throughout the day relates to the discussion in the motivations section about time. By cycling through these pitch collections over 24 hours a user can experience a wide range of what the musical system has to offer in a single day. The collections are ordered so that the movement from one to the next is smooth and will not disrupt the quality of the music when they change.

The date, or rather the month, controls the tonic, also known as the root or the key, of the pitch collection. This decision was made to provide a subtle difference in the music depending on the time of year. While all western scales adhere to a formula of whole and half steps, because of the nature of equal temperament, each key has its own character and timbre qualities. Also, because there are twelve months in the year and twelve notes in the chromatic scale, it seemed only fitting that the key was determined by the month. Below in Figure 7 is a modified circle of fifths diagram that shows the relationship between month and key.

The timestamp, in short, chooses the pitch collection that is used for all melodic and harmonic content. A probability engine driven by the environmental data, collected by the WatchDog Weather Station, controls the actual pitch selection.

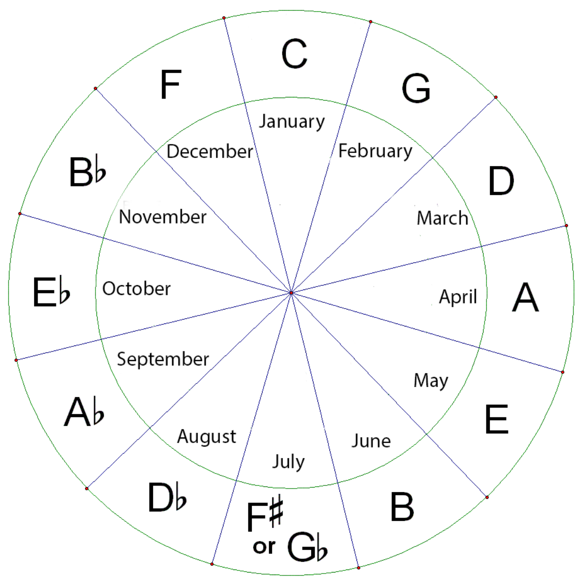

Figure 7: Tonic and Month Diagram

2.4.4 Weather Station

The WatchDog Weather Station is a small, single-purpose computer that collects information on Temperature, Humidity, and Soil Moisture, and logs them on its internal memory. Using a serial connection over USB, the host computer communicates with the weather station and requests the current conditions every couple minutes. The data received by the host computer is an array of characters, symbols, and numbers. The desired information is parsed from the message then packaged and pushed online at ThingSpeak.com[8], as well as input into the generative musical system. Each piece of collected data controls a different parameter of the musical system and can have a drastic effect on the compositions that are generated. Soil moisture is the key to the probability engine that drives the entire system, the temperature sets the rate or tempo of the system, and the humidity controls the master levels of reverb and delay.

[8] https://thingspeak.com/channels/8506

2.4.5 Soil Moisture and Probability

The soil moisture of a plant is important to monitor because soil that is too wet can promote bacteria and disease, while soil that is too dry can cause the plants to wither and eventually die. In itself, the soil moisture data cannot be used to determine the actual health of the plants, even though it can easily be the cause of health problems. In the case of a hydroponic system, the nutrients that are normally found in the soil are instead added to the water. This means the water is doubly important to the health of the plants. Still, it is difficult to determine the health of the plants based on the soil moisture metric alone. Therefore, the soil moisture was approached more conceptually than scientifically. The idea is that a plant that has been properly watered is an active plant, active in the sense of growing and blooming. A plant whose soil has been allowed to get too dry can often times enter a period of dormancy. This is part of a defense mechanism that allows plants to survive weather conditions that do not promote growth, such as harsh winters or dry summers. This idea of activity and dormancy was translated to the musical system as the threshold of the probability engine.

A probability engine generates random numbers, checks those numbers against a threshold, and if the random number is within the range of zero and the threshold, it passes through and something happens. In this musical system the soil moisture acts as the threshold. If the randomly generated number falls within this threshold, a pulse is triggered, which then triggers the three main oscillators, or voices, to pick new pitches and move to them. So, when the plants are watered and active, the threshold is high, which means there is a greater chance that the notes will move on the next pulse. When the plants are dry and dormant, the threshold is low, which means there is a lower chance that the notes will move. Sonically, this is easy to discern, as the watered plants will produce a more active piece, while the dry plants will produce a more stagnant piece. The rate at which these notes move is determined by the next piece of data, the temperature.

2.4.6 Temperature and Tempo

The idea of tying the temperature to the tempo was inspired by the sound that crickets and grasshoppers make, which is commonly known as chirping but is technically called stridulation. According to many studies, there is a direct relationship between the air temperature and the speed crickets chirp (Edes 1899). In Cultivating Frequencies this relationship is used to control the tempo of the composition. While tempo is the correct musical term, in this system, it does not act so much as a tempo, which is a little more rigid in its definition, but really acts as the rate. In fact, the temperature is scaled by a factor of .01 and used as the modulating frequency in an Amplitude Modulation (AM) synthesis operator. Scaling the temperature to a number less than one and using it as a modulator puts it in the sub-oscillating range where the waveform produced is not audible by human ears. This "frequency," then, is no longer measured in cycles per second, but seconds per cycle, which in turn means it is no longer defined as a frequency but a period. It is important to note that this period is not the carrier, but the modulator, so in this case it acts like an amplitude envelope. The modulating waveform is a unipolar sawtooth wave that helps transform the carrier wave, a sine wave, into a pulse that slowly fades away after its initial attack. This pulse is then carried through the rest of the musical system where it triggers several more probability driven oscillators that act as counterpoint and harmony to the main voice. The waveform remains a sawtooth until a user interacts with the plants, which triggers a switch to a sine wave (this is the secondary composition discussed above). When the modulating waveform is a sine wave the generated music goes from staccato notes (very short) to legato notes (very long). In essence the garden really begins to sing where it was only chirping before.

2.4.7 Humidity and Wetness

The last piece of data tracked by the weather station is relative humidity. Relative humidity differs from the more commonly known Dew Point, and is defined as "the ratio of the amount of water vapor actually present in the air to the greatest amount possible at the same temperature" (Merriam-Webster 2005). Simply put, relative humidity signifies how wet the air currently is. The idea of dry and wet have manifestations in audio signal processing too, where it denotes how much of a signal is being returned from an audio effect. Even in this context wetness has a strong association with the effects reverb and delay. This is why both of these effects are featured in this musical system over any other. Reverb is applied at the master level, which is to say at the end of the signal chain. It acts as more of a constant presence, signifying clearly how moist the air is at any given moment. Delay, on the other hand, has been applied throughout the system and is being used in several capacities. Each oscillator has its own built in delay line that helps transform the single pulses generated by the slowly moving amplitude modulation, into longer and more interesting phrases. As the delay chains on each oscillator slowly decay, the repeating notes play off of each other creating new variations of rhythm and melody by offsetting the event driving pulses and delay lines themselves.

2.4.8 The Importance of PH

In a hydroponic system the PH level is extremely important to monitor. Much of the plant's nutrients, normally found in the soil, are administered directly through the water. As the water drains, the PH is measured to determine how much nutrients the plants retain. With this metric insight is given to the plant's overall health and growth, and in the musical system, PH is given a special role that reflects this importance. When the PH is within a range of desirable levels, its affect on the system and music is unnoticeable. Similar to how one might not actively think about a plants health when it looks fine. It's only when a plant looks unhealthy that its health is scrutinized more directly. And so, in the musical system, when the PH level falls outside of the healthy range, noise is added to the system, both literally and figuratively. A noise oscillator adds white noise to the signal and, elsewhere, small deviations are added to frequencies so as to create a kind of inharmonic dissonance. This dissonance differs from normal harmonic dissonance and results in pitches that sound out of tune and create "beating" between the notes, which helps to make the PH level easily understood through the music.

2.4.9 Touch

Touch is an intimate gesture, one that brings its own knowledge and understanding of an object that is hard to communicate in any other form. Sometimes touch serves a utilitarian purpose, while other times the sensation itself brings its own rewards. Touch can serve many purposes in the context of a garden. When tending, pruning, and harvesting the gardener is constantly touching the plants, to the benefit of the plants and often times the gardener as well. As a casual observer people often go to touch plants that have certain inviting qualities, like an interesting texture that spikes the curiosity. In Cultivating Frequencies, touching the plants is encouraged and enhanced by transforming each plant into a continuous touch sensor. This touch sensitivity is achieved by way of a touch capacitance sensor, which is a kind of sensor that can detect any conductive material, including a human being. Using a touch capacitive circuit and a simple copper electrode buried in the soil it is very easy to transform each plant into its own touch sensor. Because of the water present in the soil and in the plant, the capacitive circuit extends from the electrode, to the soil, to the roots, and finally through the plant itself. In fact, this project utilizes a new technique called *Swept Frequency Capacitive Sensing*, developed at Disney

Research, that allows for richer touch data to be collected from a wide variety of conductive materials.

2.5 Capacitive Sensing

Touch capacitance is not a new technology nor is it a novelty in the context of new interfaces for musical expression (NIMEs). It has found its way into many devices and has many unique and industrial applications, from measuring micrometer displacements between silicon cantilevers, to Stud Finders, which detect wooden studs through plaster (Larry 1997). However, large-scale use has only found its way to the consumer electronics market in recent years. In the open-source community there have been few advancements in the technology driving it, most of which has come in the form of more commercially available parts such as cheaper touch screens and sensors. Because the basic form of a touch capacitance circuit is quite easy to replicate there is no need for industrial grade touch circuitry for most of the "do-it-yourself" (DIY) and experimental electronic arts community. However, until now, this basic circuit has only allowed somewhat stable detection of proximity and varying degrees of absolute touch.

This kind of sensor can be implemented as a discrete or continuous sensor depending on the circuit's design. This, plus the fact that capacitive sensing can sense conductive objects through other non-conductive materials, makes it an ideal sensor for industrial and consumer electronics alike. These facts also make capacitive sensing ideal in the development of NIMEs, allowing instrument and interface makers the ability to extract rich touch information while being able to hide the sensor from sight. This is important because the development of NIMEs is not only about technology and sensors, but about instrument design, and it can be very hard to negotiate these two objectives when dealing with sensors that have specific spatial and topographic requirements. In the field of Next Music, touch capacitance could play a pivotal role in the reshaping of human to computer interactions, allowing for novel and inventive applications, including the use of biological materials.

In this section touch capacitive sensing is discussed in detail, in the context of the Cultivating Frequencies, the development of new interfaces for musical expression (NIMEs), and Next Music. The technical details of the implementation and the greater implications of new, more advanced, techniques of capacitive touch sensing will also be presented. Expressly, two methods for capacitive sensing are analyzed and compared, and a new Arduino library, *SweepingCapSense*, is presented.

2.5.1 First, A Note on the Importance of Touch

The relationship between touch and sound has changed drastically through the development of electronics, and consequently electronic instruments. Touch has not been a requirement of musical expression since Theremin introduced his eponymous instrument in the year 1920 (Glinsky 2000). Although many electronic instruments and interfaces employ touch as the main technique for sound activation, manipulation, and expression, in many cases, the connection between the sound and the gesture is abstracted, especially when compared to acoustic instruments. For example, the same violin in the hands of a novice and a virtuoso will sound immensely different even when the same simple gesture is performed; with many electronic instruments and interfaces, it is not necessarily the quality of the sound that changes between the novice and the virtuoso, especially when examining a simple gesture. Rather, it is the understanding of the instrument and range of sounds and consequent manipulations, and even the crafting of the sound itself that is learned over extended use and practice. The gesture in this situation has been applied to the sound, arbitrarily in some cases, and in others as a consequence of the components used in electronics. With acoustic instruments, it was the gesture that physically created the sound, and even had a role in the evolution of the instrument's design.

In computing, touch dominates the interaction between human and computer. However, touch in this context takes on a slightly different meaning, and differs completely from how one touches a musical instrument. Only recently is touch steadily becoming less of a requirement in the realm of human computer interaction. As computer vision and 3D sensing technologies become more advanced and more accurate, the fundamentals of how we interact with technology are forever changing. Touch, in the computing and NIME communities, has largely been focused on screen-based interaction (Jordà et al. 2007; Hochenbaum and Vallis 2009; Yu et al. 2011), even delving into ways of expanding the sensing abilities of these types of surfaces (Hochenbaum and Kapur 2011; Gelineck et al. 2014). Touch capacitive sensing has not seen many breakthroughs outside of the application to multi-touch screens. At least that was true until Munehiko Sato et al. developed the *Touché* project and a new technique for capacitive sensing called *Swept Frequency Capacitive Sensing* (SFCS) (Sato, Poupyrev, and Harrison 2012). This new technique opens up a whole dimension of touch sensing capabilities and gesture recognition. For the NIME community, the ability to create unique sensors and combinations of sensors that can take almost any shape, will allow for the development of instruments and interfaces designed for the musical gesture, rather than for the electronic sensor. And consequently, for the evolution of Next Music, this type of custom embedded sensor design will enable artists to transform natural and man-made spaces alike, into musically interactive sites of immersion and exploration.

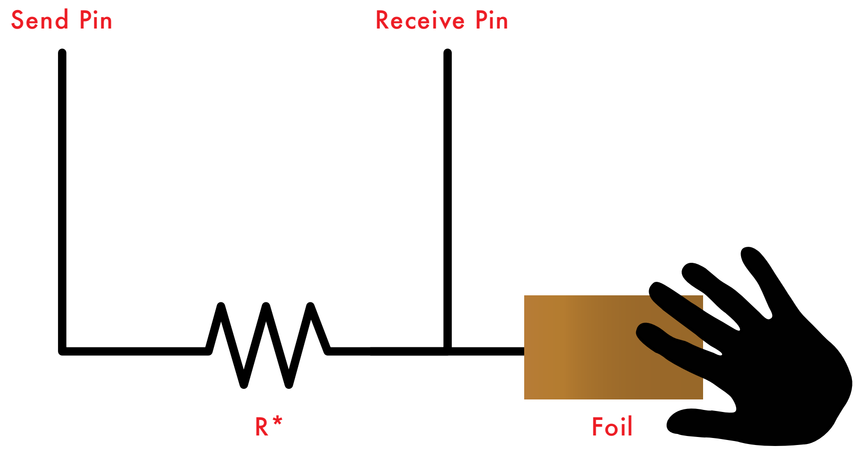

Figure 8: CapSense Circuit Diagram (* Resistor Value of 100K-50M)

2.5.2 Capacitive Sensing: Methods Old and New

In this section we examine how capacitive sensing technology has been achieved in recent years using open-source technologies. This section then introduces a new open-source technique for touch sensing using SFCS and a new Arduino library called *SweepingCapSense,* which makes the technique widely available to instrument makers and artists.

2.5.2.1 *CapSense Library*

The de facto way of creating a capacitive touch sensor on the Arduino is currently by means of the *CapSense*[9] library using the circuit shown in Figure 8. This method is simple and involves very few components, including a single resistor, some conductive material, and two digital pins on an Arduino. The *CapSense* library makes the software implementation as simple as setting up a few lines of code before obtaining a sensor reading.

Altering the value of the resistor in the circuit changes the sensitivity of the sensor, allowing a range of interaction from absolute touch to in-air proximity sensing at a distance of up to 2 feet. Optional capacitors can be used to reduce noise and increase smoothness of the sensor values.

Thus the Arduino and *CapSense* method has been popular in the open-source community due to its simplicity and ease of use. However, the data collected from this method is limited,

[9] http://playground.arduino.cc/Main/CapSense

and in many cases it is only reliable as a binary reading between a touch and no touch. With the *CapSense* method their isn't much more information that can be gleaned from the sensor data. It is possible to detect gesture and movement with an array of capacitive touch sensors, whereby moving from one sensor to another, direction and speed can be determined. Still, even in this case, the gesture is limited to directional movement. Later in this section we will test this method against the SFCS method using identical materials and test gestures to visualize the difference in quality of the data returned.

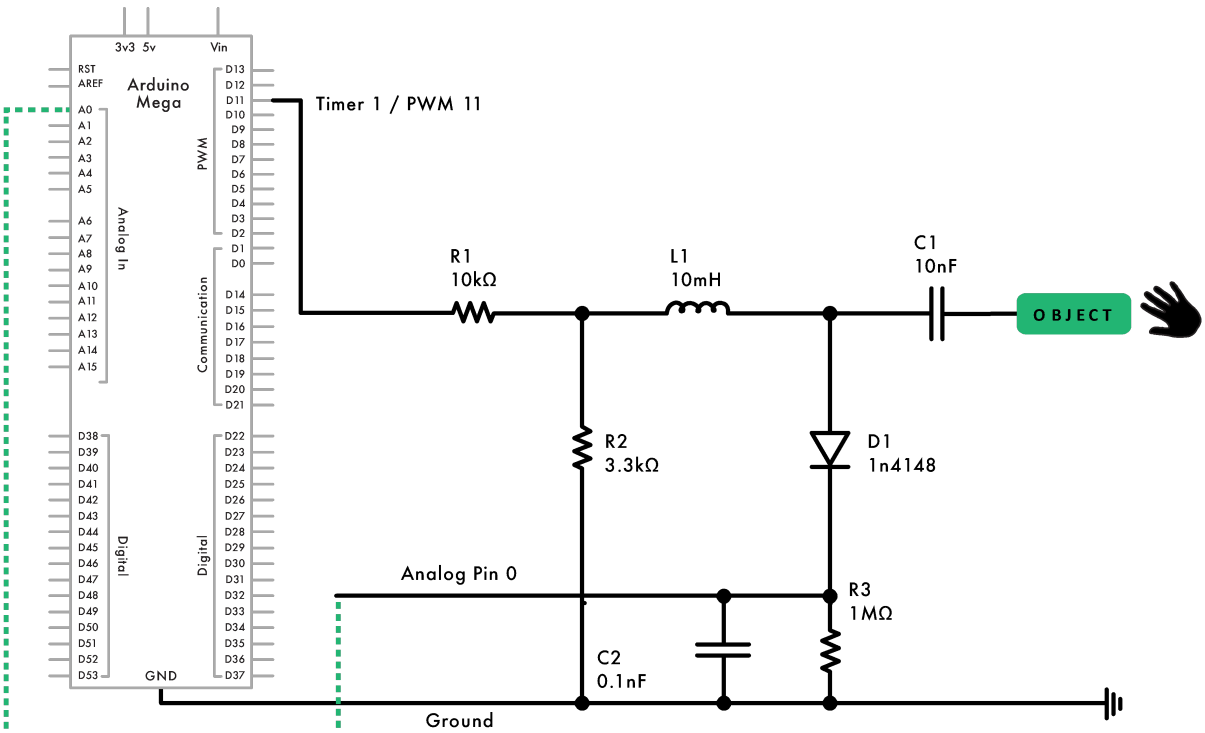

Figure 9: SFCS Circuit for Arduino Mega - Send pins change depending on Arduino model and Timer used

2.5.2.2 *SweepingCapSense Library*

While most of the progress in capacitive touch sensing has been seen in the mobile computing field, Sato et al. at Disney Research have taken a big step forward for the embeddable capacitive sensor with the *Touché* project. By exploring ways in which everyday objects can be transformed into rich touch sensors, that not only sense touch itself but also the manner of touch, they were able to implement a unique style of *Swept Frequency Capacitive Sensing* (SFCS). The novelty of this method is found in the "Swept Frequency" portion of the name. While traditional capacitive touch sensors detect conductivity at a fixed frequency, SFCS uses an entire range of frequencies. By sweeping through this range of frequencies the readings from a single sensor take on a second dimension, which, when graphed, makes it easy to actually see the difference between a single-finger touch, a two-finger pinch, and a whole-handed grab.

Below, we will analyze the difference between single and swept frequency techniques, highlighting the rich information obtainable by the latter technique. Before looking at this data, it is important to note that the *Touché* project utilized a patented and purpose-built sensor and microcontroller, which is to say, it is not available, nor will it be, to the active DIY or open-source communities. Intrigued by this work Nikolaj Møbius,[10] with the additional support of Mads Hobye,[11] both members of the interactive art studio Illutron,[12] ported this technique to the widely used Arduino platform. By acting directly upon the built-in timers of the Arduino, they were able to output a waveform and change its frequency, allowing for SFCS. Details of how they did this will be discussed below.

While Sato et al. were the first to implement the SFCS method, it was Nikolaj Møbius who first ported the method to the Atmel-based Arduino platform using passive components. The *Touché* project used an AD5932 Wave Generator IC that can produce a sine wave with very high resolution at frequencies between 1kHz and 3.5MHz. While the Arduino can also produce a signal at these frequencies, its resolution is lower and it can only produce a square wave (containing unwanted harmonics). Møbius suggests using a simple LC (inductor – capacitor) circuit, also known as a resonant circuit, to help transform the square wave into a sinusoidal wave (Nikolaj Møbius 2012). Figure 9 shows the components and schematic for Møbius' technique.

The original code provided by Møbius directly addresses the hardware timers on the Arduino and requires the user to have low-level microcontroller programming experience. This work aims to make the user workflow with SFCS and the Arduino as simple as the *CapSense* library, and thus the *SweepingCapSense* library was developed.

The *SweepingCapSense* library uses two main classes, *SweepingCap* and *Touch*, which control the sweeping frequency signal and store the touch data (respectively). When constructing the *SweepingCap* object there are two options:

[10] http://dzlsevilgeniuslair.blogspot.se
[11] http://www.hobye.dk
[12] http://illutron.dk

1. *SweepingCap::SweepingCap(180, 1)* – sets how many frequencies to sweep (up to 198) and which of the 16-bit timers to use.
2. *SweepingCap::SweepingCap(180)* – sets all four 16-bit timers at once, by only passing in the number of frequencies to use.

To use the *SweepingCap* object, requires only two other steps:

- *SweepingCap::setup*() – must be called inside the Arduino setup function, and automatically sets up the timers specified in the constructor.
- *SweepingCap::sweep(int frequency)* – this method must be called inside a for loop that counts to the number of frequencies set originally (note: passing the iterator from the for loop into the sweep() function sets the frequency).

The *Touch* object requires only one variable to set the analog input pin (passed in its constructor), however the *Touch* object is not a requirement to achieve SFCS, rather it is a convenient data structure to store and retrieve the values returned at each frequency. The *Touch* object requires minimal steps to function as well.

- *Touch::reset()* – must be called at the top of the main loop to reset the variables inside the *Touch* object.
- *Touch::readPin()* – is called to automatically read and store the value at the current frequency. This method must be called inside the sweeping frequency *for* loop.
- *Touch::topValue()* and *Touch::topPoint()* – set and store the top value and frequency, respectively.
- *Touch::interpolate()* – returns the interpolated top value for smoothing purposes.

2.5.3 Evaluation: Methods Old and New

This section provides a comparison between the *CapSense* and *SweepingCapSense* based techniques. Below, the sensor readings for two test materials with different touch positions and touch gestures for both the single and swept frequency techniques for capacitive sensing are provided. The objects used were a copper electrode (approximately 5cm x 2cm) and a small rosemary plant (approximately 15cm tall) with the electrode placed in soil. The same series of tests were conducted with both objects using the two libraries and circuits, and the results can be seen in Table 2, Figure 10, and Figure 11.

Looking at the results in Table 2 we see that using the *CapSense* library with both test materials and each test gesture only a single piece of data is obtained. While some gestures may be recognizable with this single piece of data, it is difficult to clearly understand these results as unique gestures. At least the plant's test results show that the dispersion of the signal through the soil and into the plant helps to create a wider range of results. However, these results tend to be very susceptible to outside influences that result in a lowered repeatability over time.

Object	Gesture	Results (by position)				
Electrode	Position	1	2	3	4	
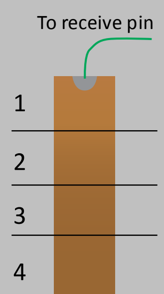	Finger Pad	1426	1430	1432	1428	
	Fingertip	1375	1385	1391	1400	
	Two Finger Pinch *(Edges)*	1450	1425	1400	1390	
	Two Finger Pinch *(Flat Sides)*	1450	1440	1435	1483	
	Two Finger Press *(Flat Side)*	1470	1475	1482	1483	
Plant	Position	1	2	3	4	5
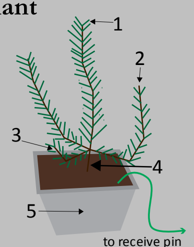	Single Finger Touch	61	59	60	900	67
	Two Finger Pinch	70	57	150	1247	389
	Three Finger Pinch	116	80	240	1040	435
	No touch *(2cm away)*	30	33	60	76	60
	Grab	118	82	163	571	730

Table 2: CapSense Test Results

Using Møbius' circuit, our new *SweepingCapSense* library, and the same objects and gestures, results in a different kind of data altogether. As seen in Figure 10 and Figure 11, the data representation of a single gesture is a two-dimensional stream that not only returns a peak value similar to the previous method, but also returns the values at every other frequency in the sweep. The resulting shape gives greater insights into the object's response to the interaction and allows a novel opportunity for gesture recognition. It is important to note that this technique is not only sensitive to gesture but to position as well. Figure 12 shows the results of the same gesture, a two-finger pinch, in all the testing positions on the rosemary plant.

The data clearly shows a vast improvement over the single frequency method in many respects. Even when the resonant frequency is similar there is enough difference throughout the rest of the values to discern gestural subtleties. Also, the repeatability of these shapes, if not the specific peak frequency and value, is very high. Ultimately, a single sensor could detect many different gestures and interactions.

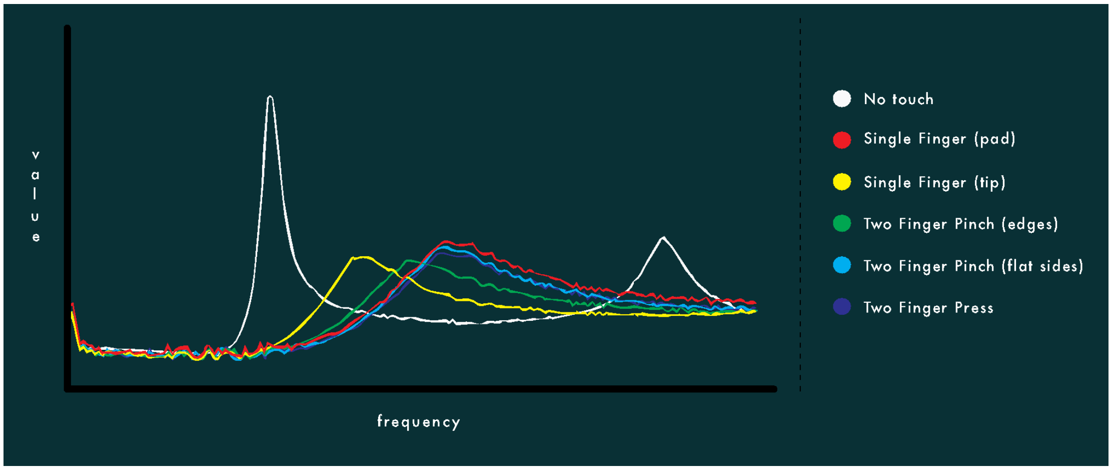

Figure 10: SweepingCapSense Test - Electrode - All Gestures - Position 1

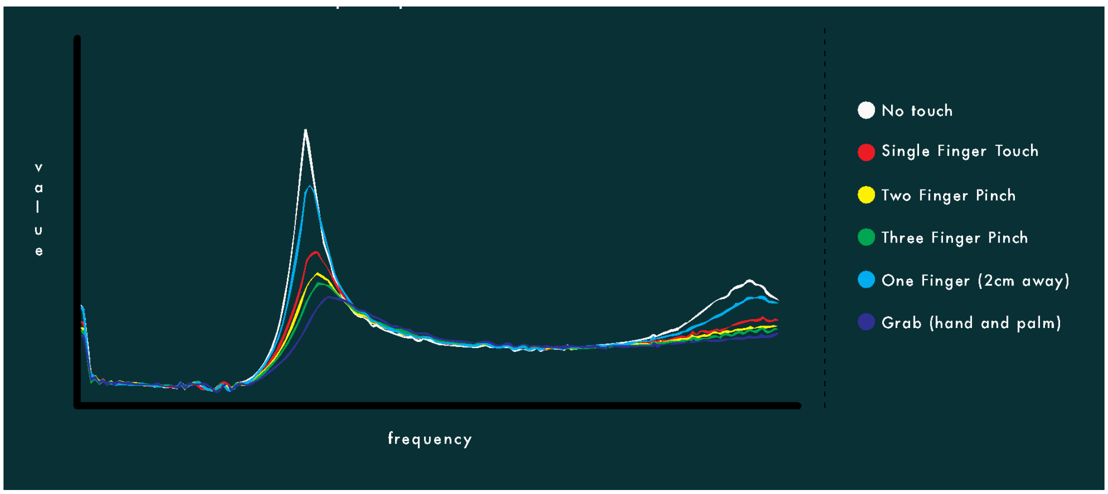

Figure 11: SweepingCapSense Test - Plant - All Gestures - Position 1

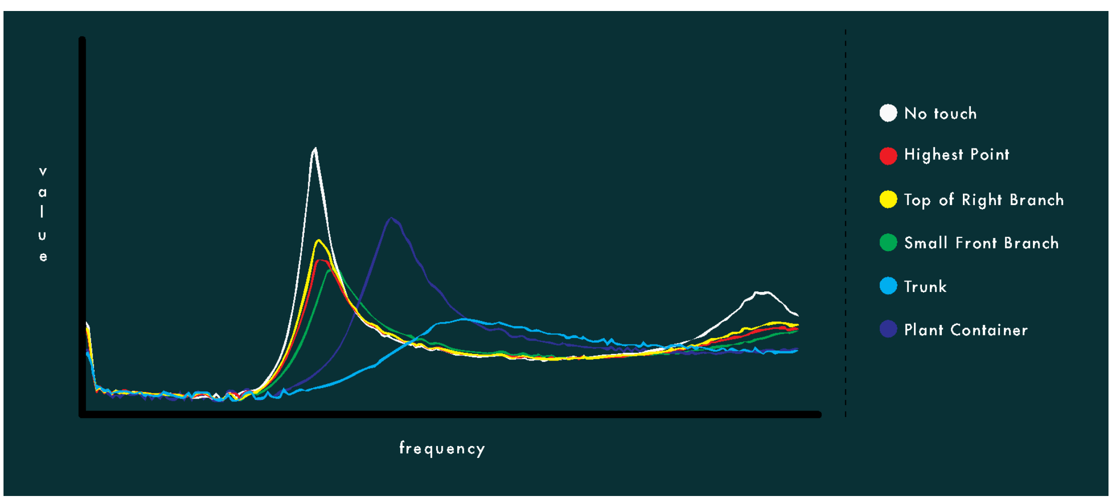

Figure 12: SweepingCapSense Test - Plant - Two-Finger Pinch - All Positions

2.6 Future Work

As a permanent installation, Cultivating Frequencies will require a lot of attention and monitoring by the creator and the owner. Unproven, especially in the outdoor conditions that it will inhabit, this project will undoubtedly require some degree of fine tuning, maintenance, and updating. To help facilitate this the author has created an "Owner's Manual" which includes much of the same information presented in this chapter, with additional detailed information about the exact hardware and software used.

Several desired elements of this project, at the time of writing, have yet to be completed. Most importantly, the actual installation of the project has yet to be completed. Upon its final installation, the author expects to return to Spain and spend a short time fine tuning and testing the installation as a whole, and to make sure it is in good working order before handing over controls to the owner. One of the important aspects of this handover is the development of an iPad application that will serve as a graphical user interface to the otherwise graphic-less interface as well as provide certain master controls. These master controls will include utility functions such as turning the installation on and off, as the computer running the software is embedded into the garden's sculptural framework and is not easily accessed. The controls will also include the ability to adjust certain settings, such as the sensitivity of the user detection, the length of the composition that is played, and even certain sound design elements for the touch interaction.

Further work is required on the project's website which will be host to the collected data and even some of the generated compositions. The last element to be completed within the original scope of this project is the standalone musical application that will allow the owner to listen to his garden when he is not physically present. This standalone application will also be available to anyone who wishes to listen to the garden and will require only an Internet connection to pull the latest data from the online host.

Outside of the original scope of this project, the author intends to continue to develop variations of this work. With the goal of a more "plug and play" setup, the engineering and design of an all-in-one wireless and solar-powered sensor is the ultimate goal, though the feasibility of such a feat is as yet undetermined. If plausible, the author intends to create a series of sensors that would allow anyone to transform their gardens into "nodes" as Cultivating Frequencies itself transforms from a single instance to a network of augmented gardens.

2.7 Closing Thoughts

Cultivating Frequencies exhibits many more unique qualities compared to the following projects presented in this document. As the last conceived of the group it builds upon ideas that were barely formed during the process of the other projects. It also includes the first real contribution of the author to the open-source community in the form of the new Arduino library, *SweepingCapSense*. And, in the author's opinion, marks a pinnacle moment, where the author's aesthetics, technical skill, and conceptual reasoning successfully converge on a level previously unachieved. Lastly, it is the first of the author's works to become a permanent installation.

Cultivating Frequencies has had the greatest influence on forming the concept of Next Music. Taking inspiration from the ideas of Next Nature and applying them to musical practice and thought, Next Music owes a great deal to the work done by the participants of the Stligat Pavilion. An unprecedented collaboration between the author, Cloud9 Architects, and a group of students from Art Center College of Design, yielded amazing work in the area of environmental design, of which Cultivating Frequencies was but a small part.

In the context of Next Music, Cultivating Frequencies exemplifies how technology can be used to augment even the most intimate and delicate of spaces. Furthermore, this augmentation acts in service to the space rather than as a novelty imposed upon it. The interaction and musicality is used as an extension of the garden itself, giving it a voice and with that voice a means of communicating with the outside world. One must learn through listening and through interaction to understand this voice, but the hope is that the gardener becomes attuned to the voice of the garden. Stretching beyond the surface conditions, the garden is also given a memory, which it only shares with the more dedicated gardeners and visitors. With a little bit of coaxing the garden will share its memory, stretching into lyrical melodies, sometimes highlighting struggle or remembering a similar day in its past.

The garden space is a rich space with a long history, many cultural contexts, both of which are fraught with dualities. It is a space made of natural elements, grouped in an unnatural way. Inherently interactive, the gardener tends to the health and wellbeing of the plants, but often for their own aesthetic or personal pleasure and not strictly on behalf of the plants themselves. These days, many garden owners do not even actively participate, opting for hired professionals and automated systems to care and propagate their garden spaces for them.

Cultivating Frequencies does not judge this behavior, but it does encourage interaction, both physically and mentally. While the owner is present, they will hopefully be drawn from their home to visit this unique space, and not just once, but over and over again. As they revisit this space, the nature of their interaction will grow as the individual plants grow, though at quite different rates. Furthermore, their understanding of the garden will also grow, allowing them to understand the coded message of each song created and played. While they are away, they will be able to think about this space, and not having to just rely on a photo or a memory, they will be able to hear their garden, just as it is, right then. Giving them instant access to what is happening now, thousands of kilometers away, on a beautiful hillside, overlooking the Mediterranean Sea, and hopefully feel as if they are there. Beyond their own personal enjoyment, they will be able to share this intimate space, at least a part of it, without overwhelming the garden or the owner.

Next Music allows for the simultaneous expansion and collapse of a space, allowing this unique musical garden to be global and local as well as public and private. Despite the mutual exclusivity of these ideas, it is possible within the augmented space to allow for such conflicting perspectives. While the experience cannot be the same for the remote and live user, the simultaneous existence of each does not impair the experience of the other. It is the author's hope that it will in fact enhance the experience of the live user knowing that others can share in this experience despite being far removed from the physical location. Disconnected by physical space but connected through the augmented space.

Chapter 3
The Third Room and the Augmented Musical Space

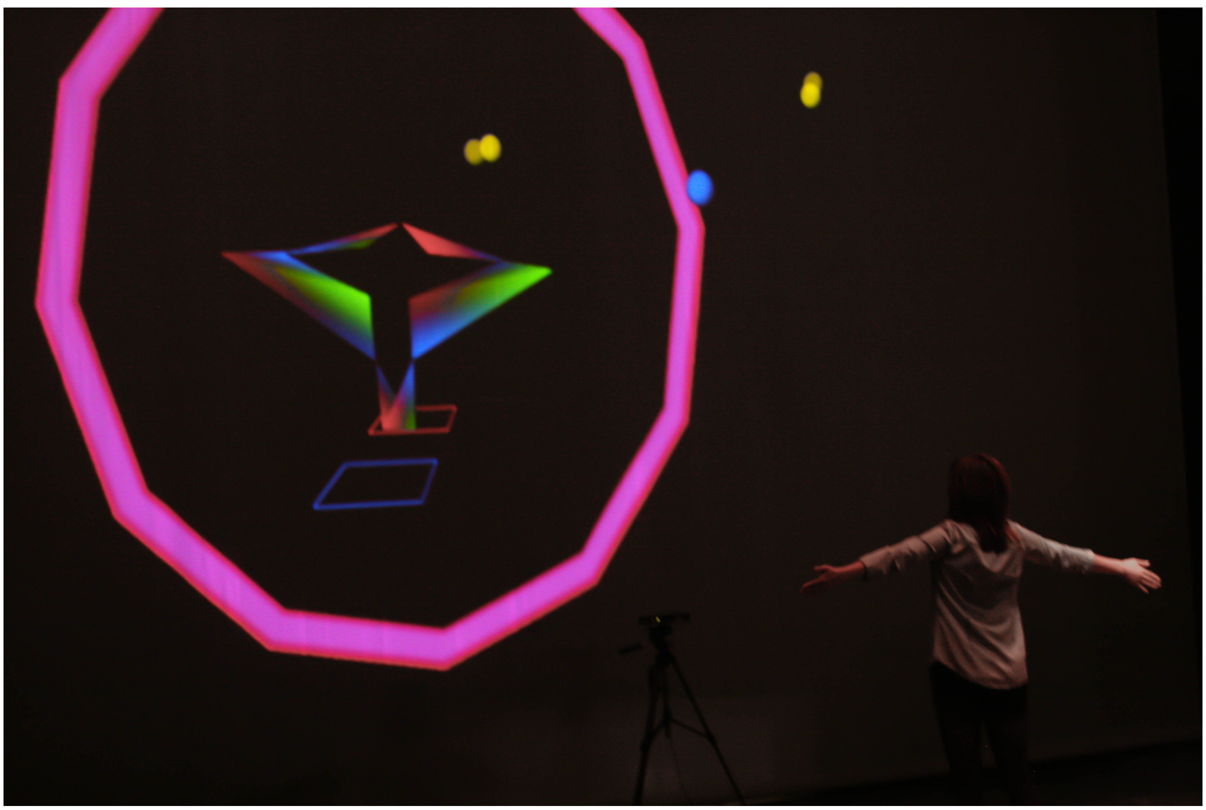

Figure 13: Third Room on display at the CalArts Digital Arts Expo 2013 (photo credit: Scott Groller)

Next Music, as defined and presented in this document, is focused on space. Moving along a scale of small and intimate spaces to large public spaces, the Third Room occupies an interesting middle ground. The "musical space" stated in this chapter's title refers not to a single kind of space, but the many spaces that are used by musicians, composers and sound artists. The performance space, the recording studio, the practice room, and even the gallery space all fall within this category. The Third Room was the author's first venture into what would become Next Music. Starting as a singular project, it took on a broader and greater significance when it was prepared for submission to the International Conference on New Interfaces for Musical Expression (NIME). The ideas expressed in the published paper "The Third Room: A 3D Virtual Music Framework" (Honigman, Walton, and Kapur 2014) were the seeds that would grow into the overarching theme and main contribution of this thesis.

The Third Room is an interactive musical installation that transforms a physical space into a Next Music style performance space. Although packaged as an installation and performance space, the same framework and sensor-based interaction can be applied to any musical space, and will be presented and discussed within multiple contexts and applications simultaneously. Augmenting a musical space allows for the integration of physical and digital instruments in new and novel ways. While all rooms are limited physically, the augmented room allows one to extend their musical control beyond the physical space, and the hardware within that space, essentially making infinite copies of the room with different sets of instruments, different rules, different generative systems, etc.

In this chapter the Third Room will be presented as a singular piece as well as the conceptual basis for further work in developing the augmented performance space. First, a brief background will be given on Kinect-based and spatially driven musical installations and augmentations (pg. 41). Second, a system overview of the installation will be explained (pg. 43). Third, a discussion on natural interaction, performance, and interactivity will be presented (pg. 45). Fourth, the idea of physicality is introduced as an important element of virtual interactions by way of the physical interfaces known as the Ember and the Fluid (pg. 49). Fifth, the idea of blending the physical and the virtual is furthered by the introduction of mechatronic sound-objects (pg. 56). And lastly, a discussion of future work and the implications of augmenting the performance space on future performance is presented (pg. 57).

3.1 Background

In the past ten years we have seen advancements in computer technology and interfacing in both music and digital media. In music, "*the Grid*" has had a huge impact on live performance and composition of electronic and computer based music. However, the grid, by design, is rigid, and not well suited for melodic and improvised expression in performance, instead relying on loops and often pre-recorded material. That is not to say it is impossible for a performer to be musically expressive within the grid, however, it seems, that the grid lends itself more readily to certain aesthetics and genres more than others. This limitation is what sets it apart from other musical practices and is also what makes it unique in itself.

In digital media, we have seen the advent of depth camera technology and how it has changed the way users interact with games, media, and art. With the release of Microsoft's Kinect[13] camera, and the subsequent hacking of the device for use with open-source software, there was an immediate explosion of musical and non-musical projects exploiting the ability to ditch the tangible controller and let the users put themselves physically into the interaction. A lot of work has been done to use this kind of technology to connect musical generation with a dancer's movements (Berg et al. 2013; Frisson et al. 2013; Bevilacqua, Naugle, and Valverde 2001), the dancer being an ideal candidate because of their high level of body awareness, physical ability, and movement vocabulary. Many other experiments have integrated this interaction into performance, however the performance tends to be more focused on spectacle, which manifests itself either as highly stylized body movements[14] or in over-the-top virtual interfaces.[15] However, very little experimentation has been done to use this interaction to create new environments for musical composition and performance. The potential of an interactive augmented space where the physical and virtual collide in an immersive reality of sonic possibility and creativity has yet to be realized.

Inside an augmented composition and performance space one is free to create all manner of "impossible" instruments and interactions that could not be achieved in the physical

[13] http://www.xbox.com/en-US/Kinect
[14] http://ethnotekh.com/project/ethno-tekh-v2/
[15] http://www.v.co.nz/#the-motion-project

world. However, it should not be our aim to replace tangibility. It is important to note, that the novelty of virtual interaction can undermine itself when it is relied upon too heavily, forcing users to completely abandon touch, one of the most important senses. The potential for combining these interactions is great, and there is a lot of work to be done in exploring new creative frameworks and approaches to the arts. The cumulative effect of integrating virtual and physical interactions reinforces the unique qualities of a new creative space.

This project was influenced by the many Kinect projects that used the interface to create virtual interaction with sound and music. Projects like AHNE, the Audio-Haptic Navigation Experiment, which is an audio-haptic user interface that allows the user to locate and manipulate sound objects in 3D space with the help of audio-haptic feedback (Niinimäki and Tahiroglu 2012). This project strives to address the problem of how to guide a user in a 3D space, providing both audio and haptic feedback, in the form of a glove, to inform the user of their proximity to the invisible sound objects. While addressing this issue in a unique and intuitive manner the actual sound created by the interaction seemed more of an afterthought to the novelty of the technology.

Other musical projects involving the Kinect camera use it as an extension of performance practice and expression, such as Qi Yang's augmented piano performance system (Yang and Essl 2012). Yang uses gestures and a small projected area above a keyboard to create new extended techniques for manipulating the sounds of the piano, allowing the performer to affect multiple controls in a way previously impossible with traditional synthesizer controls. This approach integrates the new technology into an instrument with a long-standing practice of techniques and performance, which qualifies it as a Next Music practice. Not every Next Music application needs to completely abandon the history of music and all of the wonderful instruments and practices it has produced. The augmentation of acoustic instruments is perceived as part of the movement towards the future of Next Music. The next step in this practice of creating "hyper-instruments" is to manufacture our traditional instruments to include the necessary hardware for augmentation and extended digital techniques. As is, the success of a system like Yang's augmented piano on influencing the future of performance is limited to the number of practitioners who adopt it as a new extended technique. However, if new instruments are created specifically for, with, and out of this new technology, a whole new field of music composition and performance can be realized. One that can be as rooted in, or removed from, traditional musical practices as the user/composer/performer wishes it to be. This new practice will exist in a true third space, as it bridges the gap between physical and virtual interaction, but it will do so by utilizing real spatial data to link the two spaces, augmenting and transforming it into the Third Room.

This project aims to forge a new approach to musical composition and performance by developing a system that incorporates 3D spatial data, a virtual environment populated by virtual objects, a physical environment populated with physical objects, and sound design with mechatronic sound-objects, linking both spaces, that is expressive and has the potential for a wide range of complexity. Together these elements create the Third Room, which is what we will call this new augmented space for music composition and performance. In this chapter the potential of the Third Room augmented musical environment is presented as the basis for future development of new musical systems and as a Next Music approach to composition and performance.

3.2 System Overview

When a user enters the Third Room their body is detected by the Kinect's depth camera. Their skeleton data is used to track their movement through the physical environment, displaying their position in the projected virtual environment and enabling virtual modes of interaction. In the first iteration, the virtual room was populated by different kinds of objects that the user can interact with, using their own body as the interface. However, it was found that users had a difficult time finding these objects, so in the latest iteration the users control object creation and destruction through various gestures. The virtual objects interact differently with the user and the environment, depending on their type, and each has their own sonic characteristics or affects the sound of other objects. The "Ball" is an object that a user creates by waving. It can be thrown and caught, and holds a note value that is triggered when it hits the walls of the room. The "Blob" is an object that is created when two or more users group together in close proximity. When they do, a bass note is triggered as an amorphous shape envelops their avatars, and the parameters of the tone are randomly assigned to the different joints making up the shape. The "Box" makes up the virtual environment itself and has various interaction possibilities. The boxes can be turned on and off like buttons, transforming the entire virtual room into a step sequencer allowing for a different level of compositional control. The "Taurus" is created when a user brings there hands together and grows bigger, circling their avatar, as they pull their hands apart, as seen in Figure 13. The Taurus controls the decay of tones created by the Ball objects, which is an important compositional parameter, which be explained shortly.

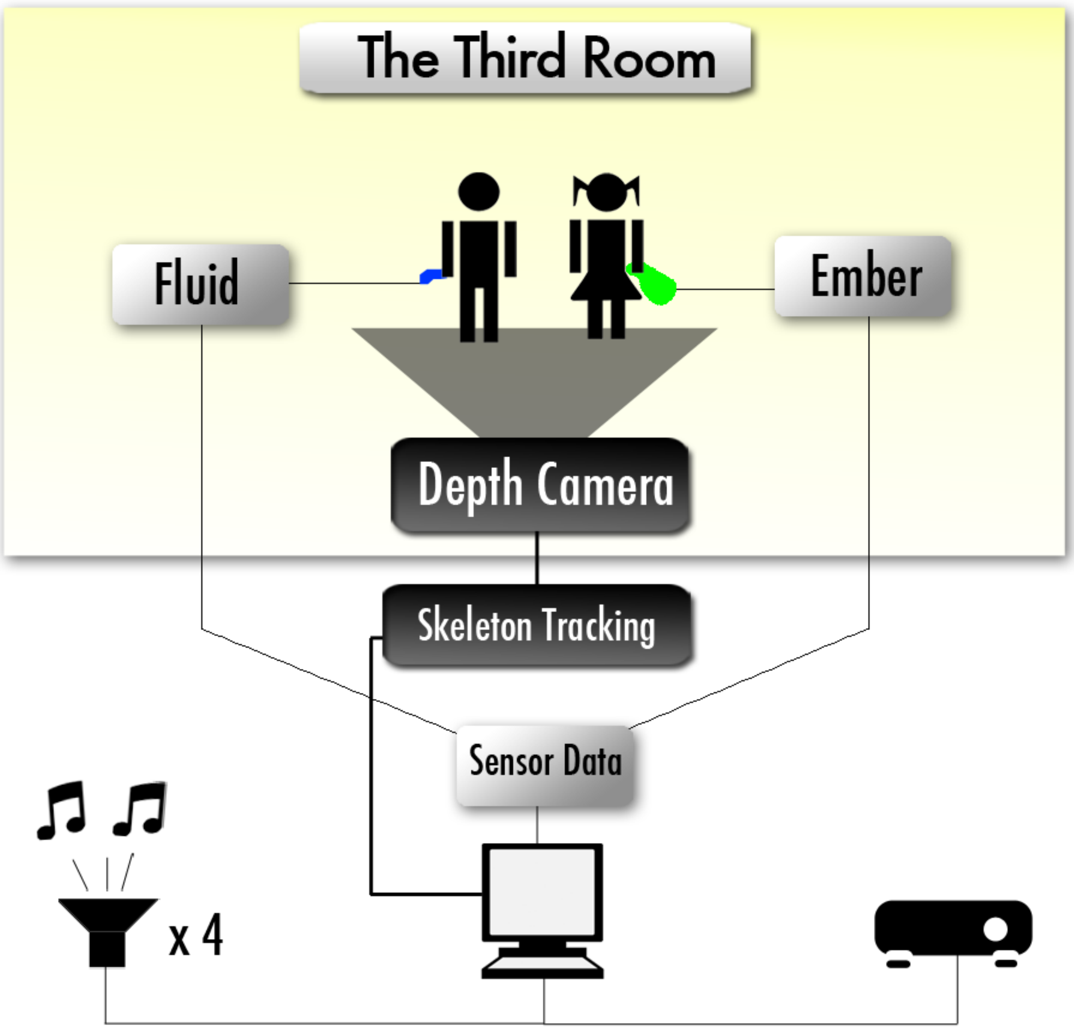

Figure 14: Third Room System Overview Diagram

Also present are a pair of physical objects that play an active role in the Third Room, known as the Fluid and the Ember. The Fluid, an abstract remote, designed as a hand sculpture, invites the user to explore different ways of holding it. As they explore the physical nature of the object they simultaneously are manipulating the sound and the projected environment, learning to control it through the process of discovering its abilities. Also present is an electronic mbira-style interface, called the Ember, which offers the possibility of melodic creation within the immersive experience. When a user plays the Ember a new ball object is created that corresponds to the played note. This new object can then be manipulated inside the virtual room by any of the users. The other sensor data from the Fluid and the Ember is collected and used to make macro level changes to the soundscape as well as the virtual environment. By combining the physical and virtual interactions we increase the depth of possibility for the subtleties of musical creation within the space. This combination of all the sensors working together allows the users to interact in both virtual and physical ways, bridging the two spaces, each augmenting the other in turn, helping to create the Third Room.

3.3 Natural Interaction and Performance

The Kinect camera allows us to transform our bodies into what is essentially a large digital sensor. In itself it is simply a combination of cameras but when used with "Natural Interaction" libraries such as OpenNI and NITE[16] it becomes a very powerful tool for interaction. The main function of these types of libraries is to recognize a body in space and capture different types of data including hand recognition, gesture recognition, and skeleton tracking. Skeleton tracking is the process of recognizing a body and then inferring the location of each major joint and limb of a person's body in three-dimensional space (Cook et al. 2013). Skeleton tracking is the main technology used in conjunction with the Kinect camera as a video game controller as well as the driving force behind this project.

The skeleton tracking for this project was achieved in Processing[17] using the SimpleOpenNI library developed by Max Rheiner (Borenstein 2012), which allows easy access to the main functionalities of the OpenNI and NITE libraries. Once setup, tracking becomes very easy, and does not require a calibration pose, which means a user can simply walk into the space and the program automatically starts tracking their skeleton data. The tracker did exhibit difficulties in detecting users whose clothes obscured their legs, like a dress, and so some users had to use the calibration pose to force recognition. The skeleton data is sent via Open Sound Control (OSC) (Wright and Freed 1997) to the virtual room's program where the user is displayed as a digital avatar. Multiple users can be tracked at once, but the program sometimes

---

[16] http://www.openni.org
[17] http://processing.org

gets confused if a user passes between the camera and another user. However, the program is robust and can easily drop and add users as they enter and exit the space.

The virtual room is a 3D computer generated space that was created using the Cinder[18] C++ library and OpenGL[19] as seen in Figure 13 below. Where the room was once prepopulated with virtual objects, now it is at the user's discretion to create and destroy these objects. As soon as the user is detected their avatar appears inside the virtual room, through which they can interact with the objects.

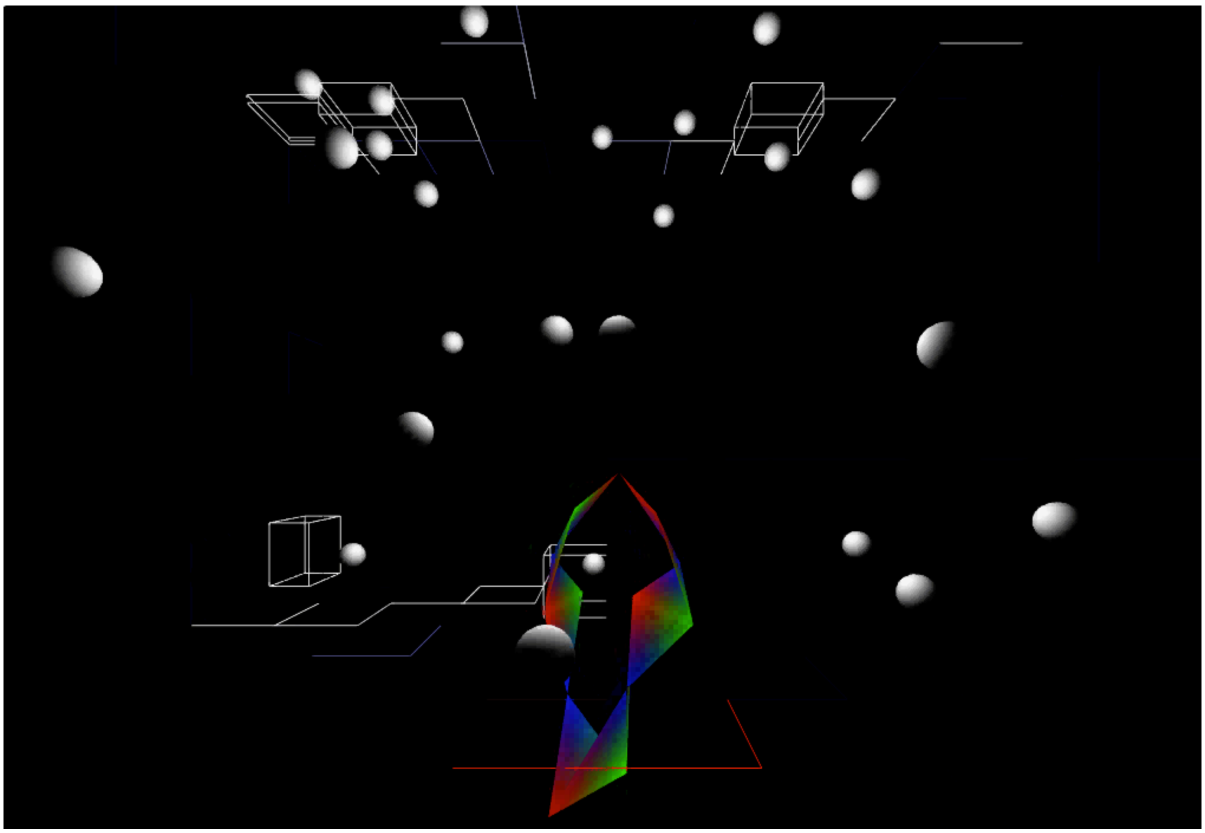

Figure 15: Third Room Virtual Rendering in Cinder and OpenGL

The main object is a virtual ball that the user can create, or catch, and throw across the room, where it bounces without friction or gravity, creating a variable rhythmic and melodic loop. Another kind of object, the blob, is a special object that requires multiple users to create. When two, or more, users come within a certain vicinity of each other their avatars become enveloped in an OpenGL mesh that triggers a different kind of sound. While inside the blob the users have a newfound control over this sound until they move apart. If new users join the blob

[18] http://libcinder.org
[19] http://www.opengl.org

the note changes and the parametric control is dispersed among them. All the virtual interaction is based on a combination of proximity, 3D hit testing, and very simple gestures, eliminating the need for advanced gesture recognition or a glove interface. Instead, to grab an object the user just finds it in the virtual room and "touches" it. To create a ball object the user simply waves at their avatar. To clear the room of objects the user reaches across their chest with their right hand and pulls it back across in a quick swipe motion. Simple gestures like these were designed in lieu of using advanced gesture recognition to facilitate a wider range of interactions.

The use of the wave as the main gesture came out of observations from the first iteration of the project. It was noticed that when users saw their avatar appear on screen they would almost always wave. This is a quite natural reaction that humans have when presented with a digital representation of themselves. It is simultaneously the quickest way to infer whether or not they are actually seeing themselves and to what degree there is a connection between what they do and what they see. On observing their digital-self waving back at them they establish an understanding of the connection allowing them to further explore other possibilities. From the frustrations born from users not being able to "find" the objects that populated the first iteration came this unexpected and quite enlightening observation. So, when a complete overhaul of the project began, the use of simple gestures for object creation and control was of the utmost importance. The added benefit of simple and naturally occurring gestures is that they have the possibility of happening without giving the user any instruction, allowing for a truly exploratory experience whereby the user discovers the means of their own interaction. Furthermore, as observed by the author, the act of discovery seemed to be a joyful one, which in turn encouraged users to help and instruct each other on their own techniques.

One of the most difficult aspects, other than the limitations of the Kinect camera itself (low frame rate, poor resolution, and limited range), was developing a system of visual cues to help the user identify their location in relation to the other objects in the room. This was the major flaw in the first iteration of this project, where the objects lived within the room and were statically created and could not be destroyed. After seeing many people struggle to find the objects a new approach was developed. Hence, we introduced the means to dynamically create and destroy objects. This simple gesture interaction gives people the opportunity to interact with the Third Room while removing the struggle to orient them within it.

The sonic aspect is, of course, a very important part of the work as well. As a combination user-driven and generative piece a lot of thought was put into the type of sounds and the degree of control the user had over them. As the creator and composer, each choice acts as a limitation on the personal creativity of the user. However, some structure is useful in creating something that is easily approachable. A compromise must be made between the intentions of the composer/creator and the will of the users. Presenting too much freedom can make it too hard for the users to create something they can enjoy, while conversely, too much limitation voids the compositional possibilities of the users, who, by interacting with the Third Room, become both performers and composers.

Furthering the exploration of space the Third Room utilizes four speakers, in quad sound configuration, to help create a surround sound experience. Each sound follows the object as it moves through the room, further integrating the two spaces. Depending on the virtual object, the visual interaction and sound is different, resulting sometimes in a one-to-one style interaction, while others are less straightforward. The more objects that are put into motion or interacted with, the more complex the soundscape becomes. The level of interaction is compounded when the physical objects are used as well.

The sound design and synthesis was created in Native Instrument's Reaktor[20] using a custom built set of instruments. Simple sine tones are used as the primary melodic content to allow for a greater dynamic range. However, melodic, timbral, and rhythmic complexity is achievable. Each ball object has a corresponding note value attached to it, and when the ball is thrown around the room it creates a deconstructed melody. As it collides with the walls it triggers the corresponding note. Depending on the decay time that is set by the Taurus object, the ball either triggers the root note, if the previous note has fully decayed, or the next note is a harmonic series, if the previous note is still sounding. There are several different sets of harmonic series' that can be activated, each providing a different palette to work with. These harmonic series' are based off simple mathematical equations ranging from a natural harmonic series and others that are related to the serialized electronic experiments of Karlheinz Stockholm, specifically his *Elektronische Studie II* (*Studie II; Nr. 3 [von] Elektronische Studien.* 1956).

Using the physical objects, the users can take more advanced control of the overall sound characteristics. The Ember allows for specific note creation as well as parametric manipulation, including setting the decay time of the note objects and choosing the harmonic

[20]http://www.native-instruments.com/#/en/products/producer/reaktor-5/

series being used. The Fluid can be used both to create sound and manipulate the environment, depending on how it is held and moved.

3.4 Accounting for the Physical in Virtual Performance

Whenever a new technology emerges there are those whose enthusiasm wins out over their reasoning. In the case of 3D sensing it seemed like many thought the age of hardware controllers was over, as there was no need to hold on to something anymore. But physical tools are an essential part of human history, culture, life, and especially music. Through these tools we discover new things about the world around us, and it is through repeated interaction that we gain a physical mastery of them. It seems less likely to achieve this mastery when it comes to waving your hands around in the air, though not impossible as showcased by the few Theremin virtuosos. Touch is an incredible sense that carries with it its own knowledge, an innate and intimate knowledge, which is both instant and learned over time. It would be a waste then, to disregard the physicality of music, and the intimate knowledge gained by exploring music through a physical instrument.

It was this idea, of the importance of physicality, that inspired the inclusion of two digital musical interfaces into The Third Room project, and in general, the inclusion of physical objects and hardware into the concept of the Third Room.

3.4.1 Ember: Digital Mbira Interface

The Ember is a handheld digital interface modeled on the mbira, a traditional African lamellophone (Berliner 1993), pictured on the left side of Figure 16. The driving force behind the Ember's design was a desire to create a new kind of digital instrument that could be played expressively and not just another knob filled or grid based MIDI controller. Furthermore, the Ember is not just a digitized traditional instrument. The functionality and form of the traditional instrument were used as the point of departure to create a new kind of handheld digital instrument. The impetus behind the physical design was to create a handheld instrument that allowed for both melodic and continuous control. Modeled on the mbira, the mechanics of the Ember are directly related to that of a lamellophone, in that the sound is generated from the vibrations of long tongue-like plates that are fixed at one end. A lot of consideration was put into the physical design of the Ember and it has gone through several revisions.

Figure 16: Traditional Mbira (left) and First Ember Protoype (right)

Originally it looked very similar to an mbira, other than the fact that it was made of cardboard and had all manner of sensors and wires sticking off of it, as seen on the right half of Figure 16. However, it became very apparent that the design of the original prototype was cumbersome to hold and to play. A combination of too many sensors and the shape contributed to awkward playing and an inability to hold it steady with one hand. So, a radical redesign was necessary if musical proficiency was to be achieved. The designer Andrew Capeluto was instrumental in leading the redesign and rendering of the second prototype and the final instrument.

First, a method of holding it securely with one hand was needed. The mbira is designed with a hole that the player places their pinky through to support the instrument. So playing on this idea a new body shape was designed that included a hand hole, so that the user's left hand can fit through it, allowing them to easily and comfortably hold the instrument with one hand. Although this restricts the use of the left hand, it adds a level of stability while playing and frees up the right hand to do other tasks, on or off the instrument. The shape of the body and placement of the handhold were designed to position the two hands at a comfortable angle. This angle allows for ease of play and also promotes prolonged playability by keeping the wrists straight and avoiding awkward or unhealthy wrist positions.

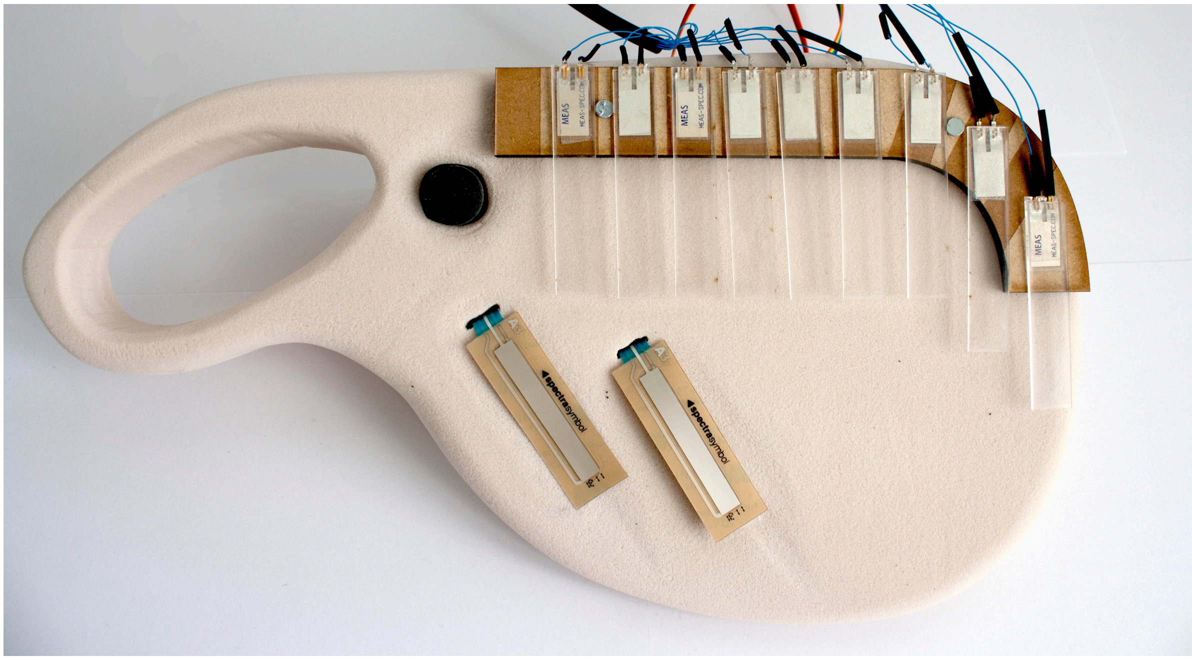

Figure 17: Second Ember Prototype

Second, the amount of sensors needed to be trimmed down to accommodate the new body shape. In the end, nine piezo powered tines, two membrane potentiometers, and a joystick were selected. The sensors are distributed between the two hand positions, the tines and the membrane pots for the right hand and the joystick for the left. In this setup the left hand becomes responsible for switching functionality, scrolling through lists, or adding expression through the joystick, while the right hand is primarily responsible for melody and continuous control through the tines and the membrane potentiometers.

When the shape was set we modeled it using high-density foam to finalize the body size, ergonomic features, and layout. The bridge for the tines was laser cut with rasterized insets to help hold the tines in place and help isolate the vibration of the individual tines. The tines themselves were laser cut out of 1/16" acrylic and don't actually vibrate like a real mbira. From the first design experiment it was concluded that a stiffer tine allowed for greater control over attack velocity and provided more consistency overall. The second prototype is pictured in Figure 17 without the top of the bridge so the piezo films are visible.

For this interface a unique method of using piezo films was developed to simulate the tines of the mbira. Sandwiched between the bridge and the laser cut acrylic, the piezos are safe from accidental triggering while picking up direct hits to their connected tine. When the keys are struck, the piezo picks up the vibration and registers it as a Note On message, even detecting different velocities depending on the pressure applied when played. Due to the body shape and hand positions it is not played like a traditional mbira, instead it is played with one hand, and it seems that striking the tines grants greater control than plucking.

The membrane potentiometers were chosen because they can be controlled with a single finger, allowing the user to manipulate them while playing the tines at the same time. In the context of the Third Room the membrane potentiometers are used to control the amplitude envelope of melodic and sequenced pitches in real time. The Ember also includes an analog joystick designed as a utility device, giving the user a kind of "mouse" on board the instrument. However, this simple sensor is also quite useful when used to add expression. The joystick features a button as well, which can be used to change functionality of the joystick itself, or used to cycle through the different harmonic series.

The combination of piezos and membrane potentiometers proved to be quite difficult to work with. Specifically it was necessary to troubleshoot a lot of "bleed" across the sensors. We found that the membrane potentiometer would create false readings across the piezos, and sometimes vice versa. We were able to control this contaminating data by a combination of resistors and diodes on the piezos, as well as a bit of Arduino trickery. As we saw the problem to be that the analog to digital buffer was not clearing itself out between analog pin reads, we decided to hook up empty pins to ground and read them in between each sensor. This acted as a sort of flush that rid the circuit of any leftover voltage and solved our data line problems very effectively. A custom PCB Arduino shield was developed to further improve the quality of the sensor data.

While the Ember is a fully functional and customizable MIDI instrument, in the context of the Third Room, it serves a more specific function. When a user plays the Ember, the notes are sounded and a new, pitched ball object is added to the virtual room. This object can then be interacted with through the virtual interface. The melodies played on the Ember are deconstructed visually and remixed by interaction inside the virtual space. Essentially, whoever plays the Ember controls the palette of the composition, but does not control the composition itself.

Below (Figure 18), the final version of the Ember can be seen. It was made from a higher density foam than the second prototype to make it more durable. It was CNC'd in two pieces and made hollow to allow for th installation of the Arduino and PCB shield. The two halves are held together by a series of neodymium magnets that were installed around the edge of the body. The membrane potentiometers were inset into the body by carving out a space on top and then their shapes were modeled into the form by using putty. The body was then painted with automotive paint and a clear coat the further protect it. The bridges were 3D printed and the tines ended being a cut from a softer style of plastic. This physical render was overseen by, and could not have been achieved without the skill and talent of, Andrew Capeluto.

Figure 18: Ember final version

3.4.2 The Fluid: Wireless Sonic Hand Sculpture

The other type of physical object is a handheld, wireless interface of sorts, called the Fluid that the user interacts with simply by holding it. Inside and out, there are a collection of sensors that capture data on how the user is holding it, and how they are handling it. In essence, it is a hand sculpture whose aesthetic quality is understood through touching. At the same time, the Fluid itself senses the user, how they are holding it, which direction their pointing it in, and whether or not they like the feel of the metal parts. This unique sensory experience is not limited to touch, as it is extended into the Third Room, effecting visual and aural aspects of the soundscape. The Fluid too was a product of collaboration between the author and Andrew Capeluto, whose previous work inspired the author to propose a musical interface in the form of a hand sculpture.

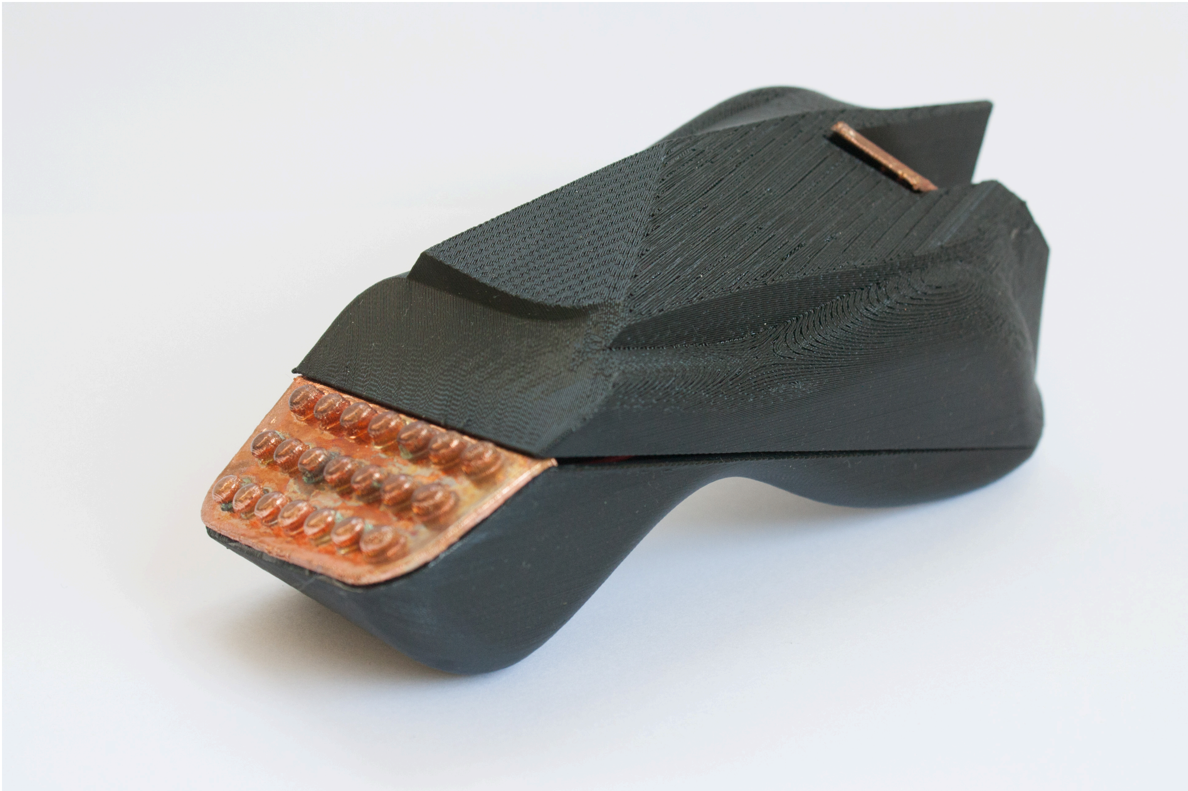

Figure 19: Fluid - Version 2, with copper plated touch points

The most important aspects of the physical designs were as follows: that it be handheld and wireless; that the shape promote curiosity in that the user will explore different methods of holding it; and that the surface has various points where touch sensing could be achieved. With these considerations the Fluid was designed so that the user first explores it in a tactile way, learning through experimentation the sonic and visual reactions to the physical interaction.

The digital design, lead by Capeluto, was achieved in Solid Works creating a basic handheld shape and then warping it into something that would be comfortable to hold in more than one way. Since the object itself is abstract, its shape should not put the user at a disadvantage, meaning that it should not intimidate them from interacting with it. By designing a shape that is ergonomic and organic the interaction becomes two-fold: the user interacting with the object itself and the user then interacting with the sound through the object.

When finished, the model had to be prepared to be 3D printed. This meant hollowing out the model, while maintaining the structural integrity, cutting it in half, and designing a method of opening and closing it that was inherent to its design. Two extra pieces were 3D printed and electroplated for capacitive touch sensing.

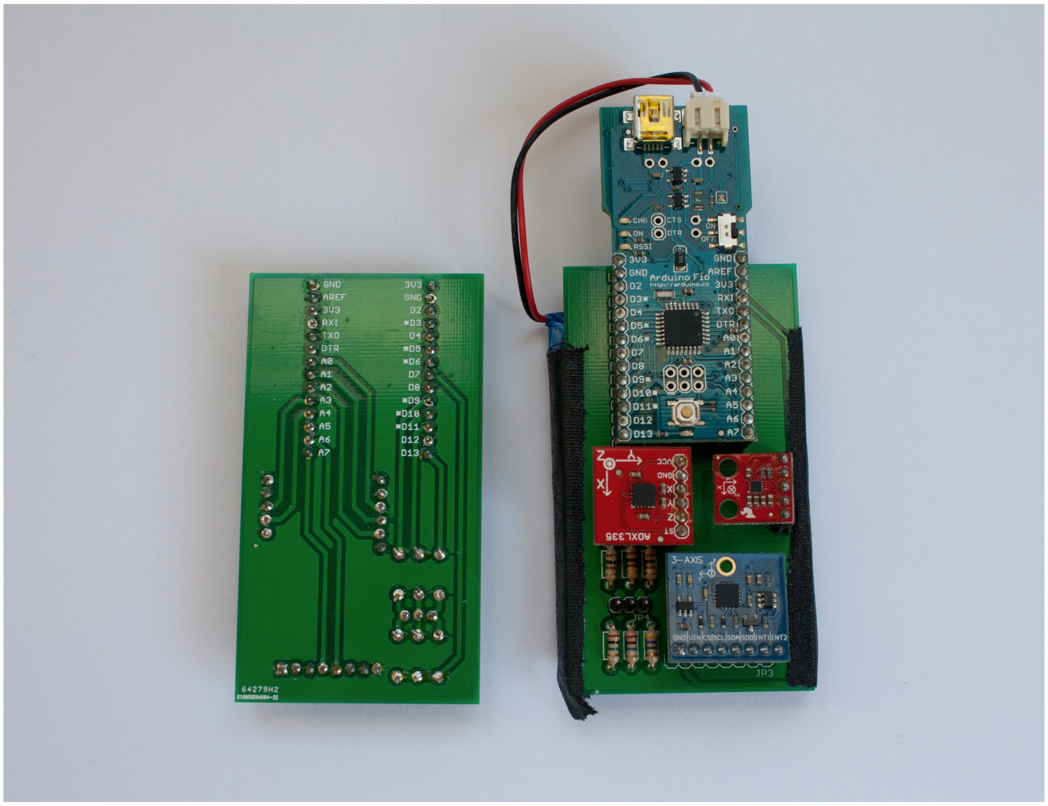

Figure 20: Custom Arduino Fio PCB Shield for Fluid

At the core of the object are an Arduino[21] Fio, an Xbee[22] WiFi chip, and a lithium battery. A custom shield PCB was designed and printed for the Fio allowing for an accelerometer, gyroscope, magnetometer, and three touch capacitance sensors (Figure 20). Two touch points, the electroplated pieces, are for absolute touch, while a third hidden sensor detects a user's proximity across a wider range. This way a user can be detected when they first go to reach for the object before they even touch it. The three positioning chips work together to return nine axes of information, which can be used to sense how the object is being held, and how the user is interacting with it. By detecting how the user is handling the Fluid, specific interactions can be encouraged through audio or visual feedback.

The Fluid has a different kind of authority over the audio and visual environment of the Third Room. Depending on how the user holds the object different interactions can occur. Hidden objects and instruments can be revealed through specific gestures. For instance, if held upright and shaken, a shaker style instrument can be accessed. However, touching the

[21] http://arduino.cc
[22] http://www.digi.com/xbee/

electroplated pieces can override this gesture, giving the user access to a different set of instruments, including an arpeggiator. Quick rotational movements across different axes can affect the orientation of the virtual environment itself or can be used to achieve score following. The combination of touch points, movement, and gesture allows for an incredible amount of depth and possibility for such a simple interface. The intention is that its use demands curiosity on the user's part to discover what it is capable of doing through tactile exploration.

3.5 Extending virtual interaction with mechatronic sound-objects

After completing the first iteration of this project it was displayed in the main classroom of the Music Technology program at the California Institute of the Arts. Commonly known as the Machine Lab, it is home to the collection of mechatronic instruments and sound objects that make up "The Machine Orchestra" (Kapur, Darling, et al. 2011). These instruments hang from the ceiling around the perimeter of the room and create an interesting spatial experience when played in this configuration.

In 2013, an array of solenoids was installed into the grid of the ceiling, so, when this project was set up as an installation it was modified to send messages to the mechatronic instruments and the new solenoids in the ceiling as well. As the ball objects bounced around the virtual environment they were programmed to trigger the instruments as if they were bouncing off the instruments themselves. When they bounced off the walls inside the program they triggered the nearest corresponding instrument, and in the case of the ceiling solenoids, they triggered the nearest single solenoid. The result is an immersive and spatial experience that further merges the physical and virtual rooms, creating a unique augmented space where the virtual is able to reach back into the physical world and activate it. This experience was so interesting and pleasing that it lead to the development of a new kind of mechatronic sound-object, designed to be hung from the ceiling and integrated into the physical environment. This new sound-object, designed by Ajay Kapur and Christopher Pheiffer, is now known, collectively, as the *Modulets*, which are 20 modular solenoid units that can be arranged in a 4 by 5 grid as shown in Figure 21. These solenoids are encased in a 1-inch square tubing with a blue LED that lights up when triggered. The *Modulets* are portable and have traveled with the Third Room to bring the mechatronic 3D audio experience to other venues beyond the CalArts Machine Lab.

Figure 21: Modulets - a mechatronic grid sound object

3.6 Future Work

The next steps in the development of the Third Room as a paradigm for the augmented musical space includes exploring other kinds of spaces. Of most interest to the author is the studio space. Not so much the traditional music studio, but more so the electronic studio, where composers and producers write and create music. In this type of setting there is usually a static arrangement of hardware and instruments set up and ready for use at any moment. This kind of arrangement itself transforms the space into an aggregate instrument for musical creativity. More and more the center of such a space is a computer that acts as a sequencer, recording device, instrument, and composition tool. Having a computer at the heart of this musical space grants the possibility of augmenting the space to include all manner of Next Music practices. Sensors and custom software could be created to allow the manipulation of hardware and software instruments; perhaps even creating a workflow that cuts down the time between thought and execution.

3.7 Closing Thoughts

The Third Room was the first large project conceived and developed by the author. Its initial version was completed as part of his final project towards a BFA in Music Technology at the California Institute of the Arts. It was then redesigned and rebuilt during his first year as a Master's student with the hopes of submitting it to the NIME conference as both a paper for publication and as an installation. It was both published and featured as an installation at the 2013 International Conference on New Interfaces for Musical Expression in Daejong, South Korea, held at KAIST University. Furthermore it has been displayed at the Newhall Sound Slam and the 2013 CalArts Digital Arts Expo.

The development of the Third Room is when Next Music first began as an idea. In the published paper "The Third Room: A 3D Virtual Music Framework" many ideas that would eventually become part of Next Music began to take shape, although they admittedly lacked focus when compared to this document. Perhaps it was that the implications of those ideas were too big and lay outside the context of the NIME paper, but it is more likely that they were still new ideas that needed further development and thought to be articulated correctly. More than that, the idea just needed to be applied to different kinds of spaces. While the Third Room created an augmented space, it was conceived as a virtual space, despite the fact that the main controlling aspect of the piece required physical space and human body input. The idea of augmentation was present but it wasn't the original focus of the piece. Augmented reality was not an entirely new idea but it seemed to be a narrow definition, limiting the kind interaction and device used to achieve (most often location sensors running on a mobile device), which didn't quite fit the idea of the original project. It was not until the author approached new spaces, both larger and smaller, that the common thread of augmenting physical space for the purposes of musical interaction that Next Music emerged.

The idea of augmenting a physical space for musical interaction is not entirely new or novel. However, the specific act of augmenting a musical space to create a new kind of space with its own set of practices is new. The Third Room attempts to begin this practice by way of example. The system, interaction, and musical output are all simplified versions of what will follow if others take up the practice. As technology continues to improve, the possibilities of what the augmented musical space can be grows. What shape or form these practices take is yet to be determined. It is very likely that many artists and musicians are attempting to conceive and create these new spaces and practices even now, and as time passes these works will be seen as the emergent musical practice of a future where computing and sensing are integrated into all aspects of life. Next Music is just one name for this new practice but it seeks to be an inclusive term. The Third Room is but one example of one aspect of this new practice, and a simple one at that.

The focus of Next Music in the musical space changes according to the context it is used. In a performance space, Next Music allows for a customizable "rig" of performance instruments and effects, allowing a performer to control, not only of electronic instruments, but environments as well (lighting, practical effects, scenic, etc.). In a pedagogical sense, Next Music allows for an enhanced practice room that could be used to automatically track the breadth and quality of a practice session, presenting an analysis to both student and teacher, and even make recommendations based on problem areas and report progress made in others. In the studio, compositional systems can act as overlays on the physical set up the studio environment. Augmenting these various musical spaces gives rise to Next Musical practices that are not possible elsewhere.

There are many applications that range from the practical to the experimental that can be explored freely once the systematic augmentation is complete. While this is entirely possible with current technology, it still requires many technical skills that most musicians do not have. The further development of frameworks and systems is necessary to bring these new practices to the average musician, composer, or producer. Until then, Next Music will more likely exist primarily on the next presented geographical scale, the global.

Chapter 4
Hear In LA: Augmenting the Acoustic Landscape

Los Angeles is a car driving culture. According to the United States Census Bureau and the California Department of Motor Vehicles, there are approximately 9.9 million people living in Los Angeles County[23] and there were over 6 million registered autos in 2013.[24] According to the 2008-2012 American Community Survey 5-Year Estimates table on "Commuting Characteristics by Sex," of the 4.3 million workers in LA County, almost 2.5 million of them commute in a car, by themselves, for over 20 minutes to and from work.[25] As Los Angeles drivers travel through the city, many of them in separate cars, they are essentially taking their private space with them wherever they go. Private space is a concept built into Los Angeles' geography. When presented with so much empty land the original developers created a "California Dream" where everyone has a front and a back yard, a place where you need not live in the over-cramped and vertical style of New York City. The caveat of having so much space is that residents need to travel from one end to the other all the time. Everything is farther apart than can be comfortably walked. While public transportation in LA was once robust enough to deliver its residents wherever they needed to go, the advent of the freeway and the closure of most of the public transportation lines, the car rose to dominance in Los Angeles. Over the years driving through LA traffic has become a trying experience, stressful and even dangerous; many experienced drivers find it necessary to counteract this stress with auditory entertainment. This desire to be entertained while driving intentionally disconnects commuters from the stress of driving, but has the unintended effect of disconnecting them from the landscape they are travelling through.

[23] http://factfinder2.census.gov
[24] http://apps.dmv.ca.gov/about/profile/est_fees_pd_by_county.pdf
[25] http://factfinder2.census.gov/faces/tableservices/jsf/pages/productview.xhtml?pid=ACS_12_5YR_S0801

Whether they know it or not, drivers readily ignore the environments surrounding their daily commutes, sometimes by choice, but also by design. By choice, they seek light distraction in the form of music, radio, and audiobooks, and by design, they are physically separated from the landscape by big walls that tend to line most of the LA freeways. Hear In LA is an effort to reconnect the mass of moving people that make up the ebb and flow of LA traffic with the invisible areas they travel through, often on a daily basis. A reconnection, not just with the physical landscape, but the diverse cultural landscape as well. Both in a historical and contemporary context the cultural makeup of Los Angeles has largely been defined geographically by neighborhood, lending each area its own ethnic identity. A multi-faceted identity made up of the close juxtaposition of different ethnic groups and cultures over long periods of time. This cultural cross-pollination is what gives Los Angeles neighborhoods their unique flavor and appeal. But for many, these neighborhoods are passed through rather than explored. Hear In LA promotes discovery by introducing users to the sounds of Los Angeles neighborhoods and encourages them to explore the spaces they normally ignore. The piece takes the form of a smartphone app that uses location data to create an interactive sound installation.

In this chapter, the project Hear In LA will be presented and discussed within the context of location-aware technology, site-specific art installations, field recording, locative media, community driven composition, and of course, Next Music. First, an overview of the project is given (pg. 62). Second, a brief background on location-aware and locative media projects is presented (pg. 64). Third, the compositional choice to use field recordings is explained (pg. 66). Fourth, the idea of crowd-sourcing as a means of collaborative and community composition is discussed (pg. 68). Fifth, the details of implementation are given (pg. 69). Sixth, the future of Hear In LA is discussed, including plans for expansion (pg. 71). And lastly, the author's closing thoughts on the piece will be presented (pg. 72).

4.1 Overview

Hear In LA is an interactive, location-aware, site-specific, and mobile device based Next Music sound installation. It is interactive in two ways, one passive and one active. Users can choose to just listen, or they can take part by contributing field recordings made through the app itself, adding them to the sonic landscape of Hear In LA. It is location-aware through the use of the location services and sensors in the mobile devices themselves. And it is site-specific in that the sounds that make up Hear In LA are directly connected to the location they were made. All of this is packaged into a smartphone app that will be available for both iOS and Android devices. Conceptually, Hear In LA touches upon ideas about transportation and culture, encouraging the user to become more acquainted with their surroundings. In the scope of Next Music, this project investigates how a sound installation can occur on a global scale and be driven by a community rather than any one artist or composer.

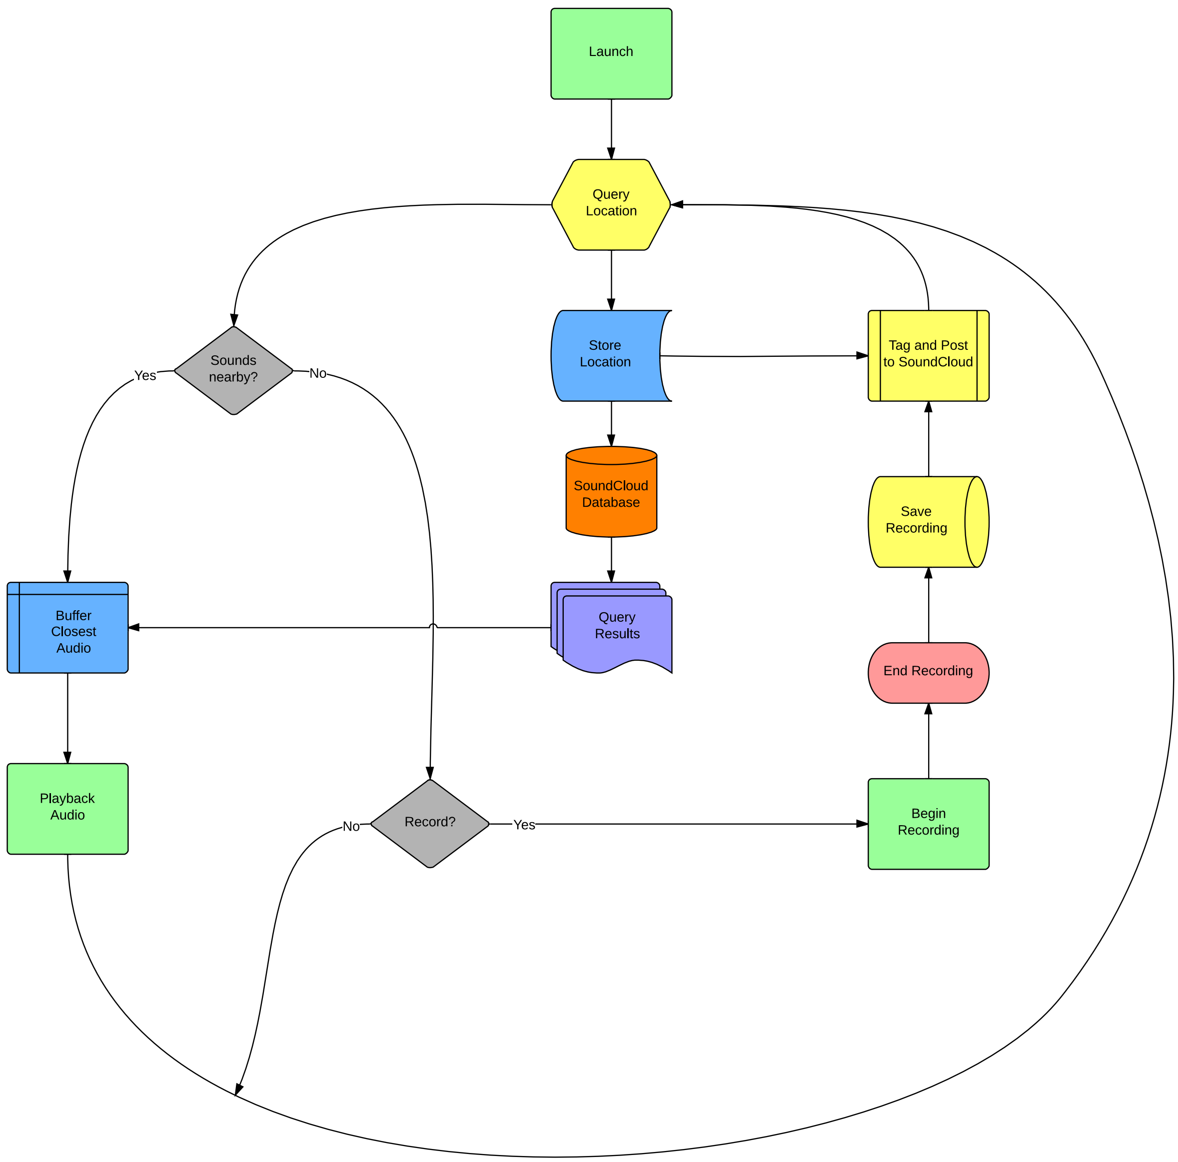

Figure 22: Hear In LA - Application Flow Chart

4.2 Background and Motivations

First and foremost, this project owes a great deal to a work by the designer Jessica De Jesus called *Hear in the North End*. This project explored a means of creating a new kind of experiential map that used field recordings to impart the local flavor of some famous and popular spots in Boston's "North End" neighborhood. The project itself was a flash-based interactive map that would be accessed via a personal computer and online. Hear In LA is a direct offshoot of this previous work and was carried out with the express permission, blessing, and design help of De Jesus herself.

Hear In LA has its roots in many experimental practices that combine art and location, and more specifically sound and location. Often classified as locative media, this kind of merging of space, experience, and sound, has been around for many years. More recently, with the rise of mobile devices that allow for real time video, sound, and location processes, this overlaying of reality with technology has become known as *augmented reality* (AR). Augmented reality exists somewhere between the real world and a virtual reality. Essentially, AR imposes virtual elements over real ones creating a mixed sensory experience. Augmented reality has many applications, some practical, others less so, and has had many representations in art, some even occurring before the term was coined. In recent years, a lot of this experimenting has occurred on mobile device platforms, focusing on the global positioning system (GPS) receivers inside these devices. But before GPS and smartphones there are examples of artists utilizing large spaces and various technologies to create locative media artworks.

Janet Cardiff has been augmenting reality since the early 1990s. Her numerous "Audio Walks" and "Video Walks" attempt to merge the physical world with an imagined narrative created by the artist (Cardiff et al. 2005). Using CD players, then iPods, and now smart phones, the user listens as Janet gives instructions on what to do and where to go. There are also prerecorded sounds which have been recorded on location which she says is very important because "the virtual recorded soundscape has to mimic the real physical one in order to create a new world as a seamless combination of the two."[26]

This seamless combination is what makes these early works "augmented reality" even though the term was not in use when she was doing this. Cardiff does not simply guide the user through the physical world; she is simultaneously the creator and the narrator who is describing

[26] "Introduction to the audio walks" - http://www.cardiffmiller.com/artworks/walks/index.html

the world they have entered. The narrative thread of her pieces is what drives the user to continue through the whole experience.

The audio walk is a popular medium for sound artists to explore the effects of location and sound. Also known as Soundwalking (Basco 2011; Drever 2009; Productions and Kirshenblatt-Gimblett 2014), there has been a lot of analysis of how these kinds of works address issues of space theory, place-making, cultural geography, and even archaeology (Adams et al. 2008; Butler 2007; Cornelio and Ardévol 2011). Most of the work done with the audio walk platform is pre-rendered field recording material created by the artist and the interaction portion of the piece consists of following the instructions or a route of some kind.

Location-aware music is a newer concept wherein a set of compositions can only be experienced if a user travels to a specific location and uses a purpose-built application. This kind of experience is dependent upon GPS location services to track a user through a large open space, revealing different tracks that have been composed specifically for each location they travel to and through. The group known as Bluebrain is well known for their location-aware music, and have released several "albums" that can only be experienced by travelling to The National Mall, Central Park, or Austin, Texas. While credited as "pioneers" by Wired magazine ("App-Album Pioneers Bluebrain Return With The Violet Crown for SXSW | Underwire" 2014) their work derives from an already established practice of exploring space through sound. While they may be the first widely known example of location-aware music the work of Cardiff and others preceded this work by several years.

Another important locative media piece is the ongoing piece most recently called *Local Report* by Robert Whitman. This piece has been reproduced several times since it's first performance in 1972 when it was called *NEWS*.[27] In its first inception people were invited to call in to the radio station where it was being performed and give a report of things going on around them. An interesting process where information is transmitted from a dispersed area, routed through a central collection facility (the radio station), and then broadcast back out into the space that generated the reports. The piece captured "the nature of a specific place and time, revealing it to the audience and participants alike."[28] This piece has been given new relevance in the age of information and the advent of the cell phone. Reimagining the piece alongside new technology, Whitman's current version of the piece, *Local Report*, invites participants to create

[27] http://www.whitmanlocalreport.net/sub_project.htm
[28] Ibid.

video and voice messages from all around the world. As the scale of the piece changed, so did the nature of the reflection it created. On a global scale it is not the nature of the global space that is revealed, instead it becomes a reflection on culture and time, as each participant offers a glimpse of the world from their perspective. The piece usually occurs in two parts, a performance, and then a compilation of the performance created after the fact. Whitman's *Local Report* bears distinct similarities to Hear In LA, both in its practice of place making and its attempt to create a candid cultural narrative through the use of open community participation. While Hear In LA focuses on the actual sounds of these places Whitman's piece always focuses on the human description and video representation.

Perhaps most relevant to the project Hear In LA is a piece called *Tactical Sound Garden* by Mark Shepard. Both a work of sonic art as well as an open-source toolkit the Tactical Sound Garden (TSG) is a community driven piece where people can "plant" sounds at specific locations using a WiFi enabled mobile device. Shepard states that this project "draws on the culture of urban community gardening to posit a participatory environment where new spatial practices for social interaction within technologically mediated environments can be explored and evaluated" (Shepard 2006). Utilizing the most advanced mobile device at the time, the PDA, Shepard used the expanding network of WiFi access points in New York City to facilitate the growing and tending of these sound gardens. Ideas of public and private space are addressed, as well as the impact of personal audio devices on these notions of overlapping space within an urban environment. Perhaps most interesting in this piece is the ability to "prune" a sound, thus opening the piece up to a sort of interactive dialogue between users. Together the community can work together to change the acoustic landscape of their city and at the same time change the way they experience travelling through it. This kind of participatory art is taking the lead of the modern art movements of cybernetics, telematics, and media archaeology, especially in its consideration and use of the community rather than focusing on individual user experience.

4.3 Field Recordings versus Compositions

The intention of Hear In LA is to focus on a collection of field recordings as a means of reconnecting the user with their physical and cultural surroundings. This occurs on multiple levels as the user chooses to simply listen, collect, or do both. Ideally, the act of listening inspires them to contribute, and the act of contributing pushes them to listen. However, a connection occurs for both the passive and active user. Not everyone who listens will choose to contribute, but they still get a sense of the invisible spaces around them. Those who collect will undoubtedly also listen, partially with the hope of hearing a sound they themselves have recorded, but also to hear what others have found in the areas they frequent.

Hear In LA could very easily have become a personal statement of the artist's own musical practices. Using the locative-media model to create a citywide composition for others to explore while only relinquishing the direction and order of sounds to the user. This model has been explored previously and provides a different experience to both the user and the artist. The choice to focus on field recordings was difficult as it forces the artist to give up their own personal desires and experience what the piece becomes just as the other users do. The artist can participate but exhibits almost no direct control over what the piece becomes as they hand it over to the public. By relinquishing control the artist allows for several things to occur.

Conceptually, this piece is about the cultural landscape of Los Angeles and addresses the divide between the average citizen and the sprawling metropolitan area that covers some 468.67 square miles.[29] The sheer size of it makes Los Angeles difficult to define, understand, and navigate. Hear In LA accepts that many simply do not have the time or perhaps even the opportunity to really explore the city and seeks to allow them a glimpse inside the many neighborhoods and communities that they pass through. If the audio component of this piece was made up of composed music the user might end up "exploring" these locations but really they would be experiencing a different kind of isolation, where in fact they were exploring the music, focusing on their singular experience. Instead, a piece that focuses entirely on field recordings is inherently about the land, the people and the community. Reconnecting a people with the landscape that, in many cases, has become but an obstacle in the way of their daily lives. When an "East-Sider" is confronted with the need to go to the "West Side" it is often received with a grunt and a sigh, as anyone who has tried to traverse Los Angeles from East to West, or vice versa, knows all too well that the endeavor could take "forever." They see their origin and they see their destination, which might be but a few miles apart, but the trip can take hours depending on traffic and the savvy navigation skills of the traveller. Perhaps too few of the residents of Los Angeles do not see these commutes for what they are, journeys. If their own perception becomes about donning an explorer's cap, to see what they can find in the space that lies in between their origin and destination, the chore of getting from one side of town to the

[29] http://quickfacts.census.gov/qfd/states/06/0644000.html

other becomes an engaging and fulfilling experience in its own right. While many commuters and daily travellers do not have the luxury of wandering off the beaten path, Hear In LA will allow them to sonically explore these “in between” spaces without requiring them to veer from the direct course to their destination.

On a more practical level, the choice of focusing on field recording and community engagement, works in the favor of both the artist and the user. To create a curated sonic experience that traverses the complex cultural geography of Los Angeles assumes an intimate knowledge of those spaces and any composition created for them implies a statement about those spaces. To attempt to know, understand, and process the complexity of the Los Angeles cultural landscape would take a lifetime, and by the time one “completed” the task, so much would have already changed that it would negate all previous understanding. It could be, that the artist decides not to focus on the cultural landscape, and instead composes a creative interpretation of their experiences and perceptions. This kind of piece, in the opinion of the artist, lacks impact and becomes little else than a novel kind of musical composition. Even if the artist embraces this novelty and exploits the technology to create an interesting and moving experience, at some point, the piece begs the question of how it reflects one person’s view of a space occupied by millions.

By creating a platform for an entire population to openly and actively create together allows for as many points of view as there are participants. Each resident has a unique experience and therefore a unique perception of their own surroundings, and this piece allows them all the same access to creatively contribute to an ever changing and expanding composition created for and by the city itself.

4.4 Crowd Sourced Collaboration and Composition

By handing over the compositional efforts of Hear In LA to the residents of Los Angeles it combines two unique practices from different fields into a single and exciting compositional practice. “Crowd Sourcing” is the practice employing large digital and physical communities to help break up the required work into many small parts. This practice takes many forms, from the more practical crowd funding platforms such as Kick Starter and Indie Go-Go, to the amazing collaborative efforts of such artistic works as the Johnnie Cash Project,[30] where a

[30] http://www.thejohnnycashproject.com

music video was created frame by frame by individuals contributing unique artistic renderings of video frames or their own original artwork. By splitting up the work amongst so many people it allows for possibilities that one artist could not conceive nor execute on such a scale. In the case of Hear In LA it would take one artist many weeks and months of working all day to find and record enough sounds to fill the geographic landscape with enough content to make the installation a worthwhile experience, and even then the limitations of their perspective and local knowledge would be a major flaw in the intended sonic landscape. By opening up participation and creative license to a community of people the artist creates something that will take on a life of its own. It becomes a new space, one that is created by the community through active participation, one that they can help shape and explore. If successful, this piece will sound a little different each time it is experienced, which qualifies the piece itself as a unique kind of generative composition. The randomness of life itself is the probability engine driving it, as the app proliferates across the population, more space becomes filled with sound and the additive effect of sounds layered on sounds will create a slowly changing composition that resounds with the activity of daily life in Los Angeles.

4.5 Implementation Strategies

In its initial form, Hear In LA was an Android based app that used Pure Data (Puckette 1996), through the library called libpd (Brinkmann et al. 2011), as its audio engine. While this made for convenient programming and helped develop the proof of concept, there were many undesirable aspects of it. For one, it required that the audio samples already exist on the device so the ability to actively participate through recording was not yet a possibility. Also, libpd required the use of WAV files that use a lot of memory. This made for a very large app that consumed far too much memory, especially considering that the initial testing only surrounded the artist's commute from Los Angeles to CalArts in Valencia and consisted of about a dozen samples. However, the success of being able to trigger sounds based on location was the main focus and the driving force behind the development of the next phase.

The second phase of development had several goals. First, the app should be cross platform for the two major market shares of smartphone operating systems, namely iOS and Android. To help facilitate this, the framework known as JUCE[31] was used as it allows cross-

[31] http://www.juce.com

platform compilation with very little extra effort. The benefits and issues of this choice will be discussed below. Second, the app needed to be small and fast so as not to consume too much memory that could deter users from participating. This meant that the recordings needed to be hosted online and only temporarily kept on the users device. The possibility of this app becoming widely used meant that there is potential for thousands of files that would have to be stored in a database for easy retrieval. This kind of database and hosting requires quite a bit of capital and time to maintain and is well out of the scope of the artists own financial and time restraints. Therefore, the decision was made to use the free audio hosting service called SoundCloud[32] to help host and maintain the audio that will make up Hear In LA. Again, the benefits and difficulties of this decision will be outlined below. Lastly, the second version needed to allow for recording and tagging of audio at a specific location so that it could be automatically entered into the Hear In LA framework.

4.5.1 JUCE and Developing a Cross Platform Application

JUCE is a C++ development framework that has become a digital audio industry standard for the development of applications on almost all platforms, mobile and otherwise. By removing the troublesome aspects of platform specific processes and coding languages, it allows developers and companies the ability to focus on the application itself rather than the platform specific issues that always arise. One of the most difficult aspects of app development can come from audio needs such as recording, playback, and synthesis. This made the decision to use the JUCE framework very easy, however this revealed itself to be more difficult than expected.

One of the major issues that occurred was that JUCE does not have location services built into it. This means that even though the audio framework needs no alteration to work across different platforms, the essential location services need to be customized on a per platform basis. This means that native development languages need to be implemented within a C++ framework. While entirely possible to communicate between languages it presents unique challenges for each language used, in this case Objective-C (iOS) and Java (Android). However, once these individual issues are addressed, the other issues that occurred came from learning to use the framework itself and interfacing with the SoundCloud API.

Although JUCE is a C++ framework it has been developed so thoroughly that it has almost become its own language. The framework requires an understanding of hierarchy and an

[32] www.soundcloud.com

intimate knowledge of the framework's objects and methods to correctly implement. Despite these issues, JUCE usually alleviates more issues than it presents.

4.5.2 Interfacing with the SoundCloud API

From the start of development of the second version, SoundCloud seemed like it would be an ideal platform to handle the audio hosting. The fact that it has a large user group, free access, and a database maintained by professional developers, all seemed like they would work in favor of the development of a full-featured Hear In LA. Unfortunately, SoundCloud's API (application programming interface) is more focused on web development than mobile devices. What this meant for the development of the full-featured application was that the process was slow and frustrating, with results that are still being tested. In the future it is quite possible that this project may require the direct help and intervention of the SoundCloud developers themselves, or the project may need to be moved to its own dedicated server and database.

4.6 Future Work and Expansion

The future of this project depends highly on its reception by the public. The hope of the artist is that this project reaches a point of popularity and use that a major overhaul would be required to keep it running. As expressed above, this could come in the form of a helping hand from the SoundCloud company directly, another large entity with the network and developers to help facilitate its growth, or if need be, a revamped system built from the ground up by the artist. Even though this would be a daunting task for a single artist/developer, it would also be a welcome challenge, as it would mean that the project is successful and well received by the public.

These technical requirements are a concern for the long term but there are still other planned developments for expansion in the future. First of all, although this project is titled Hear In LA, there is nothing in the app that actually limits its use to the city of Los Angeles or even the greater Los Angeles metropolitan area. In fact in its current form it could already be used on a global scale. The decision to focus on Los Angeles is because of the direct relevance to the commuter culture wherein a large portion of the population finds itself disconnected from each other and the city itself as well as the artist's own residence within the city of Los Angeles and his own practice of commuting to and from Valencia, some 30 miles away. Plans to create a *Hear In* brand are already in development between the artist and the original designer Jessica De

Jesus. Rather than limit the work to any one place in particular, the hope is to create a truly global interactive sound installation that pushes the limits of social musical interaction and creates a paradigm for global participation in a single musical act.

Other development plans include several extensions of the main application that would allow for customizable modes to allow users to use the app as a framework for their own designs and projects. Current ideas for new modes include allowing users to create custom curated playlists, either of the Hear In LA content or their own original work. This would allow for the development of location-aware albums for everyone who is interested. Instead of having to create their own app from scratch they could use the location services and musical framework of Hear In LA to implement their own ideas and experiments quickly and easily. The curating function could be used in many ways, to highlight a single super-user, or collection of users, or perhaps to create location-based playlists for road trips or even curated explorations of the city and surrounding areas. With enough users and interest this kind of framework could be used to develop many ideas unimagined by the artist.

Other ideas for expansion include guided navigation to a selected sound's source and a sort of "time-machine" mode activated by locations with the largest amount of sounds collected in the smallest area. In this time-machine, a user who enters this area gets access to a series of sounds each predating the previous example, wherein the user experiences the sonic history of the space as a collage of pieces that may or may not change drastically the further back they go. This kind of mode would create a kind of unprecedented archival collection of real sounds from real places over what could be a long period of time. This mode, even more than the others presented, is a "blue sky" development that not only depends on active use of the application by many people, but by many people over a long period of time.

4.7 Closing Thoughts

Hear In LA came out of an understanding that the artist's own tendencies unconsciously revolved around the concept of space. This simple observation made by a professor was really the impetus for what has become Next Music. Of all the projects presented in this document, Hear In LA fills the largest of spaces currently accessible by humans, the planet itself. While grandiose in sentiment, it is only slightly hyperbolic. The world will only continue to become smaller and as this happens communication technologies will continue to grow to fill in every space occupied by humans, and even some that are not. While the purpose of the piece would inevitably get lost in this kind of global expansion and ultimately resolve itself by actively disconnecting people with their environment, there is arguably something poetic in that thought. Truth be told, this is the least likely future for this project, but if run to its logical conclusion, without the burden of realistic probability, it could eventually devolve into people listening on headphones to the sounds that they could readily hear if they simply removed those headphones.

Hear In LA occupies an emerging space within the future of music. Digital space is not subject to the same limitations as physical space, though, on a side note, digital space requires quite a bit of physical space in the form of server farms. In a physical space, sound is subject to decay as it goes from an audible wave of air molecules traveling through the space to an inaudible wave, where perhaps air molecules are still moving with that sound's energy but on a scale so small nothing but the air would notice. Capturing that sound in a recording already creates a kind of sonic time capsule that defies the laws of acoustics in nature. The act of tying that sound, through a recording, to the space it was created in, transforms a finite acoustic space into an infinite digital space.

Next Music allows for a musical space that lives somewhere in between the finite acoustic space and the infinite digital space. Furthermore, it conceives of a musical practice created entirely out of and for this new space. How the musician, composer, and sound artist learn to work within such a conceptual space is yet to be seen. This document merely presents the possibility and hopes to inspire others to discover what music can become when created in and for the augmented space.

Chapter 5
Conclusion

5.1 Summary

This document presents the work created by Colin Honigman during the course of his MFA studies at the California Institute of the Arts. From the beginning these projects were individually conceived and completed without much consideration about their underlying connections. It was not until the compilation of this document that the idea of Next Music emerged as not only a conceptual connection between the projects but an aspirational idea about a new kind of musical practice. Through the projects Cultivating Frequencies, the Third Room, and Hear in LA, various aspects of Next Music were examined but in truth only a small portion of what Next Music is and could be was explored. As an exploration of spaces, this thesis was presented in order of geographical scale (Figure 23), moving from the smallest and most intimate space, a home garden, to a more moderate sized space, the musical space, all the way to the global space. The projects at each level display the use of a unique set of tools, or unique combination of similar tools, to help transform each space into an augmented space, or, in the case of Hear in LA, the tools to access the existing technologies creating the augmented space (i.e. mobile devices, wireless high speed data networks, and global positioning systems).

As society and culture continues to revel in the wonders of technology, information, and the network, the augmented space will also continue to grow in size and quality. The proliferation of the augmented space will undoubtedly affect all kinds of art making practices, not just music. However, music does have a special relationship with both time and space. In fact, it could be said that music is a consequence of this relationship, as what we hear is the vibration of air over time and through space. Over centuries humans have crafted ingenious ways of creating, transmitting, and listening to new sounds. The augmented space offers a unique opportunity to continue this exploration of sound and what is possible by conceiving of wholly new spaces that allow for an open-ended process of musical creativity. Next Music is the root of this idea, whether or not the name is ever used outside of this document, the notion of creating music in and for the augmented space is part of the future of music.

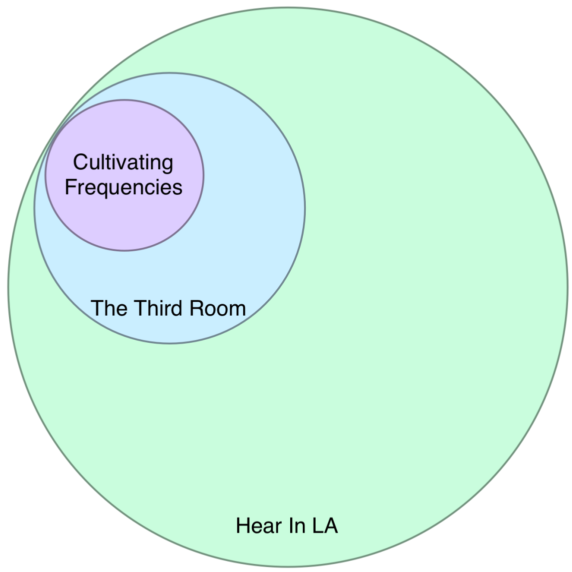

Figure 23: Geographic Scale of each project

5.2 Primary Contributions

Next Music, as stated above, is the aspirational contribution and main theme of this thesis. Providing a vocabulary and conceptual framework for future thought, development, and work, Next Music is an attempt to create a new movement, or perhaps only define an emerging movement and trend in contemporary musical practices. Outside this overarching theme and concept, the contributions of this thesis are exhibited by the precedence set by the projects themselves.

Most notably, Cultivating Frequencies saw the very real contribution to the open-source community of a new technique in touch-sensing technology. While not invented by the author, the first in-depth presentation and analysis of this new open-source technique is a unique contribution, especially when paired with the development and release of the new Arduino library *SweepingCapSense*. This contribution has already been validated by the project known as *Plant Sense*,[33] where the *SweepingCapSense* library was used to help create a plant based light installation where colored lights change upon interaction with the plants below them. The library

[33] http://www.iaacblog.com/maa2013-2014-physical-computing/2014/03/ardutechture-interactive-green-facade/

and an analysis was also published as part of the NIME 2014 proceedings under the title "Techniques in Swept Frequency Capacitive Sensing: An Open Source Approach," and was presented at the NIME conference held at Goldsmiths College in London 2014. As the author continues work on the Cultivating Frequencies project he hopes to see more developments, refinements, and outside uses of his open-source library.

In the case of the Third Room, the primary contribution is an introduction of musicians and artists to the emerging digitized reality and invitation to expand the possibilities of our musical spaces through computing and sensing. It is the author's opinion that the development of new digital instruments displays far too much reliance on the instruments and interfaces of old and requires radical new approaches to reach their full potential. Furthermore, the Third Room encourages the active removal of the wall that separates performer and audience, and the redefinition of the roles that each play within the context of musical performance and composition.

Lastly, Hear In LA seeks to provide a framework for large-scale location based music, as well as providing a new means for people to explore their cities through sound.

5.3 Final Thoughts

Next Music is an acknowledgment of the complexity that has arisen from centuries of diverging and converging, appearing and disappearing, musical practices. As stated above, Next Music was directly inspired by Next Nature and the idea that human creations can take on levels of complexity that render them as wild as nature. In music, for example, the proliferation of genres as distinct offshoots of previous aesthetic preferences is a product of continual contextualization and constant definition by both musicians and audiences alike. The creation of new musical styles is a communal practice that often includes the seemingly random appropriation of past styles mixed with contemporary technologies and trends. But, while music may qualify as an aspect of Next Nature, Next Music does not. Rather, Next Music is foremost concerned with the musical exploration of the augmented space.

The word exploration has been used a lot in the process of explaining Next Music and this is for good reason. The augmented space is emerging, through active creation and complicit use, it has already begun to influence our daily lives. The most appealing aspect of the augmented space is that everyone can seek to create their own custom versions. With the introduction of smart appliances and utilities, people are using the augmented space, on a very practical level, to automate their own lives. As ubiquitous computing becomes a reality, the augmented space will perhaps become one of the most frequently accessed spaces. Because of this, it is very important that the arts learn to use the augmented space as it affords new methodologies and practices that are impossible in other spaces. At this point in time, artists are presented with a wonderful opportunity to influence the future of the augmented space by showing the world what is possible within it. Through the practice of art and creativity, those who adopt this new space, will participate in its creation and definition. Truly exploring the depths of possibility and limited only by their own creativity. The potential held within the augmented space seems to immediately qualify it as a primarily creative space, however this qualification is dependent upon artists and creative individuals to create for and within this space.

Next Music is an open invitation to musicians and artists alike to participate in the making of this new place. It acts on the assumptions that ubiquitous computing and sensing are impending, necessary, and positive results of the technological and information revolutions of the 20th and 21st centuries. In itself, it does not seek to criticize this emergent space, but instead participate in its shaping. The augmented space itself is the artistic medium, but it is also more than that. And it is less. The augmented space is both the paint and the canvas, but it not strictly so. Similarly, it can be both the input and the output of a system. In its ideal, it is a type intangibility created through the combination of tangible objects. It is a space fraught with possibility and therefore makes it a simultaneously exciting and terrifying place to occupy. Next Music witnesses the beginning of a shift in musical and cultural practices, driven by computing and sensing technologies, that will change them both in unforeseeable ways. The possibilities, again, are vast, which makes it even more important to occupy this space as a means of protecting it. Humans have the tendency to seek control over unknown spaces, and this control is rarely for the benefit of the space itself. And so, it is rare to see the emergence of such a space and have the opportunity to protect it through occupation and exploration.

Bibliography

Adams, M. D., N. S. Bruce, W. J. Davies, Rebecca Cain, P. Jennings, Angus Carlyle, Peter Cusack, Ken Hume, and C. Plack. 2008. "Soundwalking as a Methodology for Understanding Soundscapes." http://usir.salford.ac.uk/2461.

Angel, Claudia Robles. 2011. "Creating Interactive Multimedia Works with Bio-Data." In *Proceedings of the International Conference on New Interfaces for Musical Expressions (NIME 2011)*. http://icemserv.folkwang-hochschule.de/~robles/papers_robles/Biointerfaces_CEIArtE_sm.pdf.

"App-Album Pioneers Bluebrain Return With The Violet Crown for SXSW | Underwire." 2014. *WIRED*. Accessed April 5. http://www.wired.com/2012/03/app-bluebrain-violet-crown/.

Ascott, Roy. 2002. "Behaviourist Art and the Cybernetic Vision." https://pl02.donau-uni.ac.at/jspui/handle/10002/615.

Ascott, Roy, and Edward A Shanken. 2003. *Telematic Embrace Visionary Theories of Art, Technology, and Consciousness*. Berkeley: University of California Press. http://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=90553.

Ashby, W. Ross. 1957. "An Introduction to Cybernetics." http://dspace.utalca.cl/handle/1950/6344.

Bagchee, Sandeep. 1998. *NAD: Understanding Raga Music*. Mumbai, India: Ceshwar Business Publications Inc.

Barker, Andrew. 2004. *Greek Musical Writings: Volume 2, Harmonic and Acoustic Theory*. Cambridge University Press.

Basco, Laura. 2011. "Soundwalking. A New Tool for Urban Regeneration." http://docu.iuav.it/77/1/My_Ideal_City___Basco,_Laura___86-91.pdf.

Berg, T., D. Chattopadhyay, M. Schedel, and T. Vallier. 2013. "Interactive Music: Human Motion Initiated Music Generation Using Skeletal Tracking By Kinect." Accessed February 5. http://w.tamaraberg.com/papers/kinect_music.pdf.

Berliner, Paul F. 1993. *The Soul of Mbira: Music and Traditions of the Shona People of Zimbabwe*. University of Chicago Press.

Bevilacqua, F., L. Naugle, and I. Valverde. 2001. "Virtual Dance and Music Environment Using Motion Capture." In *Proc. of the IEEE-Multimedia Technology And Applications Conference, Irvine CA*. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.92.4750&rep=rep1&type=pdf.

Borenstein, Greg. 2012. *Making Things See: 3D Vision with Kinect, Processing, Arduino, and MakerBot*. Make.

Brinkmann, Peter, Peter Kirn, Richard Lawler, Chris McCormick, Martin Roth, and Hans-Christoph Steiner. 2011. "Embedding Pure Data with Libpd." In *Proc Pure Data Convention 2011*. http://www.uni-weimar.de/medien/wiki/images/Embedding_Pure_Data_with_libpd.pdf.

Brown, Paul. 2008. "The Mechanisation of Art." http://pl02.donau-uni.ac.at/jspui/handle/10002/595.
Butler, Toby. 2007. "Memoryscape: How Audio Walks Can Deepen Our Sense of Place by Integrating Art, Oral History and Cultural Geography." *Geography Compass* 1 (3): 360–72.
Byrne, David. 2012. *How Music Works*. San Francisco Calif.: McSweeney's.
Cardiff, Janet, Mirjam Schaub, Thyssen-Bornemisza Art Contemporary, and Public Art Fund. 2005. *Janet Cardiff: The Walk Book*. Thyssen-Bornemisza Art Contemporary (www. tba21. org/).
Cook, J. S. A. F. M., T. S. M. F. R. Moore, A. K. A. Blake, and R. R. de Vasconcelos. 2013. "Real-Time Human Pose Recognition in Parts from Single Depth Images." Accessed February 5. http://www.eecs.ucf.edu/courses/cap4453/projs/cap4453fa11renan.pdf.
Cornelio, Gemma San, and Elisenda Ardévol. 2011. "Practices of Place-Making through Locative Media Artworks." *Communications* 36 (3). doi:10.1515/comm.2011.016.
Drever, John L. 2009. "Soundwalking: Aural Excursions into the Everyday." http://eprints.gold.ac.uk/7836/.
Dubost, Gilles, and Atau Tanaka. 2002. "A Wireless, Networkbased Biosensor Interface for Music." In *Proceedings of International Computer Music Conference, Goeteborg, ICMA*. http://www.csl.sony.fr/downloads/papers/2002/tanaka-02a.pdf.
Edes, Robert T. 1899. "Relation of the Chirping of the Tree Cricket (Oecanthus Niveus) to Temperature." *The American Naturalist* 33 (396): 935–38.
Franchi, Stefano, and Francesco Bianchini. 2011. *The Search for a Theory of Cognition: Early Mechanisms and New Ideas*. Amsterdam ; New York: Rodopi.
Frisson, C., S. Dupont, J. Leroy, A. Moinet, T. Ravet, X. Siebert, and T. Dutoit. 2013. "LoopJam: Turning the Dance Floor into a Collaborative Instrumental Map." Accessed February 5. http://www.eecs.umich.edu/nime2012/Proceedings/papers/260_Final_Manuscript.pdf.
Gelineck, Steven, Dan Overholt, Morten Büchert, and Jesper Andersen. 2014. "Towards an Interface for Music Mixing Based on Smart Tangibles and Multitouch." Accessed February 7. http://vbn.aau.dk/files/75884908/gelineck_NIME13.pdf.
Gies, Frances, and Joseph Gies. 1994. *Cathedral, Forge, and Waterwheel: Technology and Invention in the Middle Ages*. HarperCollins Publishers New York. http://www.getcited.org/pub/103082775.
Glinsky, Albert. 2000. *Theremin: Ether Music and Espionage*. University of Illinois Press.
Heinrichs, Christian, and Andrew McPherson. 2012. "A Hybrid Keyboard-Guitar Interface Using Capacitive Touch Sensing and Physical Modeling." https://qmro.qmul.ac.uk/jspui/handle/123456789/4241.
Hochenbaum, Jordan, and Ajay Kapur. 2011. "Adding Z-Depth and Pressure Expressivity to Tangible Tabletop Surfaces." In *Proceedings of the International Conference on New Interfaces for Musical Expression (Oslo, Norway, 2011), Jensenius AR, Tveit A., Godøy RI, Overholt D.,(Eds.)*, 240–43. http://www.karmetik.com/sites/default/files/publications/2011_nime_smartfid.pdf.
Hochenbaum, Jordan, and Owen Vallis. 2009. "Bricktable: A Musical Tangible Multi-Touch Interface." In *Proceedings of Berlin Open Conference '09. Berlin, Germany*. http://mtiid.calarts.edu/sites/mtiid.calarts.edu/files/publications/2009_berlinopen_brick.pdf.
Holmes, Aaron David. 2013. "The Evolution of Instrument Design and Its Influence on Trumpet Repertoire." http://drum.lib.umd.edu/handle/1903/14401.

Honigman, Colin, Andrew Walton, and Ajay Kapur. 2014. "The Third Room: A 3D Virtual Music Paradigm." Accessed March 30. http://mtiid.calarts.edu/sites/mtiid.calarts.edu/files/publications/92_Paper.pdf.

Jordà, Sergi, Günter Geiger, Marcos Alonso, and Martin Kaltenbrunner. 2007. "The reacTable: Exploring the Synergy between Live Music Performance and Tabletop Tangible Interfaces." In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction*, 139–46. ACM. http://dl.acm.org/citation.cfm?id=1226998.

Jordan, Ken, and Randall Packer. 2002. *Multimedia: From Wagner to Virtual Reality*. Expanded Edition edition. New York: W. W. Norton & Company.

Kapur, Ajay, Michael Darling, Dimitri Diakopoulos, Jim W. Murphy, Jordan Hochenbaum, Owen Vallis, and Curtis Bahn. 2011. "The Machine Orchestra: An Ensemble of Human Laptop Performers and Robotic Musical Instruments." *Computer Music Journal* 35 (4): 49–63. doi:10.1162/COMJ_a_00090.

Kapur, Ajay, Georg Essl, Philip Davidson, and Perry R. Cook. 2003. "The Electronic Tabla Controller." *Journal of New Music Research* 32 (4): 351–59.

Kapur, Ajay, Jordan Hochenbaum, Michael Darling, Dimitri Diakopoulos, and Jim Murphy. 2011. "The KarmetiK NotomotoN: A New Breed of Musical Robot for Teaching and Performance." In *Proceedings of the International Conference on New Interfaces for Musical Expression$}$*, 228–31. http://www.nime.org/proceedings/2011/nime2011_228.pdf.

Larry, B. 1997. *Capacitive Sensors Design and Applications*. Vol. 31. New York: IEEE Press. http://scholar.google.com/scholar?cluster=6268115552881827276&hl=en&oi=scholarr.

Manovich, Lev. 2010. *The Poetics of Augmented Space*. Springer. http://link.springer.com/chapter/10.1007/978-3-7091-0300-5_26.

McPherson, Andrew. 2012. "TouchKeys: Capacitive Multi-Touch Sensing on a Physical Keyboard." In *Proc. NIME*. http://www.eecs.umich.edu/nime2012/Proceedings/papers/195_Final_Manuscript.pdf.

Mensvoort, Koert van, and Hendrik Jan Grievink. 2012. *Next Nature: Nature Changes Along with Us*. Actar.

Merriam-Webster. 2005. *The Merriam-Webster Dictionary*. New edition. Springfield, Mass: Merriam Webster.

Murphy, DAC Jim, James McVay, Ajay Kapur, and Dale Carnegie. 2013. "Designing and Building Expressive Robotic Guitars." In *Proceedings of the 2013 Conference on New Interfaces for Musical Expression (NIME), Daejeon, Korea*. http://nime2013.kaist.ac.kr/program/papers/day4/paper12/36/36_Paper.pdf.

Murphy, Jim, Ajay Kapur, and Carl Burgin. 2010. "The Helio: A Study of Membrane Potentiometers and Long Force Sensing Resistors for Musical Interfaces." In *Proceedings of the International Conference on New Interfaces for Musical Expression*, 459–62. http://www.educ.dab.uts.edu.au/nime/PROCEEDINGS/papers/Demo%20Q1-Q15/P459_Murphy.pdf.

Nagashima, Yoichi. 2003. "Bio-Sensing Systems and Bio-Feedback Systems for Interactive Media Arts." In *Proceedings of the 2003 Conference on New Interfaces for Musical Expression*, 48–53. National University of Singapore. http://dl.acm.org/citation.cfm?id=1085726.

Niinimäki, Matti, and Koray Tahiroglu. 2012. "AHNE: A Novel Interface for Spatial Interaction." In *CHI'12 Extended Abstracts on Human Factors in Computing Systems*, 1031–34. ACM. http://dl.acm.org/citation.cfm?id=2212378.

Nikolaj Møbius. 2012. "Dzl's Evil Genius Lair: Arduino Do the Touché Dance." http://dzlsevilgeniuslair.blogspot.se/2012/05/arduino-do-touche-dance.html.

Overholt, Dan. 2005. "The Overtone Violin." In *Proceedings of the 2005 Conference on New Interfaces for Musical Expression*, 34–37. National University of Singapore. http://dl.acm.org/citation.cfm?id=1085951.

Poupyrev, Ivan, Philipp Schoessler, Jonas Loh, and Munehiko Sato. 2012. "Botanicus Interacticus: Interactive Plants Technology." In *ACM SIGGRAPH 2012 Emerging Technologies*, 4. http://dl.acm.org/citation.cfm?id=2343460.

Productions, Tourist, and Barbara Kirshenblatt-Gimblett. 2014. "Soundwalk in the City: Reclaiming Urban Reality." Accessed April 1. http://www.nyu.edu/classes/bkg/tourist/miyawaki.doc.

Puckette, Miller. 1996. "Pure Data: Another Integrated Computer Music Environment." *Proceedings of the Second Intercollege Computer Music Concerts*, 37–41.

Robert Sandall and Mark Russell Interview Brian Eno about Generative Music in 1996. 2013. http://www.youtube.com/watch?v=yJgAQ3xWMoA&feature=youtube_gdata_player.

Sato, Munehiko, Ivan Poupyrev, and Chris Harrison. 2012. "Touché: Enhancing Touch Interaction on Humans, Screens, Liquids, and Everyday Objects." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 483–92. ACM. http://dl.acm.org/citation.cfm?id=2207743.

Shanken, Edward A., In Bruce Clarke, and Linda Dalrymple Henderson. 2002. "Cybernetics and Art: Cultural Convergence in the 1960s." *From Energy to Information*. http://65.19.164.224/teaching/gen/shanken_CyberneticsArtCultConv.pdf.

Shepard, Mark. 2006. "Tactical Sound Garden [TSG] Toolkit." In *3rd International Workshop on Mobile Music Technology, Brighton, UK*. http://www.andinc.org/v2/writings/306090_09_shepard.pdf.

"Smartphone Users Around the World - Statistics and Facts [Infographic]." 2013. *Web Design Dubai - Dubai Web Design and Web Application Development Company*. Accessed July 23. http://www.go-gulf.com/blog/smartphone/.

Studie II; Nr. 3 [von] Elektronische Studien. 1956. [Wien: Universal Edition.

Taylor, Timothy D. 2001. *Strange Sounds: Music, Technology and Culture*. Routledge.

Tindale, Adam R., Ajay Kapur, George Tzanetakis, Peter Driessen, and Andrew Schloss. 2005. "A Comparison of Sensor Strategies for Capturing Percussive Gestures." In *Proceedings of the 2005 Conference on New Interfaces for Musical Expression*, 200–203. National University of Singapore. http://dl.acm.org/citation.cfm?id=1085995.

Weiser, Mark. 1991. "The Computer for the 21st Century." *Scientific American* 265 (3): 94–104.

Wiley, Meason, and Ajay Kapur. 2009. "Multi-Laser Gestural Interface—Solutions for Cost-Effective and Open Source Controllers." In *Proceedings of the International Conference on New Interfaces for Musical Expression*, 43–44. http://www.nime.org/proceedings/2009/nime2009_043.pdf.

Wright, M., and A. Freed. 1997. "Open Sound Control: A New Protocol for Communicating with Sound Synthesizers." In *Proceedings of the 1997 International Computer Music Conference*, 101–4. http://en.scientificcommons.org/26532287.

Yang, Q., and G. Essl. 2012. "Augmented Piano Performance Using a Depth Camera." *Ann Arbor* 1001: 48109–2121.

Yu, Neng-Hao, Li-Wei Chan, Seng Yong Lau, Sung-Sheng Tsai, I. Hsiao, Dian-Je Tsai, Fang-I. Hsiao, Lung-Pan Cheng, Mike Chen, and Polly Huang. 2011. "TUIC: Enabling Tangible Interaction on Capacitive Multi-Touch Displays." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2995–3004. ACM. http://dl.acm.org/citation.cfm?id=1979386.