Modeling Natural Systems in Immersive Electroacoustic Sound

Eli Michael Stine
Greenville, NC

B.M., Oberlin Conservatory, 2014
B.A., Oberlin College, 2014
M.A., University of Virginia, 2018

A Dissertation presented to the Graduate Faculty
of the University of Virginia in Candidacy for
the Degree of Doctor of Philosophy

Department of Music

University of Virginia
May 2019

Ted Coffey, Ph.D.
Luke Dahl, Ph.D.
Leah Reid, D.M.A.
Worthy Martin, Ph.D.

Abstract

How can models of natural systems be used to compose electroacoustic music? To explore answers to this question, the author presents software built in the Max programming language and multi-channel electroacoustic compositions made using that software that explore different ways to musically encode the processes present in three natural systems: flocks of birds, island shorebird habitats, and oyster reef ecosystems. The process of building and using representative models of these systems leads to their extension into novel, natural system-inspired sound production methodologies. Spatialization is privileged as a domain for both listening to systemic properties and as central to the compositional practice of telling ‘system stories’ through sound. Supplementing the presentation and discussion of these projects is an overview of relevant historical threads within the domains of natural computing, algorithmic acousmatic composition, sonification and data-driven music, live electronics, and software art, along with the introduction of an evaluative framework for work of this type. Broader topics explored include the musical potentials of different models of natural systems, the differences between how humans experience and computers encode natural systems, and how sonic re-embodiments of natural systems may be listened to.

Contents

Acknowledgments

This dissertation is both a new departure, in comparison to some of the artistic pursuits that I engaged with in the years prior to its development, and also the embodiment of a lifelong, sometimes latent, obsessive interest and lens for viewing the creation of sonic art and engaging with the world in general. It could not have been completed without the help of numerous people.

Thanks first to my music department advisors Ted Coffey, who pointed me towards cybernetics and reinforced the value of mystery in art, Leah Reid, whose breadth of knowledge in the fields of music history, analysis, and theory helped the cohesiveness of this document, and Luke Dahl, whose comments on formatting, citation, and sourcing were highly constructive. Also thanks to Worthy Martin in the computer science department for being outside reader and asking highly clarifying questions. Thanks to my scientific collaborators who contributed to the projects within this dissertation: University of Virginia Environmental Sciences doctoral candidates Alice Besterman and Martin Volaric, and also others who, while not contributing directly to this dissertation, impacted its development: computational biochemist Charles McAnany and the members of the 2017-2018 University of Virginia Ecological Methods Lab. Thanks to my funding providers: the Institute of Humanities & Global Cultures, who funded the Ecological Methods Lab, facilitating several trips to the Virginia coast, the Environmental Sciences Department's Moore Award, which funded my work with intertidal oyster reefs on the Virginia coast, the Music Department, whose personnel (especially Travis Thatcher), studios, and equipment was essential to much of this work, and the Jefferson Scholars Foundation, whose fellowship has given me time, space, and a community within which to create this document.

Special thanks my parents, Mort Stine and Tracy Donohue, the former who has willingly leapt headfirst into the vast variety of different artistic and technological endeavors I've been involved with over my lifetime and taught me of the connectedness of things, and the latter who took me to the beach, to the swamp, to the river, to the woods, exposing me to what it means to be outside in all of its form, but often forms including water. Finally, a big thanks to Heather Mease, Kevin Davis, Chris Luna-Mega, Alex Christie, and the rest of the Composition and Computer Technologies graduate students during my time at University of Virginia, who were immensely grounding and present during this entire process, acting as sounding boards to bounce ideas off of, providing emotional support, and also needed communal respite.

ELI STINE

MODELING NATURAL SYSTEMS
IN IMMERSIVE ELECTROACOUSTIC SOUND

2019

Introduction

This dissertation seeks answers to the following question: how can models of natural systems be used to compose electroacoustic music? Records or simulations of natural systems may be applied to all timescales of electronic music through a process of mapping from systemic parameters to musical parameters. For example, recorded meteorological activity of temperature, air pressure, and water vapor could control the amount, cutoff frequency, and resonance, respectively, of a filter applied to a recording over time, or the location of particles within a turbulence simulation could control the spatial localization of multiple oscillators over time (Figure 1).

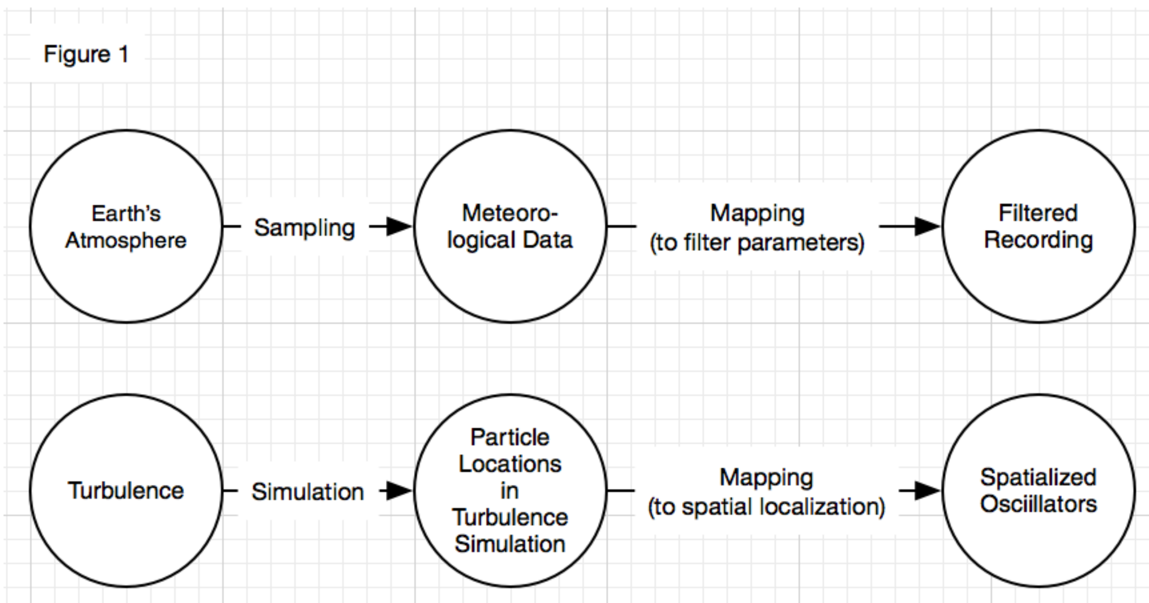

Flow chart of natural systems being modeled and used to generate electroacoustic sound

The result is electroacoustic material that is a sonic re-embodiment of modeled properties of natural systems. Careful compositional engagement with this material may result in electroacoustic music that affords the listener the ability to perceive, through sound, properties of those natural systems, including intra-systemic interactivity, gestural complexity (beauty, even), and/or self-adaption over time. Alternatively, those materials may be engaged with in a way that de-privileges a direct, indexical sonic presentation in favor of one that explores the artistic potentials of re-mixing, processing, or otherwise abstracting these natural system-derived materials.

The choice of 1) the system, 2) the design of the model of that system, 3) the mapping from parameters of that model to sound, and 4) the compositional presentation of the results of those mappings are all intimately connected techno-creative acts. Some or all of these design decisions factor into a listener's contextualization and reception of the work (Figure 2).

Figure 2

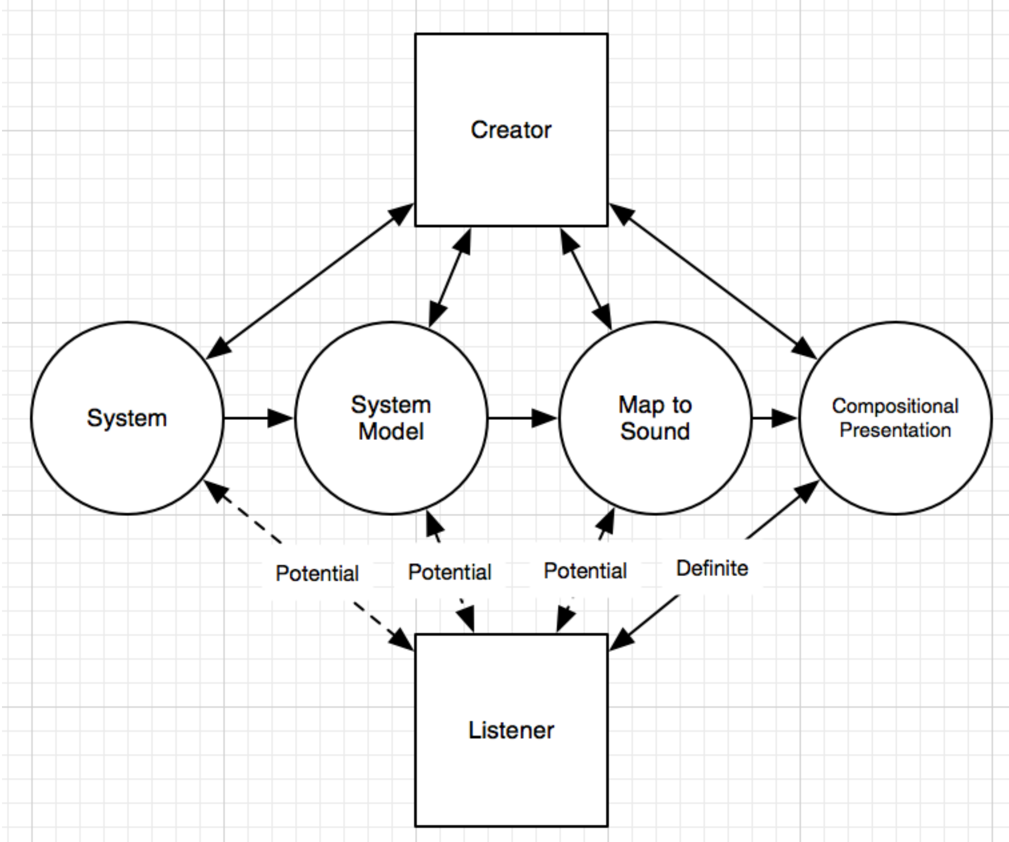

The interconnections between a system, its model, its map to sound, and its compositional presentation with respect to creator and listener

Within the medium of multi-channel electroacoustic music, sound localization as a parameter may be focused on in at least two ways. First, spatialization may be harnessed as an effective domain for listening to the activity of natural systems, particularly if the model of the natural system retains its original 2- or 3-dimensional spatial distribution. For example, spatial sonifications of brain scans allow medical professionals to detect and localize diseases in ways that could not be done with non-spatialized sound (let alone purely visual displays).[1] Second, spatialization may be positioned as central to the compositional practice of creating music using these systems, what may be described as telling “system stories” through the domain of spatialization. For example, a spatially-representative model of an avalanche may be cast as the setting for a story, with a composer zooming in on localized, cause-and-effect interactions (perhaps sonically re-embodied as call-and-response gestures) or zooming out to a sonic re-embodiment of the entire, emergent process (Figure 3).

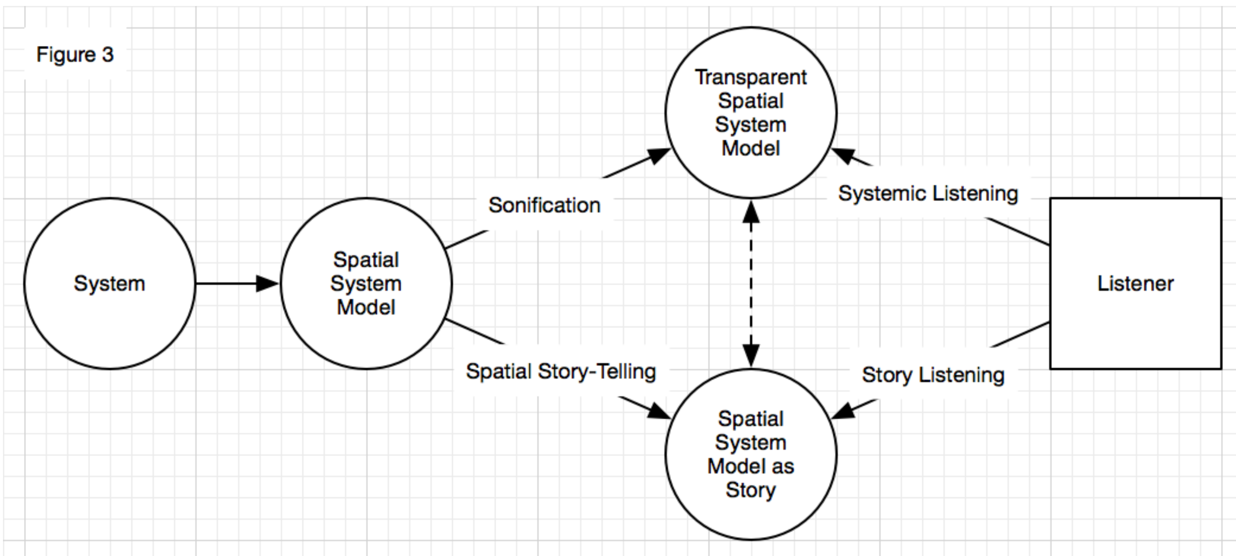

Two potential engagements with a spatialize system model

To support this dissertation, the author presents software built in the Max programming language and multi-channel electroacoustic compositions made using that software that explore different ways to musically encode the processes present in three natural systems: flocks of birds, island shorebird habitats, and oyster reef ecosystems. The latter two of these natural systems were engaged with in the context of inter-disciplinary environmental science, more specifically collaborations with Environmental Sciences doctoral candidates at the University of Virginia. These compositions and software have engaged with the world at large through presentations at the 2018 New Interfaces for Musical Expression (NIME) Conference (Blacksburg, VA), 2018 CubeFest (Blacksburg, VA), 2018 International Computer Music Conference (ICMC) (Daegu, South Korea), at Burning Man 2018 (Black Rock Desert, NV), 2018 Coastal Futures Conservatory Conference (Charlottesville, VA), 2018 Long Term Ecological Research All Scientists Meeting (LTER ASM) (Pacific Grove, CA), 2019 San Francisco Tape Music

Festival (San Francisco, CA), 2019 Association for the Sciences of Limnology and Oceanography (ASLO) Aquatic Science Meeting (San Juan, Puerto Rico), 2019 Workshop on Intelligent Music Interfaces for Listening and Creation (MILC) (Los Angeles, CA), the 2019 International Conference on Computational Intelligence in Music, Sound, Art and Design (EvoMusArt) (Leipzig, Germany) and the 2019 International Computer Music Conference/New York City Electroacoustic Music Festival (New York, NY).

In addition to engaging with representative models of these natural systems, these softwares extend the system models to hypothetical, imaginary systems, expanding a passive sonification practice into one in which the techno-artistic act of modeling and mapping to sound inspires the creation of novel, natural system-inspired sound production methodologies. The act of designing a system to spatially sonify shorebirds on the Virginia Barrier Islands suggested a spatialization paradigm within which the listener traverses an imaginary terrain populated by sounds, for example.

This expansion process leverages a (collaborative) working knowledge of the underlying natural systems, real-world, personal experience with them, and points to an artistic process which positions the development of tools to compose electroacoustic music with natural system models not as a means to an end but rather an applied exploratory method to investigate broader topics: 1) the musical potentials of different models of natural systems, 2) the differences between how humans experience and computers encode natural systems, and 3) how sonic re-embodiments of natural systems may be listened to. While the scope of this text, including its context as a final thesis within a Music Composition degree, is not large enough to engage these topics in a rigorous manner (a task which would require extensive engagement with sub-disciplines of semiotics, musicology, and psychology, among others), they are at the heart of the author's creative praxis and contextualize all the words within these pages.

This dissertation is divided into five chapters. First, expository materials, including working definitions of the words present in the title of this work ('system', 'natural', 'modeling'), are presented (Chapter 0). Next, an overview of relevant historical threads within the domains of natural computing, algorithmic acousmatic composition, sonification and data-driven music, live electronics, and software art is developed

(Chapter 1), leading to the introduction of an evaluative framework for work of this type (Chapter 2), and lastly the presentation and discussion of the projects undertaken by the author (Chapter 3). The appendix (Chapter 4) includes documentation of these projects along with several glossaries of terms and all references.

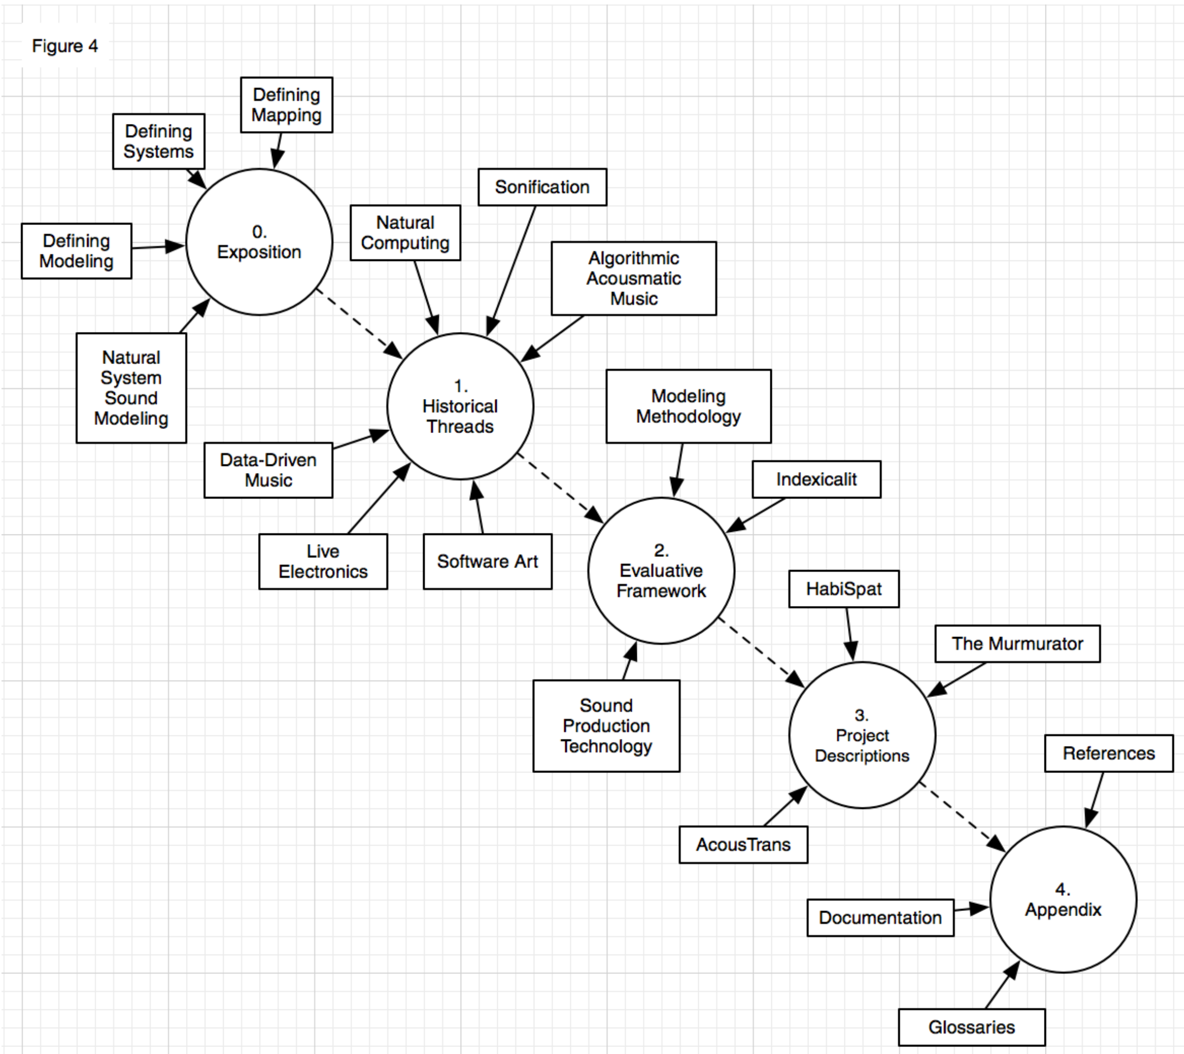

An overview of this dissertation, its sections and their contents

0. Exposition

0.1 What is a system?

When presented with the word and concept ‘system’ the first thing that may come to mind is a tool: a computer or software system. Or something biological: the body’s nervous or digestive system. Alternatively, you may think of a much larger, physically distributed entity such as a subway or banking system. Or something less physically tangible such as a philosophical, governmental, or musical system. What a ‘system’ may be defined as is a very broad concept, “so broad, perhaps, that it might seem impossible to find common ground between various definitions.”[1] Regardless, I will attempt to seek some common ground (or at least build common ground through definitional comparison) between ‘system’ definitions from several fields.

Within the context of Systems Theory, a meta-discipline which seeks to analyze and manage disorder within mostly scientific disciplines, one very broad definition of a *system* is “a set of interconnected parts which function together as a complex whole.”[2] Within the field of Environmental Science an environmental system is defined in relation to the environment as a whole: “the environment in its entirety may be regarded as a single system consisting of smaller, interconnected sub-systems.”[3] Within an applied Systems Engineering context, in particular the Introduction to the NASA Systems Engineering Handbook, a system is “a construct or collection of different elements that together produce results not obtainable by the elements alone. … The value added by the system as a whole, beyond that contributed independently by the parts, is primarily created by the relationship among the parts; that is, how they are interconnected.”[4]

It is clear from these definitions that a system consists of more than one element, that it has multiplicity. Further, each of these three definitions describes the relationship between the multiple elements as “interconnected.” The last definition, from Systems Engineering, posits that this interconnectedness is “primarily” what causes a system to be effective (“produce results”) in a way that would not be possible if the elements were disconnected. This provides an answer to the question “why?” in response to the ancient adage “the whole is greater than the sum of its parts,” and also points to another property of systems: unification via particular relationships between elements.

Thus, important properties of systems shared by each of these definitions are non-convergent multiplicity and non-divergent unity, in other words, systems are collections of *elements* that maintain their separate identities but at the same time are bound by their *relationships*.

Having broadly defined systems, **in what ways can we describe systems?** How a system may be described depends on one's viewpoint. A *hard systems viewpoint* is held by those who focus on quantitative metrics, are concerned with optimization, and typically look at systems as meeting an understood need. A *soft systems viewpoint*, on the other hand, is held by those who focus on qualitative metrics, may view systems as so complex (messy, e.g. human systems) as to be un-optimizable or do not seek to optimize systems, and may analyze systems whose purposes are entirely unclear.[5]

We first ask, using a primarily hard systems viewpoint: **how may systems be bounded?** Systems that have non-permeable boundaries may be said to be *closed systems*, whereas systems in which there is a relatively free exchange of materials across the boundary may be said to be *open systems*, with input and output. In an open system, if the output of the system informs its input, a *feedback loop* is formed (Figure 5). The world outside of a system's boundaries may be called its *environment*. An example of a closed system is that of a mechanical clock, which has a particular behavior that is self-contained and does not require external input (although it may be considered part of an open system if the clock-winder and viewer of the clock are included!). An example of an open system is a biological cell, as it receives nutrient energy through its cell wall and disposes of waste and enzymes, etc. outwards into the surrounding environment.

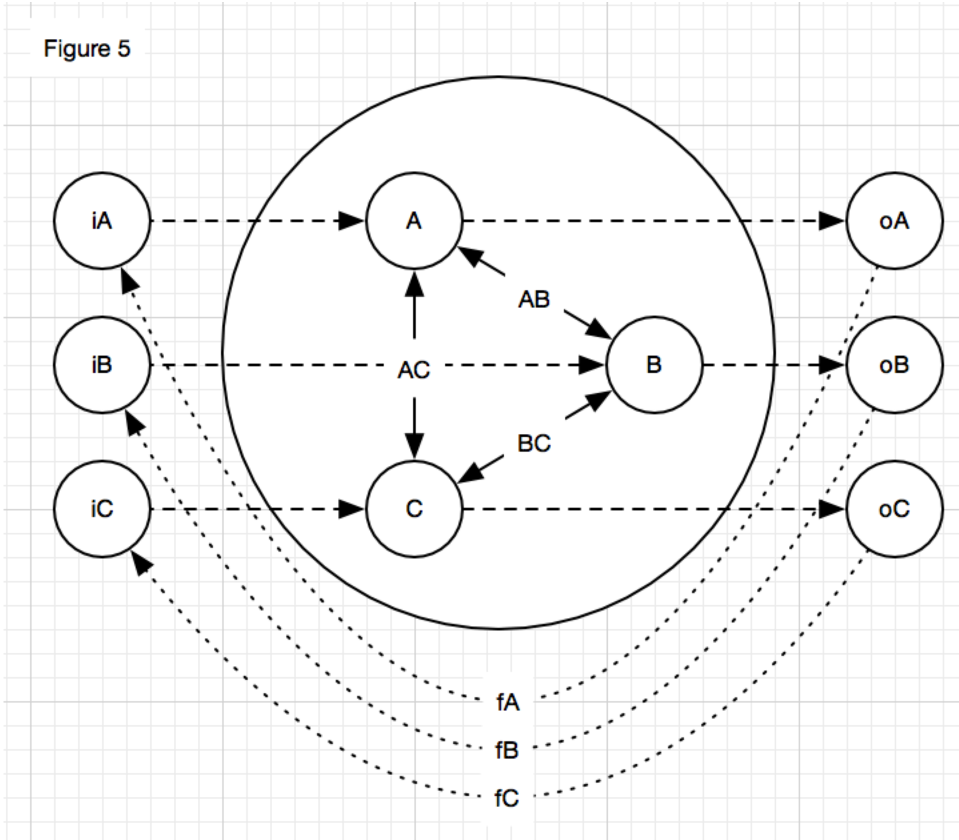

System ABC with elements A, B, and C, elemental relationships AB, AC, and BC, inputs iA, iB, and iC, outputs oA, oB, and oC, and potential feedback loops fA, fB, and fC.

The *scale* of systems is also important to consider. Environmental system scale may range from "all of the living material on earth (the biosphere), to successively smaller scales (such as individual forests), down to the level of single organisms (such as an individual tree)."[6] Systemic scale is a function of at what point elements are considered indivisible and at what point the effects of the interactions of those elements are unregarded. For example, a system defined with nerve cells as indivisible elements that only focuses on their interactions within the brain of an iguana is of an absolutely smaller scale than one in which all cells within all animals are considered.

We may also describe a system by defining its *dynamics*, how it changes over time as a function of elemental relationships and, if open, its input and/or output. Before describing these dynamics it is worth noting that our perception of the system's elements and their relationships is a model of the system (seen through the 'veil of perception', after Locke), not the system itself, and so any and all system dynamics descriptions are in reference to this perceived (or measured) version of the system. A description of a system at a discrete moment in time may be called a system's *state*. Dynamics (at least discrete dynamics) describe how a system changes from state to state (possibly in response to input and/or output), or, put another way, we know the dynamics of a system if given a state (and input and/or output) we can describe its next state.

A *simple system* is a system whose dynamics are apparent from an analysis of its elements in isolation, and, by the Systems Engineering definition above, may not even be considered a system at all, rather just a collection of elements. A *complex system* is a system whose dynamics are not apparent from an analysis of its elements in isolation. When a system as a whole exhibits properties which are meaningful only when attributed to its whole, not its parts, what is called *emergence*, that system is complex. For example, the picture emerging from a completed jigsaw puzzle, the shapes of sand dunes produced by the wind, or, perhaps, the self-awareness of human beings, are all examples of systems that demonstrate *emergence*. A system that has *chaotic behavior*, exhibiting extreme sensitivity to the initial state of its elements, is also complex. Chaotic dynamics are the result of sequences of interconnected events (events whose outcomes set the stage for the next event) that over time aggregate to produce seemingly erratic system behavior. Chaotic behavior is *not* the result of noise or randomness and may appear in completely closed systems. Examples include species population dynamics, fluid turbulence, or systems as simple as two pendulums connected end to end (Figure 6). An understanding of the complexity of a system (whether it has emergent properties, is non-linear, and/or chaotic) informs how that system can and cannot be modeled.

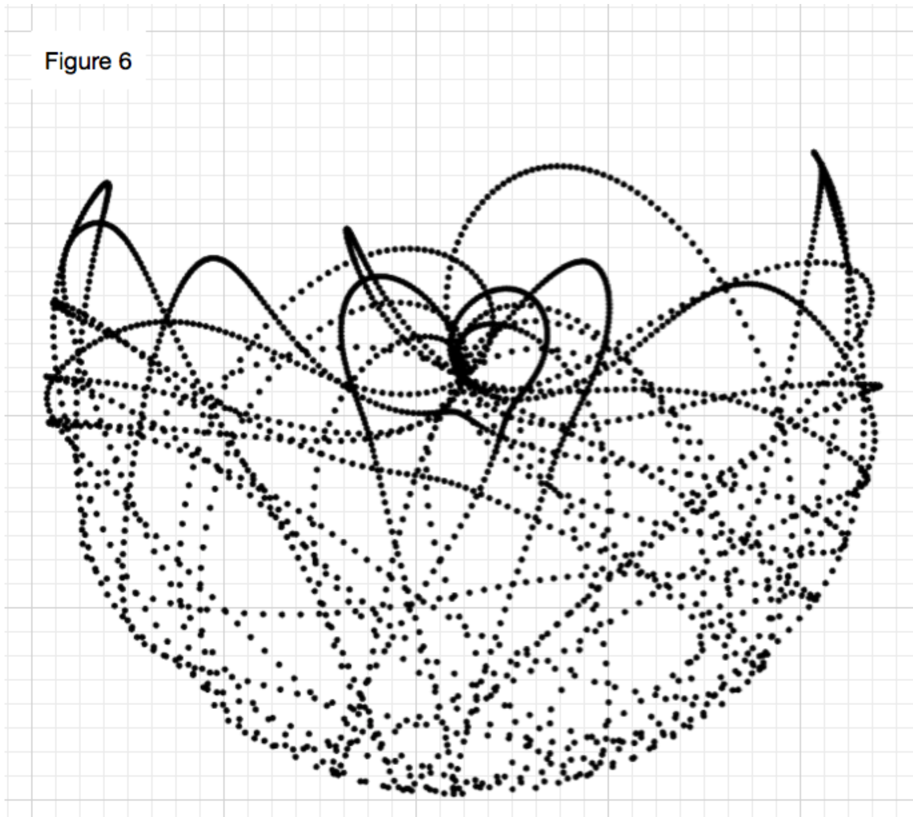

A tracing of the location of the lower mass on a double pendulum simulation, demonstrating chaotic behavior

Given a particular environment (say, Earth), open systems come about that are more or less able to endure their environment, a property called *adaptation*. All physical systems disintegrate over time (as a function of the most likely closed nature of the universe and the second law of thermodynamics), but some systems are able to last longer than others. Over time (quite a long time) a certain set of these systems was able to adapt long enough to develop the ability to *reproduce*, to pass on instructions for their organization to a next generation existing after their disintegration. These systems — what we might call living systems — may reach a level of complexity that causes them to not only be adaptive to their environment but to be *homeostatic*, to self-regulate in order to maintain *systemic equilibrium* depending on their environment and internal state. They may also engage in intra-systemic relationships: pollination, predation, or symbiosis. These systems include animals, plants, fungi, and micro-organisms. ***Natural systems***, the systems which this dissertation concerns itself with, include these living systems and non-living systems whose dynamics are comparable to those of living systems, such as atmospheric or geological systems.

The reasons for this dissertation's focus on natural systems are two-fold: first, from a hard systems viewpoint these systems are open, complex, adaptive, and sometimes chaotic, all properties that the author also ascribes to interesting musical dynamics and listening experiences. Second, from a softer systems viewpoint, you are a natural system, you exist within other natural systems, and so while they are incredibly complicated, natural systems are experientially effortless (simply go outside), relatable, and familiar to listeners. These two contrasting views of the same concept, I think, are a particularly exciting recipe for achieving the artistically sublime.

0.2 How can a system be modeled?

A *model* is an imitation of some entity that allows for analysis, experimentation, or other procedures which, when applied to the original, would be expected to produce identical results. Models are made for a number of reasons. A scaled-down model of a building allows an engineer to test its strength without the expense or resources of constructing the full-scale model. A model of a tornado, constructed using fans within a laboratory setting, allows for a scientist to test the effects of a tornado on a material without having to use an actual, real-world tornado. A model can also be trans-modal, imitating a material entity within computer circuits, for example (or, less often, materializing a computational process with physical resources).[7] Computational models harness the speed, parallelism, and precision of computers, and are used extensively within modern scientific research and business.

What types of models are there? For the purposes of this dissertation, two primary types of models will be considered: *sampled* and *simulated*. Sample-based models imitate an original entity through (possibly trans-modal) records of that entity. For example, the recordings of the meteorological data applied to different filter parameters in the introduction of this text model a meteorological system through samples, a finite set of captured, real-world values (often stored on digital media) that imitate the original. Alternatively, simulation-based models imitate an original entity through (again, possibly trans-modal) simulation, constructed via observed behaviors, measured elemental relationships, and/or known environmental catalysts. For example, the computational turbulence simulation applied to oscillator panning in the introduction emulates the activity of turbulence within a computer (perhaps through approximations of the *K*-epsilon equations, which were devised by optimizing towards measured turbulence behavior), providing a stream of data that models the activity of turbulence.[8] Using the right methods, the boundary between sampling and simulation is permeable. Sampled data may be used to train a model (using machine learning) that simulates the original, and simulation can provide a (possibly infinite) stream of (second-hand) 'sampled' data (Figure 7). Beyond machine learning to create simulations is an entire field dedicated to simulation design, a process which necessarily requires an exceptionally comprehensive understanding of the entity to be simulated, which, after construction and analysis, may lead to even more systemic comprehension.

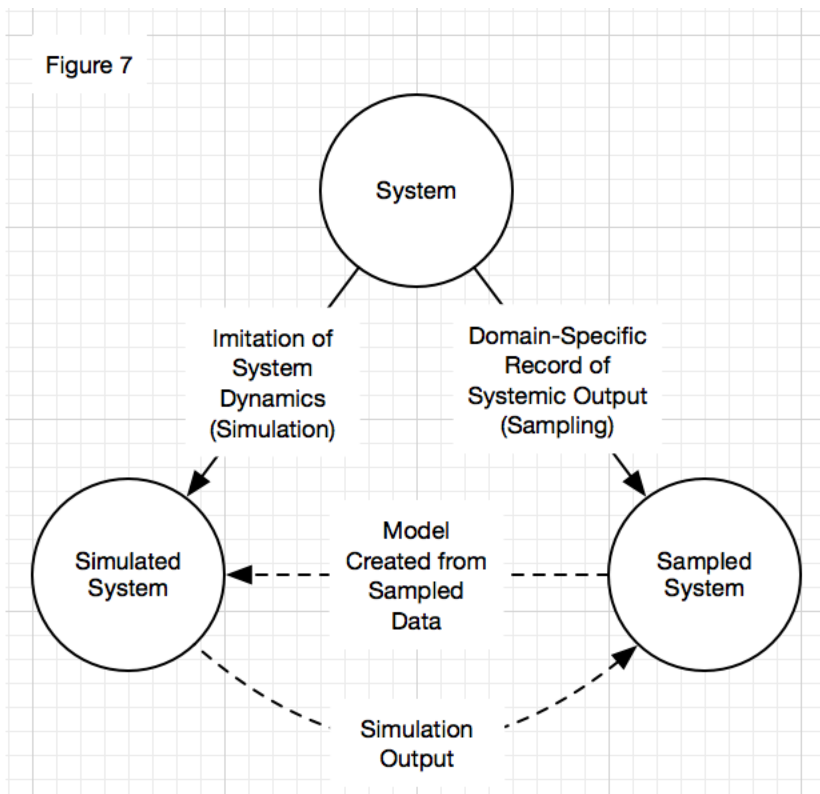

A system and its simulation- and sample-based models, with potential permeability between these two types of modeling

Regardless of whether or not it is sampled or simulated, creating a model of any entity necessitates *modeling decisions* that best represent desired properties of the original. For example, consider a scaled model of a skyscraper that has exceptional detail but is made of one material and a scaled model of that same skyscraper that is less detailed but is made of scale-proportionate materials of the skyscraper. Both are models of the skyscraper, but the former may be more usable within a display or measurement context, whereas the latter might be useful to test the resiliency of the skyscraper to earthquakes or hurricanes. There are an infinite number of potential models of an original, and there are always differences between an original and its model (otherwise the model is just a complete re-creation) (Figure 8). It is up to the modeler(s) to guide these differences into areas that will least negatively affect the way in which the model will be used, or alternatively, to highlight these differences (especially within an artistic context): to use the process of modeling to reveal a particular way of encoding or knowing an entity (Figure 9).

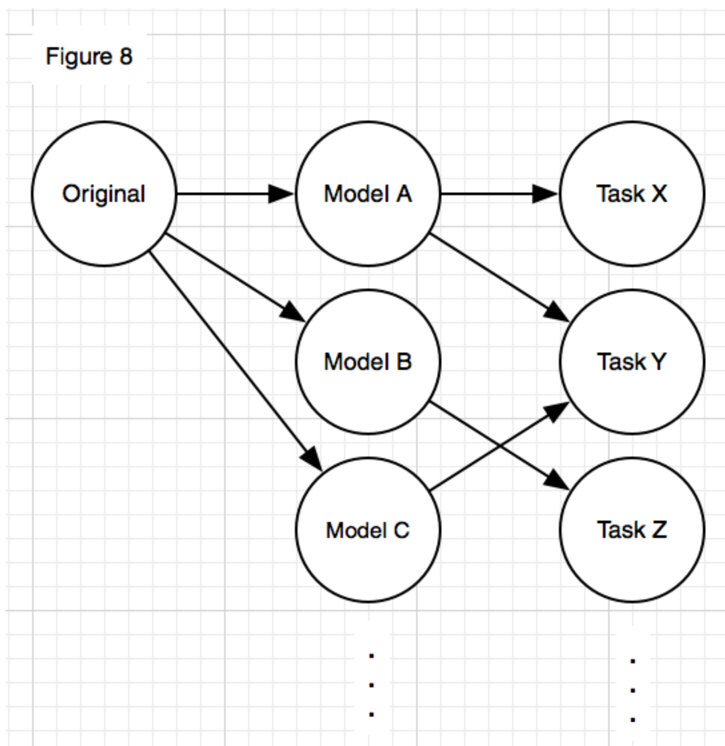

An original entity, its infinite potential models, and then an infinite set of tasks that those models are/aren't suited for

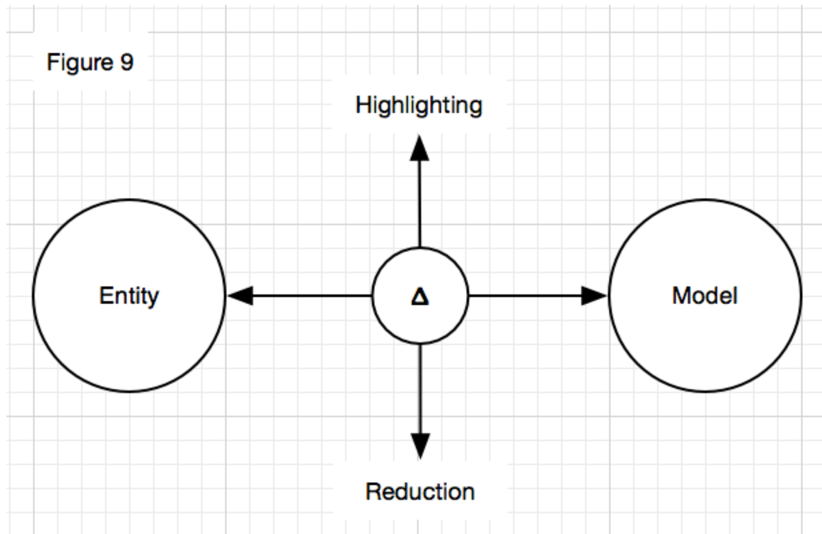

The difference (Δ) between an entity and its model recast as something to be reduced or something to be highlighted

Within the context of systems modeling, the curation of differences requires an understanding of how best (whatever "best" means for a particular context) to retain the behavior of a system, something that was shown in the last section to hinge on an understanding of that system's dynamics: the type and quality of connections between different elements within the system and how it handles inputs (if they exist). Sufficient models of complex systems, including natural systems, should model their properties of emergence, adaptation, non-linearity, and/or chaos. Complexity within a computational simulation setting can very often lead to *computational intractability*, the inability of a computer to efficiently (within a reasonable amount of time and/or resources) compute the next state of a particular simulation, for example. Simulating natural systems on computers thus requires enough computer science knowledge to avoid intractability, or the re-use or adaptation of tractable simulations, to arrive at usable, reliable models. Additionally, within the context of this work, modeling of natural systems is a pre-cursor to a mapping to musical parameters, discussed next, and so the format of the output of the model and/or the way in which the relationships between elements in the model may be accessed should be carefully considered. Put another way, our natural system models should not only represent the dynamics of the original system, but encode the output of those systems, and possibly their elemental relationships, in a way that can be effectively mapped.

For example, a bird's syrinx (the avian equivalent to the larynx, our voice box) is a natural, sound-producing system. Shaped like an inverted Y and located within the throat of a bird, the syrinx produces sounds via the vibration of its walls and a small piece of cartilage called the pessulus (Figure 10) . The syrinx has been proven to demonstrate self-oscillatory non-linear dynamics, allowing some species of birds to produce more than one sound at a time and/or to mimic human speech.[9]

There are a number of different ways to model this complex system. A sample-based model of the syrinx might be a sound recording (a sonic domain imitation) of a bird call or a video of the syrinx in action, perhaps recorded as a sequence of 2-dimensional pictures made with high-field magnetic resonance imaging (MRI). A simulation-based model of the syrinx might be rendered using physical materials, electronic circuits, or through a computer simulation. These simulations necessarily require an understanding of the elements and elemental relationships of the syrinx system and a way to map these to different mediums: non-living matter, connected electronic components, or computer instructions, respectively.

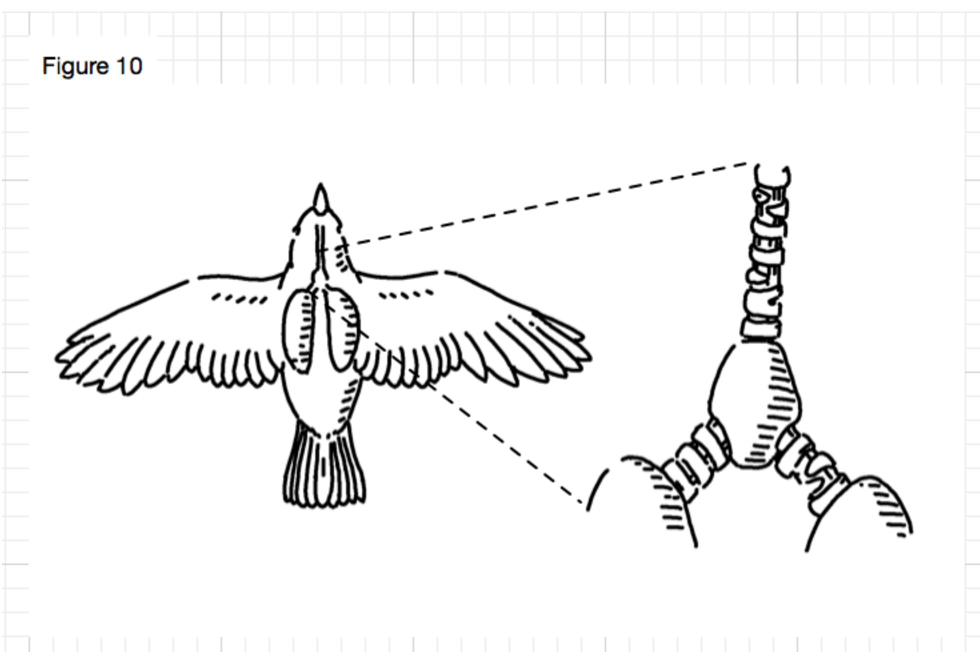

Respiratory system of the song bird, including the syrinx, within the body and detail (after Boswall)

As previously discussed, the efficacy of a particular model depends on its context, including the systems viewpoint of the modelers and model users. Within the context of a scientific discipline such as ornithology, a sampled-based model of the syrinx of a particular species of bird may not be sufficient until its sample corpus includes representations of the aforementioned demonstrations of self-oscillatory nonlinear dynamics. A model used by those with a softer systems viewpoint may be much more lenient: a corpus including only several samples of bird calls might suffice. A simulation-based model of the syrinx within a scientific context might be insufficient unless it make quantitative use of the field of syringeal biomechanics, accurately simulating the intricacies and complexities of the elements and elemental relationships within the syrinx.[10] From a softer systems viewpoint, a simulation that outputs a coarse mimicry of the sound of a bird may suffice. Models may even be purposefully designed to be more distant from the original than they could be, particularly within the context of art-making. These system model distortions or *systemic caricatures* may be constructed to highlight or critique a particular quality of the system, to synthetically develop the system (suggest what it could be), or to amplify, rather than minimize, differences between the model and the original.

Regardless of its context, each syrinx model affords different outputs, different *handles* to re-present the behavior of the modeled syrinx system. In addition to the handles granting access to the output of the system, a model may also afford access to the underlying elements and elemental interactions that determine that output. With a sound recording syrinx model, the recording of a bird's call may be reproduced over a loudspeaker, filtered, time-stretched, or otherwise processed as time series data. Those recordings alone do not, however, grant the users of those models information about how those calls came about. In contrast, a computer simulation of the syrinx allows the output of the system to be used like a sample-based model, but in addition the variables, underlying functions, and other computational processes which define the model may also be accessed. Not only the behavior of the system, but how that behavior comes about, is available. How system models may be mapped to electroacoustic sound is discussed next.

0.3 How can a systemic model be mapped to electroacoustic sound?

What is a map? When the word 'map' is brought up it typically refers to a piece of paper that represents in two dimensions some three-dimensional real-world place

(Central Virginia, for example). “Map” in this context is synonymous with a model: the inked piece of paper (or a digital surrogate) is a model of the real-world locale. A *mapping* is the particular procedure by which an original entity is adapted to a map, and is a concept that can be extended to many contexts. With the map of Central Virginia, a particular scaling down of the real-world locale is chosen and is assumed to be uniform. Different colors are chosen to indicate different elevations. Text is added at locations to indicate different counties. Solid and dotted lines are added to indicate the boundaries of those different counties. And so on.

This mapping is a byproduct of 1) the map maker’s resources on the real-world place to be mapped and 2) the functional and aesthetic choices they make. Functional choices are often informed by mappings that have come before. For example, a color scheme that is consistently used to indicate topographical information is preferred over another color scheme, so as to potentially harness a map reader’s previous experience. Just in case this mapping meme is not present in some map reader’s mind, a *key* to the mapping is nearly always included beneath the map. Without this key, anyone who engages with the map has to deduce the mapping or assume some (possibly incorrect) mapping, a common process that will be returned to soon.

A mapping may be dynamically altered through an *interface*. The turning of a knob to change the volume of a speaker system, clock-wise being louder, counter clock-wise being quieter, is dynamically altering the mapping of volume of sound being played back over the speaker system. Changing the gear on a bicycle alters the mapping of which sprockets and chainrings the chain goes around (determining the bike’s gear ratio). Clicking and dragging to change the size of a window within a WIMP user interface is altering the mapping of that window onto one’s screen real estate. Interfaces which alter the mapping of another interface may be called *meta-interfaces*. In the case of the volume knob, a switch that swaps the clock-wise/counter clock-wise directionality of the knob or a slider that changes the sensitivity of the knob are both such meta-interfaces, interfaces that alter the mapping of another interface (Figure 11).

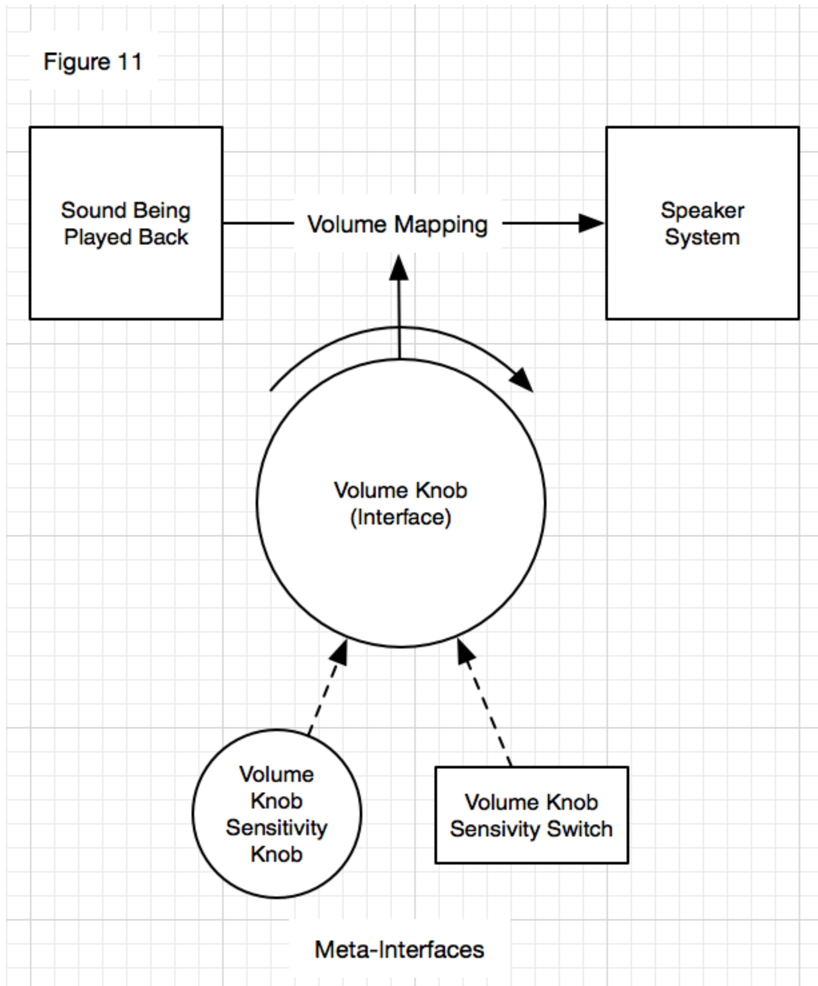

Visualization of the volume knob as interface to a mapping,
with additional knob-affecting meta-interfaces

How does one map to electroacoustic sound? Electroacoustic sound is sound made using loudspeakers, electroacoustic transducers that introduce vibrations into a medium (often air) that are correlated to an electrical signal, that, if within the human-audible range (often cited as 20Hz to 20kHz), are perceived by humans as sound. Within the context of mapping a system model to electroacoustic sound, the loudspeaker (or a group of loudspeakers, which, when placed around a listener, creates a 360° sound field), is analogous to the blank sheet of paper on which the map-maker draws their map. The map-maker's straight edge, compass, and, more recently, computer-assisted design (CAD) softwares and/or Geographic Information Systems (GIS), are analogous to a large set of electroacoustic sound-mapping tools, methodologies developed in electroacoustic sound-related fields (audio production, sonification, algorithmic music, electroacoustic composition, among many others) over the past century and a half.

In contrast to applying these tools to a sheet of paper (or its digital surrogate), electroacoustic sound is an immaterial mapping destination experienced over time (although it may be said that viewing and experiencing a map is a time-based procedure!). Just as with map-making, however, the resources available from a system model (its handles), along with functional and aesthetic considerations, determine how a model is mapped. A map to electroacoustic sound may take advantage of previous maps to sound, harnessing listener experience with electroacoustic sound (and music). Mapping may be done in non-realtime, producing fixed media, or controlled live, via interfaces and meta-interfaces. Pre-existing interfaces for controlling the mapping from a system model to sound may be used or custom interfaces may be created.

Electroacoustic sound-mapping may be applied, and its effects may be perceived, over what Curtis Roads calls the different time scales of music: the macro (overall form, minutes), meso (phrases, several seconds), sound object (events, sub-second), and micro (timbre, the thresholds of auditory perception (0.00005 to 0.05 seconds, or 20Hz to 20kHz)).[11] Examples of mappings to electroacoustic sound that occur at the micro scale include playing back the output of models at audio rate over loudspeakers. Examples at the sound object scale include electroacoustic events triggered by certain behaviors of a model. Examples at the meso scale include curated phrases of musical material triggered by certain behaviors of a model (necessitating a harmony between artistic intervention and systemic dynamics). Examples at the macro scale include large-scale changes in a model affecting the sound world's evolution over time (perhaps minutes or even hours) (Figure 12). Real-world examples of mappings at each of these scales will be discussed in detail in the next chapter.

Of course, musical experiences take place at all of these timescales simultaneously, and the time scales in between these time scales, for example, gestures at the threshold of timbre and rhythm, are often particularly interesting areas of musical expressivity. There exists a rich body of research within the scope of the psychology of music and cognitive science that explores these different temporal layers and their boundaries. For more information see Chapters 6 and 7 of Tan and co-author's

Psychology of Music: From Sound to Significance (Routledge), Section 1.2 of William Sethares' *Rhythm and Transforms* (Springer Science & Business Media), and for a classic electronic music composer's perspective Karlheinz Stockhausen's *The Concept of Unity in Electronic Music* (Perspectives of New Music).

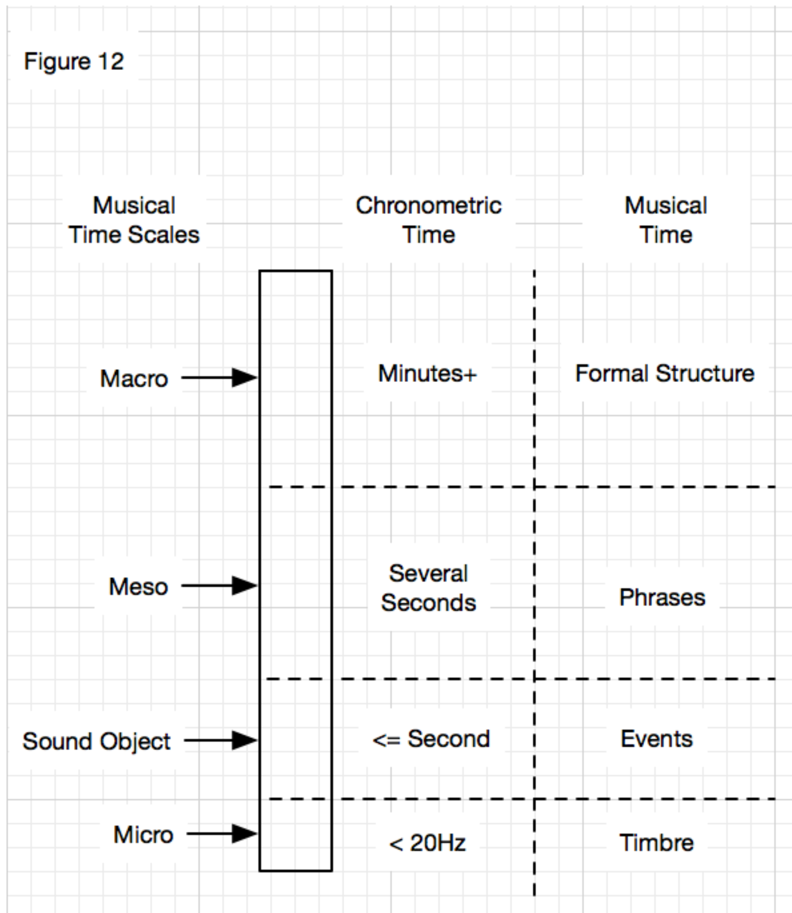

Musical time scales, after Roads

Sound is not only experienced over time, but also within **space**. While a smooth, one-dimensional scalar spectrum for time is not as nuanced as its experience by humans, one for space seems even more artificial. This is because space is experienced through a vast network of objects, events, and their environments, each of which contributes to a perception of 'space'. Some criteria that may be described quantitatively are localization, size and distance, either instantaneously from the perspective of a listener or over time, corresponding to angular velocity, perceived growth/shrinking, and change in distance, respectively, all of which contribute to perception of movement (Figure 13). Electroacoustic sound, especially multi-channel electroacoustic sound, can replicate real-world spatial experiences to greater-or-lesser degrees (depending on a number of factors, discussed later in this text). Ways to describe different spaces and space-related dimensions of electroacoustic sound include Simon Emmerson's spatial frames, which are, at progressively smaller scales: the landscape (the entire acoustic horizon), the arena (a sub-part of the landscape), the stage (within the arena, separating a listener from an event), and the event.[12] Dennis Smalley also discusses different scales and types of spaces, some defined via size and distance, others in relation to movement, and yet more that are intimately intertwined with extra-spatial musical dimensions. These include Perspectival Space (space perceived from the listener's vantage point), Enacted Space (space produced by human activity), Nested Space (the embracing of one space within another), Immersive Space (perspectively and spectrally filled-in space) and many others.[13] Electroacoustic space is also an important characteristic within the context of natural systems and their models because all natural systems (living or non-living) take up real, physical space, and some may be experienced directly through immersive sonic experiences (fauna vocalizations in the rainforest, crickets chirping within a field, or bullfrogs croaking around the perimeter of a pond, for example).

Figure 13

Spatial Criteria	Instantaneous Space	Space-over-Time
Localization	Where?	Changing Localization
Size	How Big?	Growing/Shrinking
Distance	How Far?	Approaching/Receding

Different spatial criteria of electroacoustic music, after Emmerson, Smalley, Wishart, others

The flexibility of electroacoustic sound processing (particularly using computers) allows us to not necessarily have a one-to-one mapping between system time/space and musical time/space. For example, within the time domain, long-term system dynamics may be sped up to the time scale of timbre, or minuscule, incredibly fast-paced elemental relationships within a system may be slowed down to the scale of minutes or hours. Within the spatial domain, 3-dimensional data that takes place over miles can be reduced to the size of a living room, or sub-atomic level spatial interactions may be diffused within headphones. Additionally, relationships over time may be translated to relationships over space and vice versa. Distances between particles in a turbulence simulation may be re-cast as a waveform (which, when looped at a rate within the frequency domain, results in a particular timbre) or harmonic intervals or rhythms or section lengths. The magnitude of meteorological measurements may inform the spatial size, perceived distance, or localization of electroacoustic sounds.These types of non-one-to-one and/or cross-domain temporal and spatial mappings allow for an enormous amount of temporal and spatial sound-mapping plasticity, creating unique experiences of systemic properties and potentially interesting sound worlds and musical dynamics (Figure 14).

Figure 14

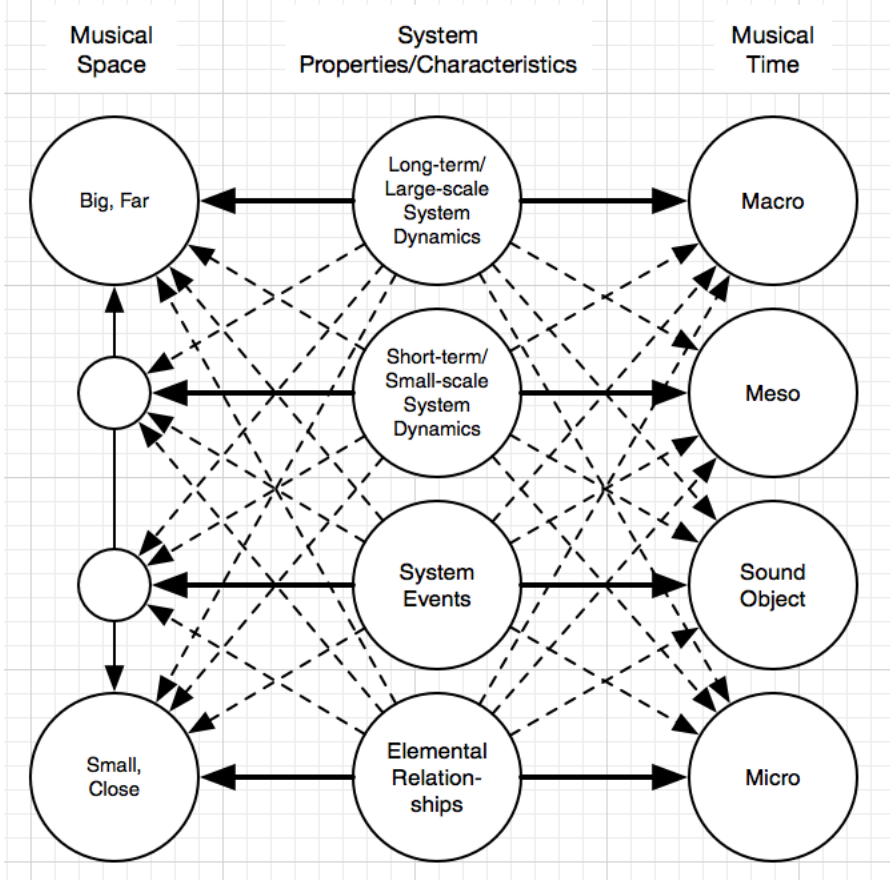

Mappings from temporal and spatial systemic scales
to temporal and spatial musical scales.
Dotted lines indicate non one-to-one mappings.

0.4 How can mappings of a natural system model to sound be used to make electroacoustic music?

There are innumerable potential mappings from systemic model handles to musical materials. These materials take many forms: they may be audible or they may be functions, lists, sequences of events, or other data. Consider the meteorological data (temperature, air pressure, and water vapor) mapped to the parameters of a filter (amount, cutoff frequency, and resonance) applied to a recording described in the introduction. This mapping consists of a set of system model handles (in this case, a three-dimensional time series), a set of sound processing parameters (here the settings of a filter), and then some recording that is to be filtered. The handles are functioning as streams of data that alter parameters over time, in turn affecting the recording, resulting in electroacoustic sound that is a recording filtered by a model of a meteorological system (Figure 15).

Within the process of composition, it is possible that a first mapping will be deemed successful, but more often than not an iterative sequence of trial-and-error, a repeated looping of execution (mapping) and evaluation (listening), is required. For example, within our example context, the range at which the cutoff frequency of the filter is driven by air pressure, say, from 200Hz to 5kHz, may need to be widened to 100Hz to 8kHz to get a particular musical result. After re-mapping and listening, the water vapor handle may be determined to be best mapped to cutoff frequency instead, with the air pressure handle instead being mapped to resonance. A particular change in the data (perhaps a rain front moving in) may also suggest a corresponding change in mapping. Or a mapping itself may be constructed in such a way that it alters its own definition, adapting to what is being mapped (the higher the average air pressure, the larger the cutoff frequency range, for example). A composer may also determine a set of successful mappings and use them as constraints for an interface deployed within a live performance setting.

The more a mapper maps, the more sensitive they may get to the nuances and affordances of different types of data, data handles, and mapping destinations. The field of sonification (discussed in the next chapter) also engages with quantitative strategies for optimizing mappings in terms of data transparency, which might also inform a composer's mapping strategies.

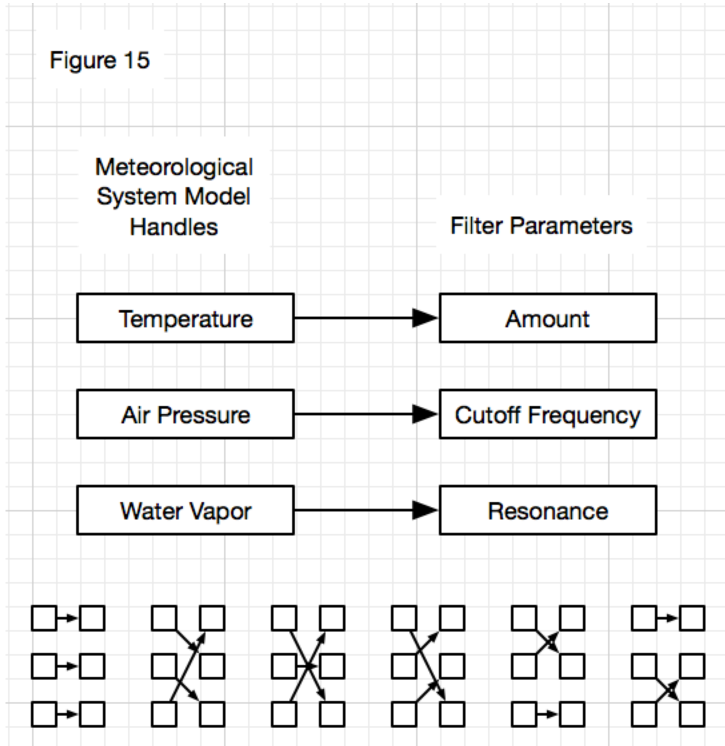

Multiple handles being mapped to multiple parameters, and the numerous possible configurations of them, even excluding the particularities of the mappings

After mappings have been determined, certain temporal segments deemed more musically interesting than others may be extracted through a process of *curation*. Those materials may then be *abstracted*: reordered, layered, time-stretched or compressed, processed with effects, or otherwise altered using electroacoustic sound-processing techniques. Both curation and abstraction affect the *indexicality* of these materials: how much they are direct indices to their system model handle origins. Indexicality, and its reception-oriented counterpart *comprehensibility*, will be discussed later in this text, but suffice it to say that the ways in which these musical materials are altered is consequential.

For the author, choices of mapping, curation, and abstraction serve **telling a story through the medium of electroacoustic sound.** This story might be created through artistic engagement with the system's behavior, reflection on the way in which it is modeled, exploration of how the systemic model handles are mapped, or a striving for sublimity in the aesthetics of the musical materials they are mapped to. These processes might not necessarily be clear or observable, nor need the process of creating the story be uni-directional. Rather, the act of electroacoustic story-telling is a techno-creative experience which might suggest re-modeling, re-mapping, and deepened levels of abstraction at all points within the process. The composed stories are either non-realtime, fixed electroacoustic media, encoded as one or more digital audio files to be played back within a particular speaker configuration, or, alternatively, are performed: existing as a set of interfaces (and possibly meta-interfaces) that control system model maps, or even models themselves, in real-time.

There are as many ways to listen to these electroacoustic stories as there are listeners. Composers and theorists including Pierre Schaeffer, Michel Chion, Dennis Smalley, and Simon Emmerson have each defined different sets of “listening modes”. These sets diverge in a number of ways, but also share several modes. These include listening that focuses on sound sources (Schaeffer’s Mode 1 (‘Ecouter’), Chion’s causal listening, Emmerson’s heightened listening), listening that engages with a language or code in order to interpret sounds as messages (Schaeffer’s Mode 4 (‘Comprendre’), Chion’s semantic listening), and listening that engages with the traits of a sound itself (pitch, rhythm, timbre, etc.), separate from its source or meaning (Chion and Schaeffer’s reduced listening (Mode 3, ‘Entendre’)) (Figure 16).[14] Additionally, these authors and others have developed more nuanced electroacoustic listening frameworks, focusing on the work as a whole, as in Emmerson’s Language Grid, for example[15], or exploring a wide range of context-specific and emotion-provoked/evoking listening modes or styles (e.g. ‘sing-along listening,’ ‘fault listening,’ and ‘distracted listening’) within the work of David Huron[16]. A listener may switch between different listening modes over the course of a work, or use two at once (listening both causally and semantically, for example). For a comparative study of listening modes see Tuuri and Erola’s “Formulating a Revised Taxonomy for Modes of Listening” (Journal of New Music Research).

Figure 16

Listening Focus	Schaeffer	Chion	Emmerson
Sound Sources	Mode 1	Causal	Heightened
Messages via Sound	Mode 4	Semantic	
Traits of Sound	Mode 3	Reduced	

A list of pertinent electroacoustic listening foci,
after Schaeffer, Chion, and Emmerson

In addition to a multiplicity of listening modes, how stories, and potentially narratives, are received within electroacoustic music has received attention from a number of composers and theorists. These theories emphasize how this medium privileges the agency of the listener, going so far as to say that the act of listening to this kind of music is a process of auto-ethnographic composition, according to Katharine Norman.[17] The context of the electroacoustic story-telling experience is also analyzed: the periods before (the sounds and sound experiences the listener comes to the work with), during (what is heard while engaging with the work), and after the listening experience (the cool-down period, where the work as a whole is retroactively analyzed) are addressed within the writings of Natasha Barrett.[18] James Andean, extending analyses of listening modes, outlines ten different 'narrative modes' of listening to acousmatic music, ways in which listeners actively engage in creating stories while listening.[19] These modes include material narrative (story-telling via the interactions of real-world, recognizable sounds over time, what Smalley calls first-order surrogacy), mimetic narrative (where structure and sound materials are organized in a manner recognizable from our experience of the world; the gravity-induced rhythmic trajectory of a bouncing ball, for example), extra-musical narrative (including what we will define as *keys* in the next section), and spatial narrative (the acousmatic work as a series of spaces or spatial transformations) (Figure 17). Just as with listening modes, different narrative modes may come and go over the course of listening and several may be co-present at a given time.

Figure 17

Story Medium	Activated Listening Mode(s)
Sound Sources	Causal
Formal Structure	Reduced, Semantic
Mimesis (System Dynamics)	Causal
Spatialization	Causal, Reduced
Sound Processing Technology	Technological
Extramusical Materials (Keys)	N/A

A list of pertinent mediums (or channels) for story-telling with electroacoustic sound after Andean, Smalley, Gaver, others

These listening modes and theories of electroacoustic story-telling have special properties within the context of compositions that map natural system models to sound. A natural system model's dynamics over time might function as a type of mimetic narrative, perhaps called a 'systemic narrative'. An environmental or everyday listening mode might be activated, a listening focus akin to how one listens in a non-musical or non-technologically mediated context. This mode can be enhanced through multi-channel loudspeaker configurations, affording a physically immersive, perspectival listening experience that contrasts with a traditional two-dimensional, front-oriented loudspeaker experience.

Additionally, a composition that maps natural system models to sound might attempt to foreground comprehensibility of the natural system(s) at play. This understanding of the mapping from natural system model to musical materials might be called a *key*, in the same way that a key on a physical map indicates a particular mapping used within it. Knowledge of these keys may be gleaned from particular types of listening or extra-sonic resources. Ways of understanding keys through listening include acousmatic or live spoken descriptions of the mapping used (as in Westerkamp's *Kits Beach Soundwalk*, where sound processes are explicitly defined through speech), engaging semantic listening, or sounds that a listener associates with a particular source (mapping procedure, model, or natural system), engaging causal listening (as is the case with a sound being sourced as emitting from a bird's syrinx, for example). Extra-sonic channels for conveying keys include text within a program note or the work's title. Video projected during performance or a physical, material process unfolding on stage during performance (a visible biosensor strapped to a performer's arm, for example) might also act as keys. How present or hidden a work's keys are to an audience depends on how much a composer chooses to foreground their mapping methodologies in combination with each audience member's experience.

Regardless of how keys are interpreted, the reception of the compositional presentation of these electroacoustic stories is **a conversation in sound between the listener, the creator, and the natural system(s) in play**, encompassing their modeling, mapping, curation, and abstraction.

0.5 Introducing Natural System Sound Models

This exposition is concluded by an applied condensation of all of the discussed topics and methods into the concept of *natural system sound models* (NSSMs) (Figure 18). A NSSM may be defined as a constellation of entities and methods acting on those entities, namely:

One or more *natural system(s)*, which act(s) as source for…

one or more *modeling procedure(s)*, which construct…

one or more *natural system model(s)*, which act(s) as source for…

one or more *mapping(s)*, which result in…

a set of *musical materials*, which act(s) as source for…

a *compositional process*, the result of which is…

a *musical presentation*, live or fixed.

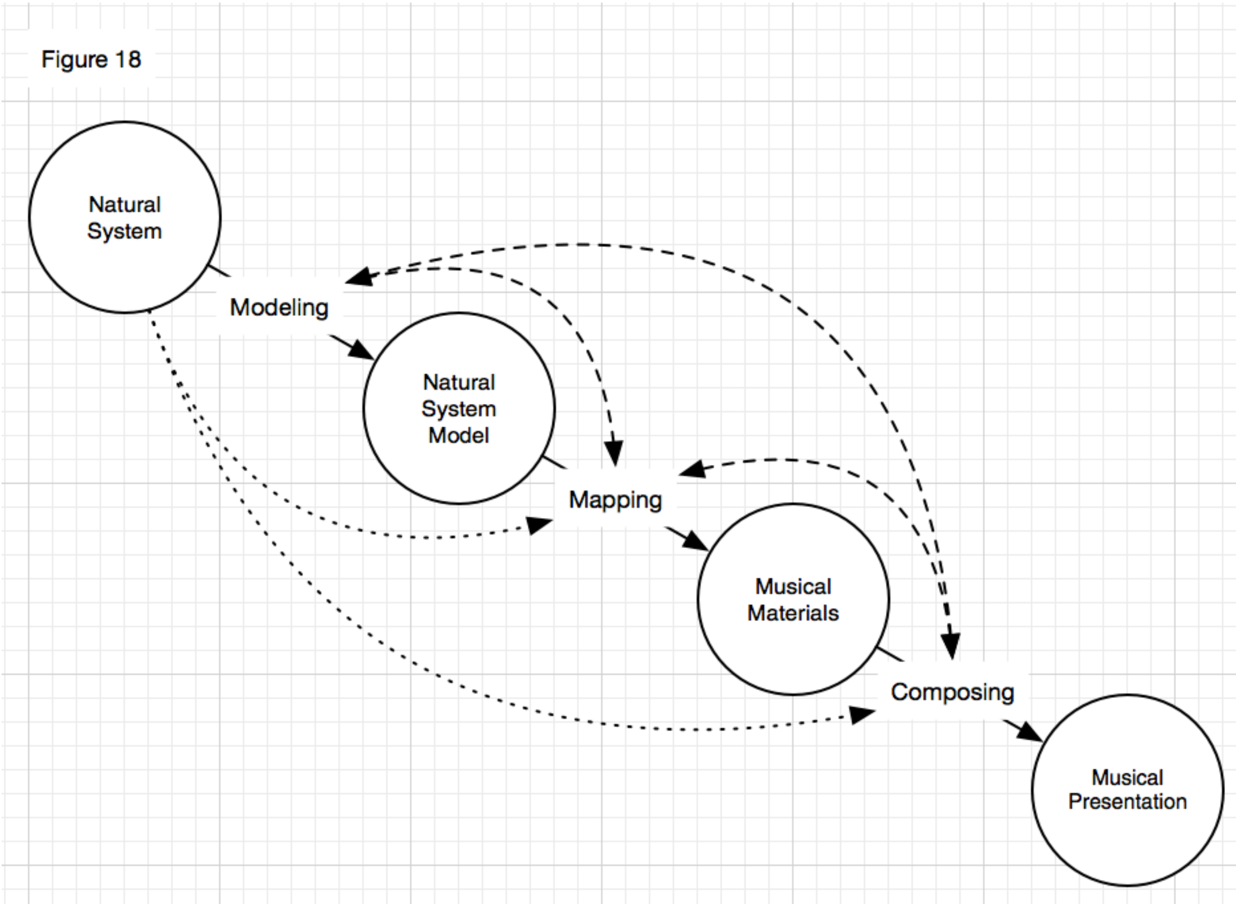

As discussed in the previous section, this constellation is replete with potential self-affecting interconnections. Each of the procedures (modeling, mapping, and composition) may inform the other. For example, the act of composing a desired musical effect may require changes to a mapping or even the model itself. Additionally, the natural system, as experience, as concept, as system, is essential not only to the modeling procedure, but also the mapping method, composing process, musical materials and, ultimately, the musical presentation. This concept foregrounds these interactions as meaningful and unavoidable.

Work that seeks to sterilize a model of a natural system and deal with it objectively is opposed to this concept. Work that doesn't give attention to the intricacies of the modeling process, regardless of how soft or hard a systems viewpoint is used, or how technologically engaged the composer is, is opposed to this concept. Work that forces a natural system model into a particular context, rather than letting it define its own context, mapping schemes, and emergent sound world, is opposed to this concept. Work that focuses on the simplest, top layer, rather than the more complex, inner-workings of natural system models, regardless of how quantitative or qualitative[20], is opposed to this concept. There are many works that engage with natural systems, their modeling, and their mapping to sound (or other media), but only a subset of these works resonate with the concept of NSSMs as defined in this text. The next chapter explores this subset through discussions of algorithmic acousmatic music, sonification and data-driven music, live electronics, and software art.

1. Historical Threads

What are the historical contexts of using natural system sound models? An answer requires engagement with a number of different artistic and scientific practices, along with answers to sub-questions related to ‘what’, ‘how’, and ‘why’: given a particular artistic practice or field of work, what natural systems, models, mapping procedures, and aesthetic contexts are in play? How are these systems modeled, mapped, and composed with? Why might the paradigm of NSSMs be preferred over some other paradigm?

This chapter traces historical threads within the fields of algorithmic acousmatic music, sonification and data-driven music, live electronics, and software art. There are no necessarily stark boundaries here, as praxis within each of these areas might bleed into the others. However, each field uses different models, mapping procedures, and/or has different sonic and socio-cultural contexts (perhaps even world views). Outlining these differences allow us to understand how this paradigm has lived and lives in the world.

1.1 Algorithmic Acousmatic Music

Algorithmic music may be defined as music that uses formal procedures to generate or manipulate musical material, and includes a wide variety of algorithm-centric musics over multiple centuries of music-making. Which procedures are chosen, how they are formalized, and which musical materials are generated or manipulated are all up to the *algorithmicist(s)*. Many classes of musical algorithms exist, ranging from those consisting of formulations of rules (a top-down approach) to those generated from analyses of corpora (a bottom-up approach, often used to imitate particular styles of music). The degree of control over the resultant music that the algorithmicist(s) or other involved agents has may vary from minimal (entirely automated) to substantial (interactive, perhaps even symbiotic). Musical algorithms may be applied to music at all of its temporal and spatial scales. Much algorithmic music research deals with symbolic representations of music: notes within scores or, within the computational domain, corresponding digital surrogates (e.g. Musical Instrument Digital Interface (MIDI) representations). Musical algorithms may also encode music as continuous parametric control: multiple dimensions in parallel changing fluidly over time. This representational choice may be more fitting for certain acoustic musics and the electroacoustic domain, including control of sound synthesis and processing techniques.

The history of algorithmic music is inextricably intertwined with historical developments in music theory, technology, and the cultural role of music. The use of instructions and formal processes to create music dates back to the time of Pythagoras. The ancient Greeks believed that music was inseparable from numbers: that systems of intervallic relationships, harmony, and rhythm corresponded directly to systems within the cosmos. For example, Ptolemy and Plato gave mythical form to this world view through their concept of the 'Music of the Spheres', the idea that the distances between the planets and their movements correspond to different musical modes and certain notes.[1] During the first half of the second millennium algorithmic thinking was expanded to automate compositional and performative musical processes.This includes the work of Guido of Arrezo, who created a system for generating melodic materials from texts, the improvisatory development of organum to accompany Gregorian chant within the church, as well as the 'composition machines' of Atanasius Kirchner.[2] The idea of automating composition reached maturity in the combinatoric 'musical dice games' of the 18th century, one of which is attributed to Wolfgang Amadeus Mozart.[3] These games involve rolling dice to choose how to assemble a number of small musical fragments that, when pieced together formed a new, magically melodious and harmonically successful composition.

Within the 20th century systems for automated composition are central to new compositional styles and methodologies. John Cage's engagement with chance and natural systems in the 50s and 60s to drive compositional choices redefined the role of the composer, performer, and audience. The twelve-tone method and integral serialism, techniques pioneered by the Second Viennese School from the 20s onwards, aimed to completely control all parameters of music in objective, abstracted ways. Schillinger's unique system of musical composition sought a balance between creative freedom and scientific rigor through a set of algorithmic composing tools (which, while failing to be widely adopted, resonated deeply with many disciples). Iannis Xenakis used meticulously researched stochastic processes to compose acoustic and electroacoustic music beginning in the 1950s, and was also an early adopter of computers in his compositional practice. These musics took advantage of algorithmic procedures including combinatorics, probability and aleatoric (chance) procedures, and rule-based systems (grammars that define sets of iteratively applied procedures to generate and alter materials).

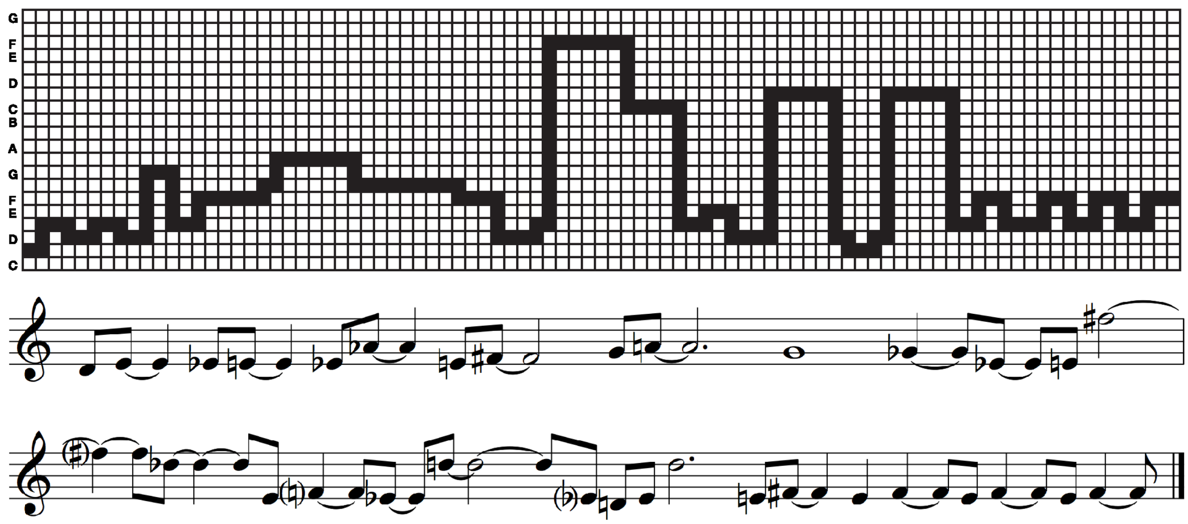

Application of Schillinger's Melody-Generating *Millimetrization* Technique to the Richmond Skyline

Using a rule-based system of iterative musical material generation, modification, and selection the *Illiac Suite*, the first computer generated composition, was performed by a string quartet at the University of Illinois in 1956.[4] From the 1970s onwards, computers not only got faster but their means of encoding data became more sophisticated. The computer became a tool to design and test novel, computer-specific algorithms, rather than simply a higher speed alternative to analog electronics or pen-and-paper calculations. Properties of these computer-specific algorithms include non-linearity, complexity, and high density (many of the same properties we ascribed to natural systems in the previous chapter).

Engagement with computer science, mathematics, and other sciences and humanities (notably linguistics and biology) during this period led to the musical use of many new classes of algorithms. These include Markov models and generative grammars (from linguistics), transition networks (from systems theory), as well as algorithms that attempt to model the behavior of natural systems on computers (from chaos theory and biology). Algorithms within this last set, deeply pertinent to this text, may be positioned within the field of *natural computing*. Briefly, natural computing (also called unconventional computing or biologically-inspired/bio-inspired computing) may be defined as the process of extracting ideas from nature to develop computational systems, or using natural materials (e.g., molecules) to perform computation.[5] This field is primarily focused on using natural system models to solve computer-based problems, for example, *artificial neural networks* developed to model the activity of the brain on computers being applied to the task of image recognition.[6] In addition, this field also focuses on simulation of natural systems by means of computing, examples including *Lindenmayer systems (L-systems)*, originally developed to model the growth processes of plants but that effectively simulate self-similar processes. Lastly, natural computing includes computing with living materials, for example, using slime mold rather than silicon to build logical gates, an area which has some interesting recent musical applications.[7] Other examples include *genetic* or *evolutionary algorithms*, which use probability and search techniques to model Darwin's theory of evolution, and *cellular automata*, which model simple interactions of agents that in turn may lead to deeply complex behavior.

While strategies for algorithmic music applied to the acoustic domain developed in the 20th century, developments in analog electronics and computers also affected the acousmatic medium. Early on in the history of computer-assisted composition a set of engineers and composers sought to make music not with the computer, but *for* the computer, using it as a source of sound rather than a symbolic music representation generator. This lineage of digital acousmatic music begins with Max Mathews' creating the *MUSIC* program at Bell Labs in 1957. *MUSIC* was then extended into the interactive command-line *CARL* System built in UNIX at UCSD in the 1970s, which was expanded into *CSound* by Barry Vercoe and others at MIT Media Labs in the 1980s.[8] These systems (and the many others derived from them) differ greatly in their capabilities, but share the ability to synthesize sound (using synthesis techniques such as additive, subtractive, and later, frequency modulation synthesis) and/or process digital recordings (re-arranging, time-stretching, adding effects, etc.). These capabilities allow these systems to engage with sample-based models of natural systems (digital audio recordings) and/or be the mapping destination of natural system simulation handles, receiving and making sound with the output of generalized natural computing models (such as those discussed above) or natural system-specific simulations.

In addition, digital surrogates to analog studio environments, called Digital Audio Workstations (DAWs) were developed from the 1970s on.[9] DAWs allow a user to call sound synthesis, sequencing, and processing functions through a recording studio-skeuomorphic graphical user interface (GUI), applying the speed and flexibility of the computer to a familiar music production environment. These systems may be used to provide a graphics-assisted methodology to curate and abstract musical materials derived from natural system models. The 1980s also brought the first real-time computer music applications (such as Perry Cook and Gary Scavone's Synthesis Toolkit (STK)), a topic which will be returned to later. For more information on the history of computer music see Douglas Kiesler's *A Historical View of Computer Music* (The Oxford Handbook of Computer Music) and Ge Wang's *A History of Programming in Music* (The Cambridge Companion to Electronic Music).

Having discussed the history of algorithmic music, computer algorithms and natural computing, and very briefly the development of computer music and its repercussions for NSSMs, what follows are examples of natural computing algorithms being applied to the creation of acousmatic music.

1.1.1 Simulated Acoustics: Physical Modeling Synthesis

A digital sound-based application of natural systems that doesn't involve the typical trans-domain mapping necessary for NSSMs (for example, the non-sonic meteorological data of the introduction needing to be mapped to sound-affecting parameters) is *physical modeling synthesis*. Physical modeling synthesis engages non-linear computational models such as simulations of masses and springs and systems of delay lines and filters to create expressive and nuanced emulations of the sounds of physical systems.[10] Typical applications of physical modeling synthesis include recreating the physical systems that go into making the sounds of acoustic instruments — a brass player's lips, the vibrating reed and resonant bore of a saxophone, or the hammer of a piano — as well as constructing systemic caricatures that simulate fantastical, physically impossible sound-producing entities: a thousand foot saxophone, a microscopic snare drum head, or ten foot long vocal folds. There exists a large body of scholarly research dedicated to physical modeling synthesis, which will not be outlined in its entirety here. Instead, several projects pertinent to NSSMs within the context of electroacoustic music-making will be noted.

In Chris Chafe's article 'Case studies of physical models in music composition' (2004) works by a number of artists using physical models are described. These include Claude Cadoz's *GENESIS* system, a NSSM constructed of a deeply polyphonic and multi-dimensional physical modeling toolbox that affords the creation of fantastical virtual instruments. An ecosystem of springs, masses, and other physical models produce acousmatic sound that ranges from other-worldly to recognizably instrumental, with the complex dynamics of this ecosystem creating its own emergent behavior over time in concert with large-scale physical parameter automation composed by the user. A visualization of *GENESIS*, provided a particular physical modeling ecosystem isn't too complex, provides the user keys to the mapping process used.

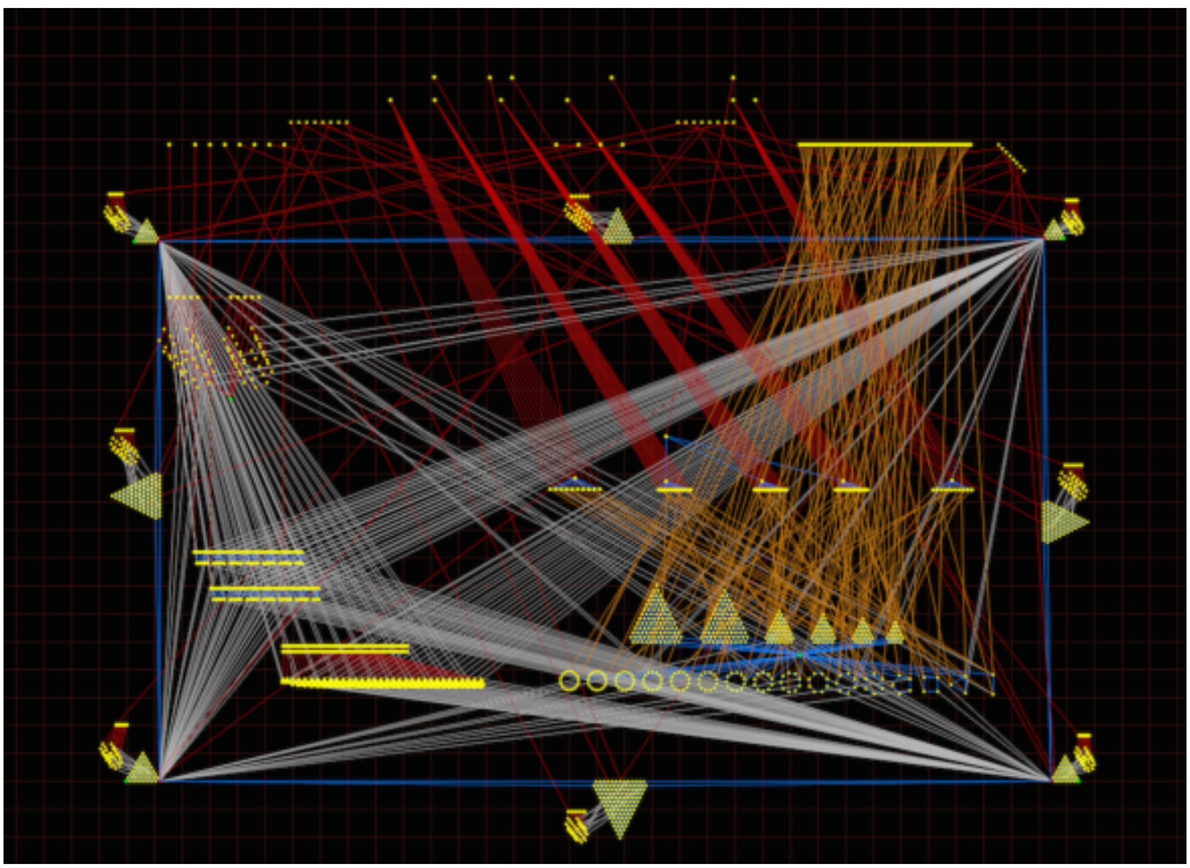

GENESIS system visualization

Another, more philosophical engagement with physical modeling synthesis is outlined by Juraj Kojs in 'The Language of Action and Cyberaction-based Music: Theory and Practice' (2009). This article situates itself within the work of J.J. Gibson and W. Gaver and frames uses of physical modeling synthesis methods as a toolbox for the conceptualization of actions and instruments in the digital domain, yielding emulations of extended, hybrid, and abstract cyberinstruments and cyberactions. Here the complex systemic interactions of physical modeling are mapped to different temporal and spatial scales, affording NSSMs that use computational models of physics and acoustics not just to synthesize timbre, but to define methods of interaction and virtual instruments.

Lastly, a more state-of-the-art commercial example is *Wavesolver*, a sound generation software system that provides high-quality offline sound synthesis through the resolution of animation-driven physical models.[11] Computer models of paper cups, cymbals, and dropping spoons are used to generate eerily realistic foleys, suggesting a future of automated sound replacement in the context of 3D motion graphics.

1.1.2 Simulated Growth, Movement, and Interaction: L-Systems, Flocking, and Cellular Automata

The distribution of natural systems over time and space, as embodied in the growth of plants and the movement and interaction of animals, is modeled in L-systems, flocking simulations, and cellular automata (CA). Because these systems have intuitive maps to visual and, to a lesser extent, sonic domains, there exists a large body of work using each algorithm in a wide variety of contexts. The NSSMs outlined here focus on applications that interface with mostly electroacoustic music composition contexts.

An L-system is an iterative rewriting process used for generating the fractal patterns of plants and trees, first formalized by Przemysław Prusinkiewicz and Aristid Lindenmayer in the late 1960s.[12] There exist many different types of L-systems (context-free, stochastic, parametric, timed, and more) but all L-systems consists of an alphabet of symbols that are used to define strings, a set of production rules that expand each symbol into another, possibly empty string of symbols, an initial string ('axiom') from which to begin, and, optionally, a method for translating the generated strings into geometric structures.[13] These geometric structures, if visualized, emulate the branching growth of trees, plants, or algae on a pond, or, if sonified, result in music whose timbre, gestures, rhythm, or overall formal structure emulates such structures (recursively defined structures) in sound. For more information on L-systems see Daniel Shiffman's *The Nature of Code* (Section 8.6), Jon McCormack's *Evolutionary L-systems* (in Design by Evolution, 2008), and Gerhard Nierhaus's *Algorithmic Composition* (Chapter 6).

A plant-like L-system generated with the rules
X=C0F-[C2[X]+C3X]+C1F[C3+FX]-X and F=FF applied over 6 generations

Perhaps because of their intuitive mapping to symbolic structures, the vast majority of musical applications of L-systems are for instrumental music. These include R. Luke DuBois' dissertation work with L-systems ('Applications of Generative String-Substitution systems in Computer Music'), which include L-systems with alphabets of only two symbols which are creatively mapped onto different intervals and rhythmic durations, as well as experiments with parametric mappings, where different symbols correspond to different musical parameter changes: for example, setting a lower voice up an octave, reducing the duration of the next note by a sixteenth note, etc.[14]

Stelios Manousakis is a Netherlands-based artist who has incorporated L-systems into purely acousmatic works such as *Do Digital Monkeys Inhabit Virtual Trees?* (2006). This NSSM uses what the composer calls 'L-System Digital Sound Synthesis' within a patch created using Cycling '74's Max. The composer explains: "The L-systems control a granular DSP engine in a multitude of different levels by generating a constantly varying number of musical agents with hierarchical relations that move, act and interact with and within their environment. The behavior of the agents is used to develop the initial musical seed causing complex contrapuntal patterns to emerge in all the time-scales of the composition."[15] In short, *Do Digital Monkeys Inhabit Virtual Trees?* is a NSSM that involves the mapping of the activity of L-systems to many temporal scales of acousmatic sound simultaneously, affecting the timbre (the quality of grains of the granular DSP engine), gestural activity (when grains are emitted), emergent rhythms, and formal structure of the work.

Flocking algorithms simulate the movement (and possibly environmental interactions) of many (hundreds, possibly thousands or more) elements within virtual physical space. These simulations are complex: the emergent behaviors of the flock derived from the individual interactions of its many agents are naturalistic and chaotic. Often, as with L-systems, this complex emergent behavior comes from a set of simple procedural constraints. For example, the original Boids algorithm simulates the movement of flocks of digital birds ('boids') through three simple, boid-centric rules[16]:

```
Separate: move away from Boids that are too close
Align: move towards the average heading of Boids nearby
Cohere: move toward the average position of Boids nearby
```

Flocking algorithms suggest not only a means for NSSMs to engage with naturalistic group movement (to inform parameter spaces, rhythms, formal structure, etc.) but also, quite intuitively, immersive sound spatialization: the mapping of the positions of the agents in the flock to localized sound sources within a multi-channel loudspeaker array. As an example of this, Enda Bates and Dermot Furlong employ the Boids algorithm in a NSSM that generates spatial data which are mapped to the locations of sounds produced by granular synthesizers.[17] Bates and Furlong also model the Doppler effect, early reflections, and global reverberation, which, combined with the naturalistic movement within an immersive loudspeaker context, heightens the sense that the flocking is taking place in physical space around the listener.

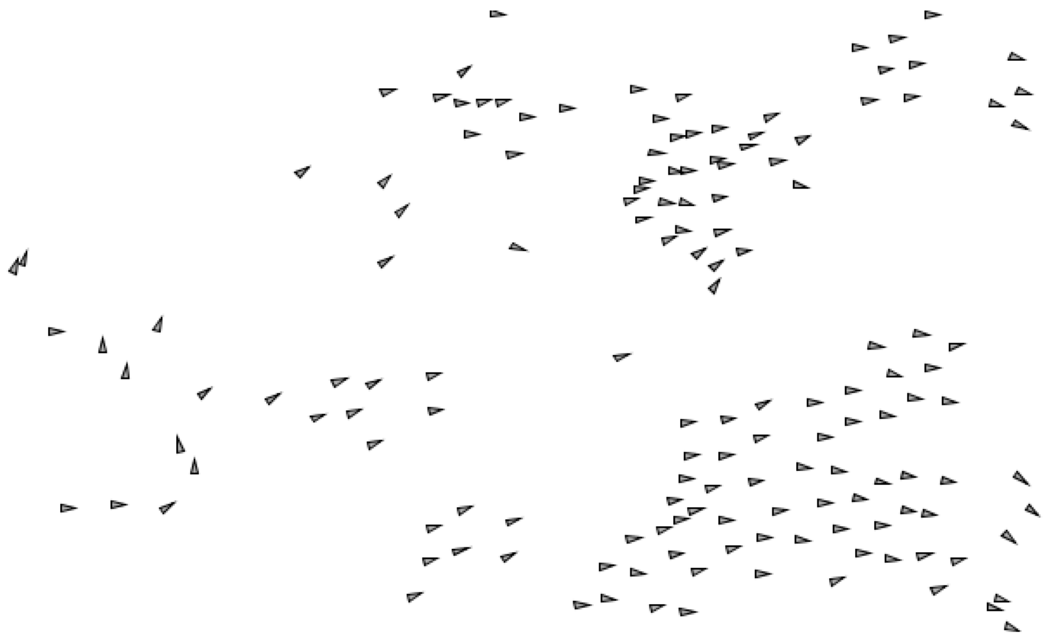

Simulated Flocking Behavior

Schacher and co-authors use Boids-like swarm algorithms to generate sound and produce visuals in a manner similar to that of Bates and Furlong, but use another model to automate control of the flock. The different constraints of the flock are managed using a Finite State Machine (FSM) which automates large-scale changes in the system over time, an example of multiple models being used in tandem within the context of a NSSM. In their 'Future Work' section, Schacher and co-authors also express interest in relating the movement of the flock to its sonic output: "endowing agents with the capability to perceive aspects of the acoustic output," suggesting a system within which agents have 'self-awareness': changing their behavior in response to the sound they produce.[18] The author's own *Mumurator* system is an example of a flocking algorithm-based NSSM which does just that (among other things), with full details outlined in the Project Descriptions chapter.

Like flocking systems, Cellular Automata (CA) produce complex global behaviors based on the interactions of simple elements. Unlike flocking systems, however, these individual elements do not 'move' like agents; instead, each element within a grid (or torus, or sphere) of elements may be in one of several states: 'on' or 'off', 'alive' or 'dead', or a wide range of other values in multi-valued CA. The state of an element depends on the states of its neighbors in accordance with a set of prescribed transition rules: for example, a rule might state that if three of an element's neighbors are 'on' then that element should transition to an 'off' state. Different types of CA as defined by these transition rules may be organized into classes which range from those that tend towards all cells being in the 'off' state, to those that tend towards fixed or oscillating patterns, on to those with chaotic and complex patterns.[19] One famous cellular automaton is The Game of Life (GoL), created by John Conway in 1970, which produces a wide range of anthropomorphic patterns, some of which have been named: 'blinker', 'glider', and 'toad', for example.[20]

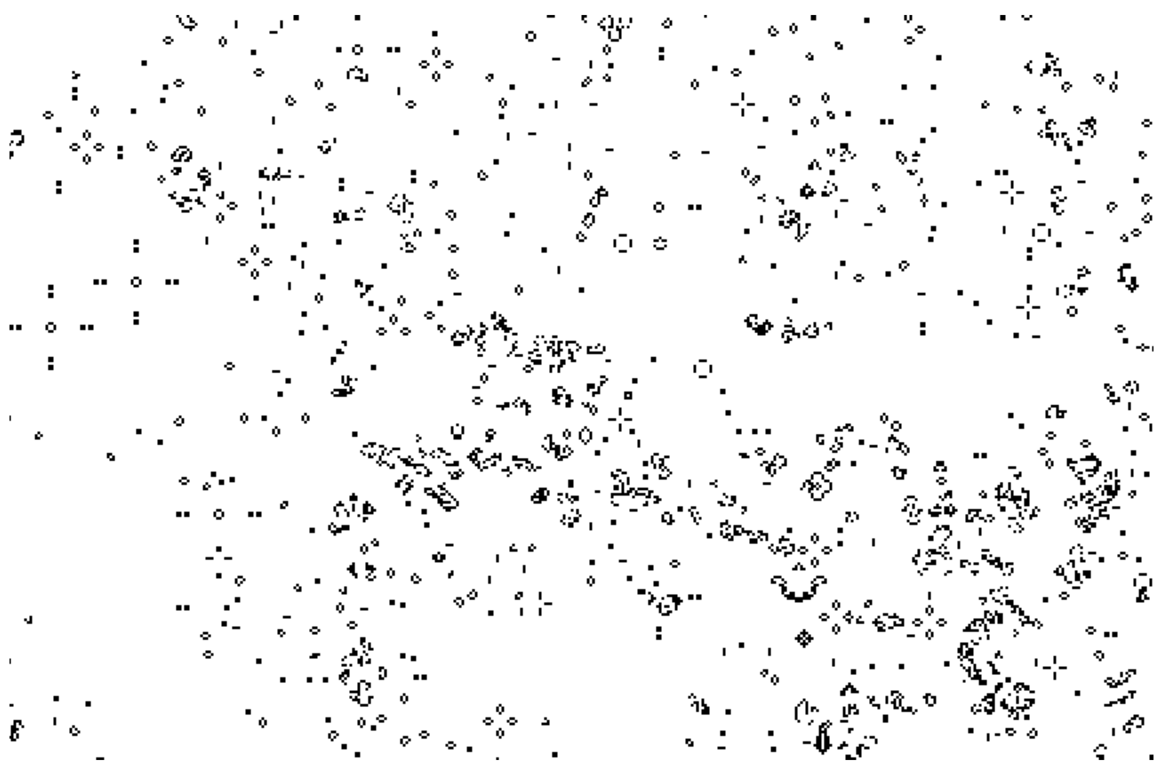

A Cellular Automaton In Action

Acousmatic sound NSSMs of cellular automata include Eduardo Reck Miranda's *Chaosynth* (1995), which uses granular synthesis and a modified CA to generate complex sound spectra. More specifically, the CA within ChaoSynth is a model of a neurophysiological phenomenon known as a 'neural reverbatory circuit', made up of a 2-dimensional grid of models of 'nerve cells'.[21] The map of this model to sound is a modified version of additive synthesis that involves many parallel oscillator banks in which the states of the nerve cells are mapped to a frequency value and oscillators are associated with a number of nerve cells. The amplitudes of these oscillators, along with a number of other sonic and CA settings, may be set by a user or performed over time through a GUI. Chaosynth is thus a NSSM in which a cellular automaton controls the overall timbre and timbral dynamics of the system, in concert with a meta-interface affording performative gesture and tweaking by the user. In addition to Chaosynth, Miranda also created CAMUS ('Cellular Automata MUSic generator'), an algorithmic composition system which uses cellular automata to drive a symbolic (MIDI) music generating process.[22] In contrast to Chaosynth, CAMUS is a NSSM within which rhythms, gestures and formal structure, rather than overall synthesized timbre, are controlled by CA.

Bill Vorn's library of external objects for Cycling 74's Max created in 1996 called *Life Tools* contains CA models for one-, two-, and three-dimensional GoL. The *Life Filter*, constructed using the *Life* object, maps cell movement to control the frequency of a 1024 band filter, which was also incorporated into the (now defunct) Cycling 74' Pluggo Harmonic Filter.[23] *Life Filter* is an acousmatic NSSM that processes (filters), rather than synthesizes, acousmatic sound, using the natural movement of a CA activity-driven filter bank to color a user-chosen sound.

1.1.3 Simulated Life: Genetic Algorithms and Artificial Neural Networks

Another class of computational algorithms is concerned with the emulation of processes of life that are less physicalized and more difficult to detect: those of evolution, via genetic (or evolutionary) algorithms (GA), and thought processes (including learning), via artificial neural networks (ANNs). GA simulates evolution by first starting with an initial population, evaluating its fitness by comparing the characteristics of the population to some characteristics criteria, simulating reproduction by culling the members of the population whose characteristics are distant from that criteria ('survival of the fittest') and introducing small mutations into the next generation of the population, and then repeating this cycle as many times as desired.[24] GA have been deployed in many NSSMs because of their ability to emulate musically relevant processes at a number of different temporal scales and for a variety of musical purposes. These include musical material and variation generation (i.e. simulated development) and interactive computer-assisted composition.

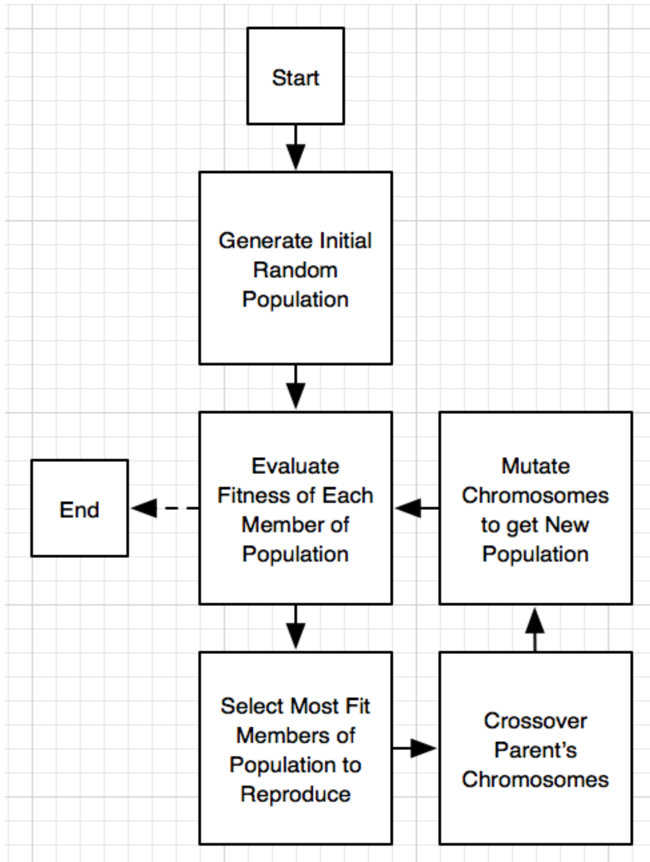

Genetic Algorithm Systemic Diagram

Cristyn Magnus explores direct application of GA to acousmatic sound materials in her *Evolutionary Musique Concrète* system (2006).[25] In this system, GA is applied directly to digitized waveforms, whose samples encode their own genotypes. More specifically, members of the population are time domain waveforms and segments of waveforms bounded by zero crossings are treated as genes. Music created by this system starts with the playback of an initial waveform population, whose fitness is measured before creating a new set of waveforms, which are played back next, who are then evaluated, mutated or replaced accordingly, and so forth. The resulting musical presentation is a NSSM whose formal structure is a product of the evolutionary process. Not only is this a fascinating acousmatic application of GA, but also acts as a clear example of a NSSM that emphasizes mimetic listening and systemic narrative while foregrounding its key.

Other acousmatic GA applications put the onus of the fitness function onto the composer's ears. Such is the case in Sonic Charge's *SynPlant* (2009), a software synthesizer that encodes different synthesizer parameters as genes.[26] Instead of creating patches through setting knobs and dials, Synplant uses a metaphor of planting seeds, which a user can choose to water or not, that grow into different plants (synthesizer timbres). This interactive process is also found in *MutaSynth* (2001), a tool for evolving the settings of programmable hardware synthesizers that was integrated into Clavia's Nord Modular as *PatchMutator* (2006). As with SynPlant, the settings of a set of control knobs are encoded as genotypes, which are then varied (mutated) and presented to the user, who listens to them and picks their favorites. Repeating this process, the system then creates new synthesizer settings from the chosen mutations, which are then auditioned and a subset of which are chosen, an iterative process which the user can repeat at liberty until ending up with a synthesizer setting which is deemed successful.[27]

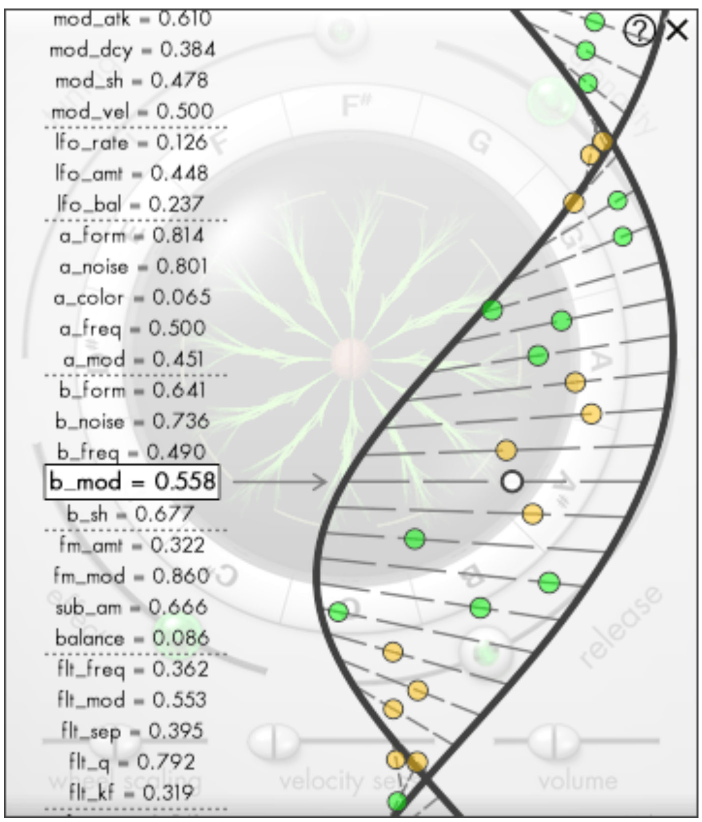

SynPlant synthesizer genotype user interface
With kind permission of Fredrik Lidström.

ANNs are simulations of adaptive, self-defining open systems. ANNs have many variations (depending on the design of their propagation and feedback networks) but in general consist of a set of artificial neurons connected together in a neural net. Some of these neurons (input neurons) take values from the outside world and others (output neurons) output. Inside the neural net, each artificial neuron receives input from some set of other neurons, accumulating their outputs through differently weighted connections.[28] Once the accumulation of input to an artificial neuron reaches a certain threshold that artificial neuron activates and sends its output to another set of neurons. Many layers of these neural nets, hierarchically combined, are often used. ANNs have been used in the context of musical genre emulation, within digital musical instruments to control and learn different parameter spaces, and in NSSMs to control various synthesis and sound processing techniques.

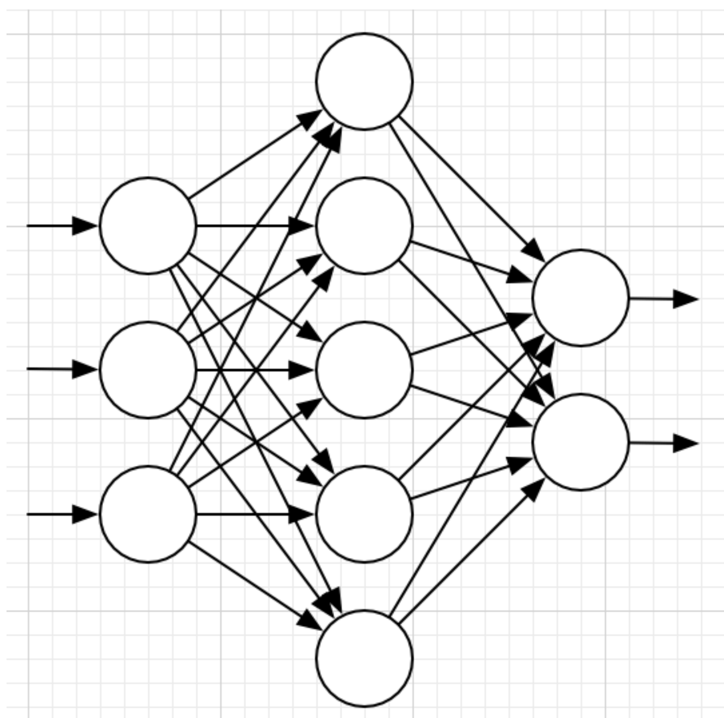

Artificial Neural Network

One such application is 'smart' gestural recognition, where ANNs are trained to recognize gestures from complex data streams in real-time. For example, in Lee and co-authors *MAXnet* project (1991), gestural data from the Max Mathews Radio Baton, the Nintendo Power Glove, and a standard MIDI controller are fed into an ANN which handles the process of gestural recognition.[29] The same concept may be found in the *Wekinator* (2010), a generalized interactive real-time software system for machine learning[30]. Here ANNs are not used to generate sound within a NSSM, but rather to inform its mapping: to learn to recognize and categorize often messy and inconsistent human gestures from digital controllers which are in turn used to generate acousmatic sound.

An example of a NSSM where ANNs are used within the sound generation process is Google's *NSynth (Neural Synthesizer) Super* (2018), a hardware synthesizer that uses ANNs to intelligently combine the timbres of instrumental samples. To create *NSynth* Google's developers first created a corpus of over 300,000 separate notes made up of 1,000 different instruments. Those notes, and more specifically the digital samples of the digital audio files (sampled at 16kHz) of those notes, were then used to train an ANN that output a multitude of transitions between the timbres of the notes, called 'embeddings'.[31] A user of *NSynth* Super can then transition between different embeddings on the fly, resulting in a NSSM with complex, dynamic real-time timbral control as a product of an acousmatic sample-driven ANN.

Another example of a project using ANNs acting on digital audio samples is *WaveNet* (2016), a deep neural network for generating raw audio applied towards creating high quality text-to-speech (TTS) systems. In addition to training ANNs on speech, the creators of *WaveNet* also trained the system on a corpus of solo piano recordings, resulting in fascinating, chimeric acousmatic sounds that range from realistic piano sounds to flurries of jumbled piano-esque timbres.[32] A last application focuses not on ANNs but on the design of a sound synthesis technique using their nuclear elements: neurons.[33] Snyder and co-authors developed a software instrument constructed by mapping the output of a computational model of a neuron directly to digital audio samples. This neuron-modeled synthesizer is thus a NSSM that affords listening to the actions of a building block of learning at the temporal scale of timbre. In addition to it being released as a plug-in for people to use, this nueron-modeled synthesizer has also been featured in the context of a laptop orchestra composition.

1.1.4 Piece-Specific Models: Bespoke Simulations and Parallel Perceptual Models

A last set of acousmatic music algorithms stems not from generalized natural computing processes but rather from system- or composition-specific methodologies. Here models are hyper-specific, focusing on a particular natural system's dynamics and/or its experience and perception by the composer.

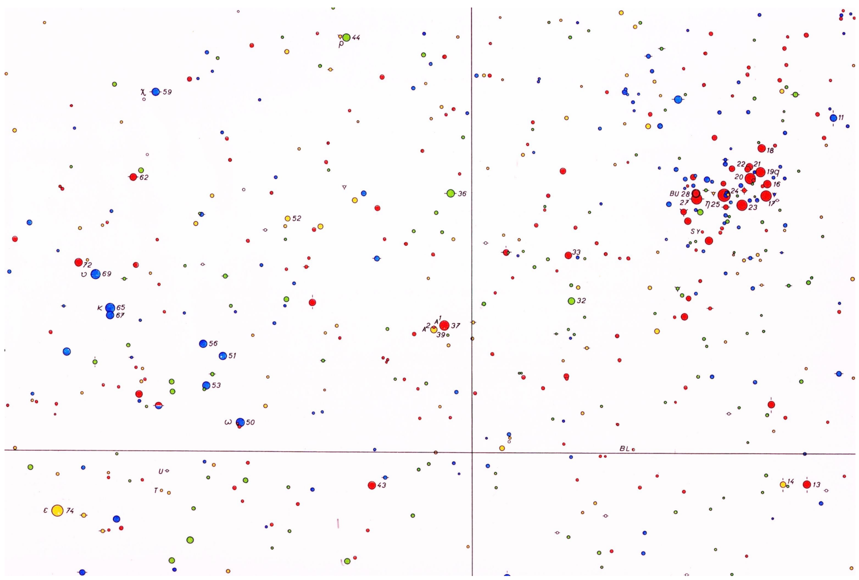

Star chart from Anton Becvar's Atlas Eclipticalis.
Sourced from the Astronomical Institute at the Slovak Academy of Sciences.

A non-acousmatic example of this type of natural system model hyper-specificity may be found in John Cage's *Atlas Eclipticalis* (1961-62), a composition for orchestra whose musical materials are derived by placing transparencies containing staves over star maps and mapping the star's positions to notes.[34] The musical composition techniques of *spectralism*, which take as musical material and organizing principle structures derived from specific sound recordings or acoustical phenomena more generally, may also be seen as engaging hyper-specific models of natural systems. For example, the opening material of spectra list composer Gérard Grisey's chamber ensemble work *Partiels* (1975) is derived from a sonogram analysis of a recording of a low E on a trombone, approximated within the ensemble to the nearest quarter tone.[35] Similarly, the work of zoomusicologist Francois Bernard-Mâche (and, more famously, Olivier

Messiaen) expands outwards from instrumental music representations of birdsong to other animal vocalizations and sounds of natural phenomenon: rain, thunder, avalanche, etc., mapping to musical contexts a variety of natural system sound recordings and experiences.

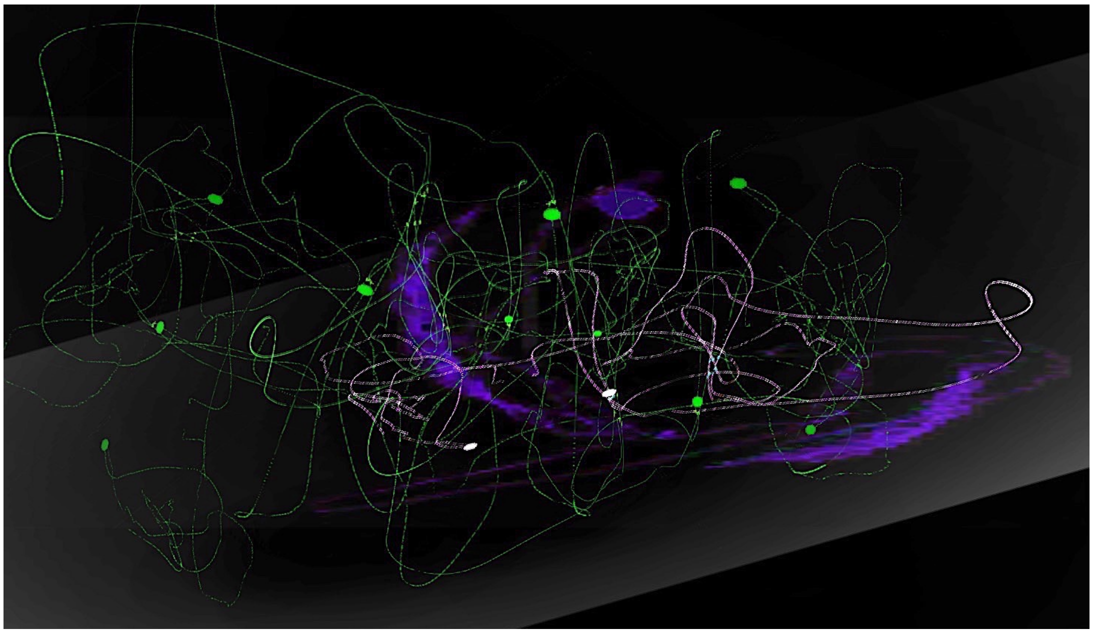

Visualization of the physical micro-movements used in *Involuntary Expression*
With kind permission of Natasha Barrett.

Examples within the acousmatic sound domain include natural system simulations deployed in the works of electroacoustic composer Natasha Barrett. For example, Barrett's 2000 composition 'In The Rain' from *Three Fictions: Northern Mix* uses a statistical simulation of rain drops falling onto a 2-dimensional surface in CSound to dictate various acousmatic sound parameters. In particular, the X coordinate of falling drops is mapped to stereo panning and the Y coordinate is mapped to simulated distance as well as pitch shift.[36] More recent works by Barrett deploy models made from the sampling of physical movements via 3D motion capture camera systems which are then mapped to sonic and spatial gestures. For example, *Involuntary Expression* (2017) first involved capturing the physical micro-movements of crowds attempting to collectively stand still, a cellist, and a drummer, and then mapping those movements to electroacoustic gestures spatialized over immersive high-density loudspeaker arrays (HDLAs).[37] Barrett's uses of trans-modal mappings to timbral and spatial parameters of sound in both of these works are examples of NSSMs that explore how the design of heightened levels of abstraction may be used within an artistic context. Put another way, the level of abstraction of the natural system models with respect to their sonic re-embodiments within these works is neither totally opaque nor totally transparent, but rather is incorporated in and, perhaps, central to Barrett's musical aesthetics.

Within a less quantitative, but still deeply related, space, the musical methodology of Hildegard Westerkamp and soundscape composers more generally engages what I will call *parallel perceptual models* of real-world sound recordings. In these works, the listener-composer's experience and extra-musical features of a soundscape function as a model parallel to acoustic models (recordings) of that soundscape. Some soundscape works, then, may be described as NSSMs that parse out the differences, similarities, and interactions between these two models (the actual and the perceived), perhaps in an active way by using filtering or other sound processing techniques to get the sound of a recording closer to its perception by the listener-composer. Ultimately, one of the purposes of soundscape composition, according to its originators, is to cause the listener-composer to re-evaluate their original, perceptual model of the soundscape. [38] [39]

A similar methodology may be found in site-specific ecological sound art, which sonically intervenes into a particular environment for the purposes of inviting audiences to re-evaluate their sonic or extra-sonic perceptions (perceptual models) of that space. For example, sound artist Graciela Muñoz's 2014 work *El Sonido Recobrado* involved recording the sound of Chile's Baker river and re-presenting those recordings through a multi-channel speaker array positioned at the bottom of a dry river bed in Petorca, her home town, whose own access to the water supply had been dammed in the 1990s.[40] In this work, the modeling of the natural system is done through recording, but another perceptual model is emphasized by drawing attention to the 'negative space' (lack of water) within the dry river bed through acousmatic sound playback. This and other NSSMs that engage natural systems in conjunction with models of natural systems are multi-layered and complex, suggesting fascinating new avenues for creative expression engaging natural systems.

1.2 Sonification + Data-Driven Music

Auditory Display (AD) may be defined as “the transformation of data relations into perceived relations in an acoustic signal for the purposes of facilitating communication or interpretation.”[1] Its genesis as a field can be traced to Sara Ann Bly’s 1982 dissertation, titled ‘Sound and Computer Information Presentation,’ which showed “how non-speech sound could be successfully used to communicate specific information.”[2] AD engages with topics of psychological research in sound perception and cognition, the development of sonification tools for research, and sonification design and application.The International Community for Auditory Display (ICAD) was formed in 1992 and hosts a conference each year showcasing research related to these pursuits.

For the purposes of this text, AD may be seen as a particular set of modeling and sound mapping strategies. As AD focuses on the ways in which sound can be used to communicate or interpret data, within these strategies the aesthetics of the sonic output (their non-semantic qualities) are secondary. For example, within AD a source of sampled data might be abstracted (e.g. through filtering or averaging) to create a set of pre-processed, mappable handles, which are then mapped to sound in such a way that optimizes a listener’s ability to receive and interpret the underlying data source. There exist several larger classes of modeling and mapping strategies in AD that may be defined by their level of modeling abstraction and mapping indexicality.

The use of a sample-based model in conjunction with a completely indexical mapping, i.e. mapping time series data directly to sound pressure levels, is called *audification*. For example, seismic data audification has been used to differentiate between events caused by bombs and by earthquakes.[3] A model whose output only indicates when that model gets into certain states (events), which are then mapped injectively to different audio files or sound synthesis gestures facilitates *earcons* or, if using real-world sound sources, *auditory icons*. For example, the sounds of personal computing systems such as emptying the trash, sending an e-mail, or entering an error state are high-level computer system events mapped to sound objects in real-time. Other modeling and mapping strategies, in particular those that involve higher levels of abstraction, deeper data-to-sound metaphors, and flexible mapping schemes fall under the category of *sonification*. Examples of sonification include data exploration (for example, the sonic navigation of brain scans mentioned in the introduction) and artistic and entertainment applications.

Within a scientific context the sonification of natural system data is particularly prevalent, as the complexity of such data may be more readily comprehensible through aural rather than visual or physical domains. This is due to the fact that our ears are incredible sensors: through sound alone we can track multiple streams of information at once, allowing sonifications to effectively represent several streams of data in parallel.[4] Our ability to detect and differentiate subtle variations in timbre and sound patterns, combined with a sonic memory that allows us to compare and contrast new sounds with old, allows sonification to be used to categorize and detect different complex phenomena, such as disease. [5] [6]

While sonification as a scientific discipline seeks to optimize data transparency, a decidedly hard systems viewpoint, artistic applications of sonification methods are often from a softer systems viewpoint, suggesting a descriptive alternative to sonification might be necessary. Carla Scaletti, a composer, software and hardware developer, and early pioneer in data sonification, uses the term *data-driven music* to differentiate her creative works that use sonification methodologies within an artistic context from scientific applications. Data-driven music thus positions the adaptation to music of data, and any accompanying information-obscuring abstraction and curation procedures, as distinct from the scientifically oriented sonification. Paul Vickers and Bennett Hogg explore the difference between sonification ('ars informatica') and sound art ('ars musica') through the development of a multi-dimensional sonification-music framework (the 'Æsthetic Perspective Space'). Within this framework, artistic sound research is positioned on a spectrum of abstract vs. concrete sound and evaluated on its relatedness to sonification, while scientific sound research is positioned on a spectrum of direct vs. metaphorical mapping and evaluated on its relatedness to music. The purpose of this research is to shed light on the inability of sonification to escape sonic aesthetics, and to call on sonification designers to develop their listening skills and draw on the skills of composers, sound designers, and recording engineers to "fully maximize the communicative potential of the auditory channel." [7] This analysis methodology, which seeks to understand a diverse range of sonification-related works, also informs the NSSM classification system described in the next chapter.

While the principle of data-to-sound-mapping inherent to sonification is fundamental to NSSMs, not all natural system sonification works are pertinent to this discussion. Projects that map to sound natural systems, but that are positioned entirely within a scientific domain, or projects that eschew the musical aesthetics of the mapped sonic materials, are of less interest here. Instead, works that privilege the sensuous experiences of models of natural systems, and their interaction with electroacoustic composition practices, will be discussed.

1.2.1 Musical Applications of Sonification

Examples of sonifications that engage NSSM concepts include Bob Sturm's *Surf Music* (2002), a sonification of ocean buoy data. This project maps the distribution of ocean buoys changing over time to the frequency domain—changes in buoy data being expressed as different timbres—and also localizes the significations of the buoys in stereo space in accordance with their physical locations.[8] This is a fascinating example of a project engaging a non one-to-one temporal and cross-domain spatial mapping. More specifically, the spatial data of ocean buoys over 26 years is mapped to an audio rate frequency domain for ten minutes, a significant temporal dilation, and the physical locations of the buoys informs the spatialization of those different shifting timbres over that duration, mapping miles of distance to the space between two speakers.

Another example of a NSSM sonification is Timothy Schmele's 2012 work on mapping brain data to 3-dimensional audio, for both artistic and scientific purposes.[9] In this work, brain activity data measured by functional Magnetic Resonance Imaging (fMRI), typically resulting in noisy data that is difficult to visualize, is sonified in a manner close to audification. The output of this project includes a number of rich-sounding binaural acousmatic compositions using high-dimensional brain data sonification as material, emphasizing, rather than attempting to reduce, the complexity of this system as a source of music, seeking to "explore the aesthetic potential of brain sonification not by transforming the data beyond the recognizable, but by presenting the data as directly as possible."

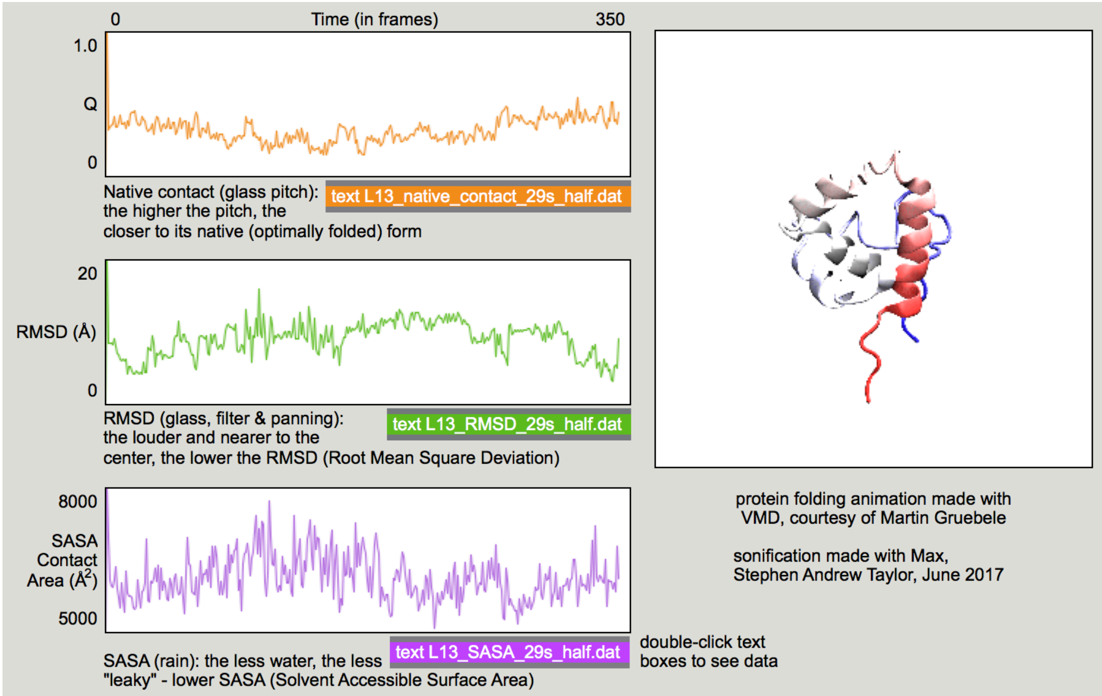

Protein Folding Sonification Visualization
With kind permission of Stephen Andrew Taylor.

Lastly, a project that engages the heuristics of a simulation and the ear's ability to distinguish multiple layers of musical information at once is Stephen Taylor's Protein Folding Sonification (2017).[10] This research exists as a Max patch that takes three heuristics of a DNA protein folding (the process by which proteins form their 3-dimensional structures) simulation in real-time and maps them to three timbrally-distinct granular synthesizer outputs. The result is a multi-layered acousmatic work that tells the story of protein folding, emphasizing mimetic listening and a systemic narrative.

1.2.2 Data-Driven Music

Outside of the context of sonification, there are a number of fascinating examples of natural system data-driven music. Contexts for these works include installation settings, music for dance, and multi-channel fixed media contexts. Further, these works might use real-time streams of (possibly local) data, emphasizing the essential tethering to time and place of natural systems.

For example, John Luther Adam's work *The Place Where You Go To Listen* makes use of a multi-dimensional set of real-time meteorological, seismological, and magnetometric data sources, along with off-line, pre-calculated data on solar and lunar cycles to inform the timbres, lights, and spatialization of an immersive installation that has been housed in the Museum of the North since 2006.[11]

The meteorological data, visibility and cloud coverage, is scraped from Internet-sourced local METAR reports from the National Oceanic and Atmospheric Administration (NOAA). The seismological data, three-axis data sampled at 60Hz, is taken from five seismometers the project was granted access to from the Alaska Earthquake Information Center. The magnetometric data—information on solar activity—is taken from five stations in the region the project was granted access to from the Alaskan Geophysical Institute. Each source of data is massaged, processed, and deployed in different ways. The seismic data is audified through the following algorithm: for each of the five seismometer inputs, X and Y channels are combined into a vector, and an FFT analysis is done on that vector. The Z axis is kept in the time domain and used to bin shift the FFT of X and Y over time. The output of this process is diffused through two subwoofer channels, placed on the sides of the room, with the five locations of the seismometers linearly displaced as phantom images along the line in between the channels. The way in which the magnetometric data is mapped is more complex. The locations of the stations are determined in relation to the installation location, and their distance, along with the height of the ionosphere taken from that data, is used to control the bandwidth and center frequency of bands of noise. These are then diffused, again being spatialized according to their real-world locations with respect to the installation location, over an array of speakers in the ceiling.The meteorological data makes no sound itself, but rather alters characteristics of sounds corresponding to the sun and the moon, namely the width of a bandpass-filtered noise ('sun band') is determined by cloud cover and the frequencies of narrower bandpass filters applied to that noise band is affected by visibility. All of the sound programming for *The Place* is done in Cycling '74's Max.[12]

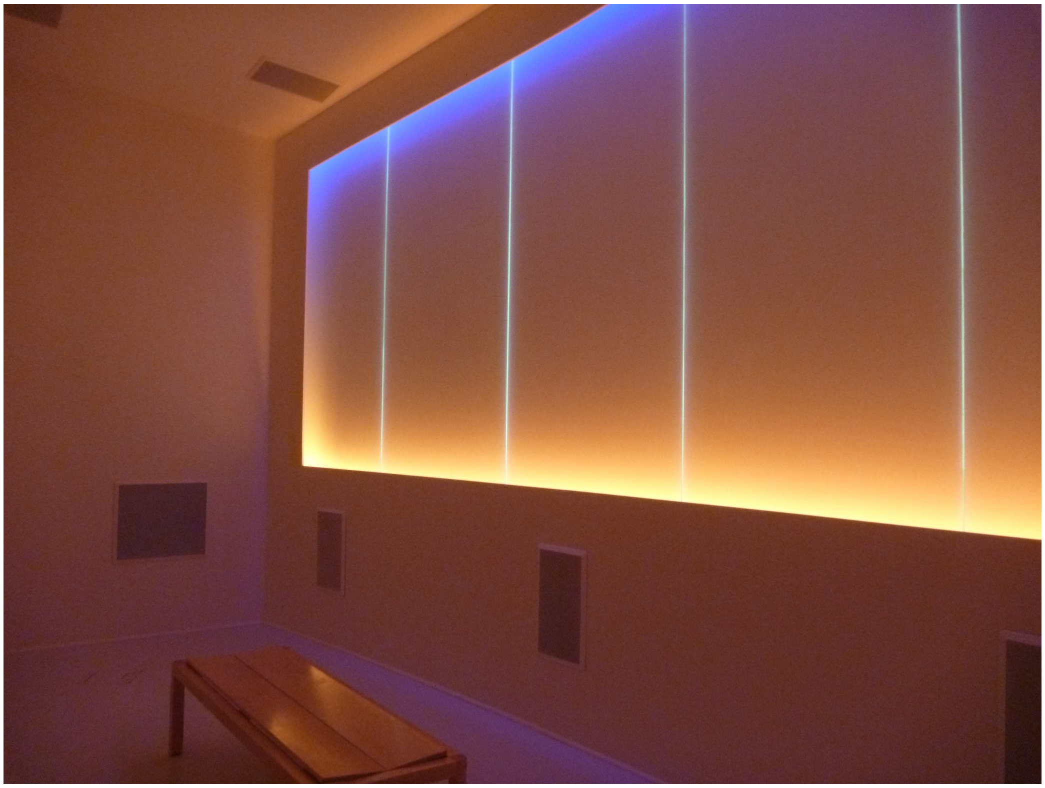

The Place Where You Go To Listen
With kind permission of Alice Doughty.

Composer Natasha Barrett's *Cheddar* system (2013) is a collaboratively developed Max patch made with geoscientist Karen Mair that sonifies geological data in real-time. A user of Cheddar can modify various temporal, spatial, and sonic parameters during a real-time audition process, potentially uncovering patterns and new interpretations of the data than might be hidden when using non-real-time, non-interactive methods. In addition, spatial information is sonified in higher order 3-dimensional ambisonics and the user can move their virtual listening position within the 3-dimensional geological data sets, adding another, spatial audio-specific, way in which different perspectives on data can assist in its interpretation. Outside of a scientific context *Cheddar* has been used as a sound generation tool by Barrett in her multi-channel acousmatic compositions as well as within *Aftershock* (2014), an interactive sound art installation that explores the environmental listening and systemic story-telling potentials that this NSSM has to offer.[13] Describing the experience of *Aftershock*, Barrett states "We are immersed in a wonderful cacophony of sound. Sustained and intermittent pings, cracks, burrs, plops and tingles jostle for position in our heads. High-pitched delicate cascades contrast starkly with deep thunder like rumbles that seem to permeate our entire bodies. We are exploring a fragmenting fault zone from the inside, a dynamic geological process brought to our ears through sonification and science–art collaboration."[14] Additionally, *Viva La Selva* (2002) is an acousmatic work by Barrett whose formal structure, events, and sound spatialization were quantitatively derived from multi-channel field recordings made in the Amazon rain forest.[15]

Lastly, within the context of electroacoustic music composition for dance, Carla Scaletti's *h->gg* (2017) is a 4-channel acousmatic work of data-driven music that takes as gestural material data from the ATLAS Experiment at CERN's Large Hadron Collider (LHC).[16] From Scaletti: "The parameters of the sounds you hear in the piece were modulated (or controlled) by variables of collision events recorded at CERN—in a sense, making the LHC the world's largest data-driven instrument." In combination with choreography, this work explores how NSSMs can engage with gestures from many different sizes and kinds of space simultaneously: from the collisions of sub-atomic space, to the gestures of spatial sound, to the relationships between sonic gestures and the gestures of dancers physically engaging with the space of the performance environment.

1.3 Live Electronic Music

Live electronic music may be roughly defined as electroacoustic performance in which the music-making (whether that be done by human, electronics, or computer) is not fixed or absolutely determined before-hand. This includes instrumental performance with electronics, the design and performance of interfaces to control electroacoustic processes, and elements of chance, chaos, and/or emergence. While live electronics has a rich pre-computational history (outlined briefly below), much of the work focused on here engages with the computer as primary electronic component.

Live electronics interacts with NSSMs in a number of ways. First, it can involve the performance of the modeling and/or mapping of natural systems in real-time. Second, it can enact a view of human performers or other living things as natural systems (with quantitative measures of them being models of those systems). And lastly, some live electronics practices (along with purely acoustic practices) engage collaborative improvisation as a kind of ecosystem model.

The history of live electronics is a history of continual development in at least two domains: the interfaces for controlling sound materials—from re-purposed knobs and switches from radio stations in the 1940s to elaborate bespoke digital musical instruments (DMIs) in the present—and the sound materials that can be controlled live—from the frequencies of analog oscillators to complex, multi-dimensional digital sound processors. Models of natural systems interact with these two developmental fronts in a variety of ways which will be outlined after a brief tracing of the history of live electronics.

Live electronics begins in earnest near the end of the 19th century, with the creation of the first electric sound synthesizer by Elisha Gray (1876). The beginning of the 20th century brought with it the invention of the Theremin (1919), a touch-less monophonic instrument for use in concert music contexts, the Ondes Martenot (1928), an electrical keyboard-theremin hybrid, and the Trautonium (1929), which used a resistor wire over a metal plate instead of a keyboard to dictate electronic pitch (famously used in Hitchcock's *The Birds*). Each of these systems, to greater or lesser extents, had form factors and interfaces that emulated or embellished acoustic instruments or, alternatively, strived to interact with a traditional instrumental music context, as was the case with the theremin.

A contrasting non-electric engagement that explored both timbres and interfaces outside of (and opposed to) traditional acoustic instruments is found in Luigi Russolo's *intonarumori* ('noise machines'), machine-like boxes built between 1910 and 1930 with cranks and various other non-standard musical instrument interfaces that produced 'noisy' sounds (roars, explosions, crashes, etc.).[1] While engaging with a view of a post-industrial soundscape that is far from what could be called 'natural' in a classical sense, Russolo's engagement with sounds and interfaces outside of the concert hall opened the doors for views of music that focused on interfaces and music-making that foregrounded systemic processes.

Instead of inventing their own electronic instruments, other composers and engineers re-purposed the old. John Cage's 1939 composition *Imaginary Landscape No. 1*, which included two turntables equipped with test tone records, asked performers to manipulate the pitch and rhythm of the tones by changing turntable speed, spinning the platter by hand, and dropping and lifting the needle, inventing the Disc Jockey as a stage performer.[2] Some of the gestures and sounds within this work are the byproduct of enabling, rather than resisting, the unfolding dynamics of physical systems (the turntable's tone arm bouncing, for example), positioning the materiality and self-reflexive cause-and-effect of physical systems, rather than their careful, sound-centered control, as musically viable.

An ethos of electroacoustic device re-purposing also encompasses a cybernetics-inspired DIY circuit-building culture that started in the 1960s with groups such as the Sonic Arts Union. Works by these artist-engineers positioned bespoke and re-purposed technology (amplifiers, custom circuits, guitar pickups, microphones) as a way to incite and/or mediate electroacoustic feedback loops. For example, in Gordon Mumma's 1967 work *Hornpipe*, the resonance of a french horn in a reverberant environment is amplified through a custom-built 'Cybernetic Console'. This concept is further explored and extended in the work of David Tudor, including his 1973 composition *Rainforest IV.* In this half performance, half installation work, an array of objects are attached to both transducers (contact loudspeakers) and contact microphones, the output of which are diffused through speakers within the performance space. The simultaneously amplified and resonated objects mediate the electroacoustic feedback loop enacted by the mic and transducer, attenuating, filtering, and otherwise transforming it over time. The result is a sonic ecosystem of zoomed-in microscopic sounding objects that deeply resonates as NSSM. In the mid-1960s, the introduction of voltage-controlled synthesizers (and later modular synthesizers that afforded live performance) also changed how electrical signals could automate analog synthesis processes.

During the 1940s and 50s tape music studios arose and groups of composers and engineers (including *musique concrète* originators Pierre Henry and Pierre Schaeffer) began exploring their use as an environment to make music with models of natural systems: real-world sound recordings. While initially stemming from a tradition of pre-meditated prescriptions of editing actions borrowed from film editing and exclusively engaging acousmatic practices, over time the studio became both a performance space and a live electronics instrument.

The studio as performance space involves viewing a composer's or engineer's performance of the multitude of different sound generation and processing devices available in the studio as a concert-without-audience that is simultaneously recorded onto fixed media. The metaphor of studio as performance space is particularly apt when performative actions are directly audible in the resultant musics. For example, the performance of sound-processing gestures in a pre-control voltage studio environment, where envelopes, gestures, and other musical actions had to be performed by hand, speaks to such a practice (for example, listen to early works by Karlheinz Stockhausen and Pauline Oliveros' early tape works).

The studio as a live electronics instrument involves taking studio equipment—tape delays, mixers, sound panners, etc.—into a performance setting and engaging with their real-time, rather than non-linear or offline, properties. Examples include Terry Riley's 1969 work *Poppy Nogood and the Phantom Band,* which used a 'time lag accumulator' made from two tape machines, looped audio tape, and a patch cord (the 'phantom band') to overdub and accumulate multiple copies of Riley's saxophone and organ playing live.[3] Stockhausen's *Mikrophonie I* (1965), which explores the live amplification, filtering, and gain control of prescribed microphone movements around a tam tam (large gong), re-frames the studio process of close mic'ing as live sonic microscopy: amplifying the complex, physical systems of the tam tam's resonance to generate a sound world that is traditionally unheard within a live context.[4] Another example that focuses on spatialization may be found in the late 1950s with the Philips Pavilion, a building installed at the 1958 World's Fair designed by Corbusier and Xenakis.[5] Within the Philips Pavilion the music of Varese and Xenakis encoded in 11 channels was played back over 450 speakers installed within the building. The materials for playback were fixed, but their diffusion over the multi-channel loudspeaker array was performed live via a mixing console, a live electronics practice of *sound diffusion* that continues to live on throughout the world today.

A 2016 Performance of Stockhausen's *Mikrophonie I* by Hélène Colombotti, Maxime Echardour, *percussion,* Ève Payeur, Vincent Leterme, *microphones,* Jean-François Charles and Stéphane Sordet, *sound processing*

The incorporation of computers into live electronic music practices began in the 1970s and posed some interesting questions and re-evaluations to composers and engineers alike. First, the continuous control of analog interfaces, and their resulting fluid expressivity, was replaced by the necessarily discrete data streams of computers (and corresponding interfaces). Within a natural systems-focused context the analog-to-digital conversion process (and its inverse) also implicitly adds at least one more layer between physicalized systems (including human performers) and their sonic output. Second, the sonic potentials of the computer as instrument (discussed briefly in the context of acousmatic music earlier in this chapter) and the procedural abilities of the computer, especially its memory, decision-making, and high-speed trans-domain mapping capabilities, became available to live electronics practitioners.

The impacts were significant and resonate to this day, with a wide variety of digital musical instruments (DMIs) being developed to engage the computer in a live electronics context. How these different systems may be described and classified has been debated significantly (see Miranda and Wanderley 2006, O'Modhrain 2011, and Birnbaum et al. 2011), but in general DMIs may be described as instruments which include a physical interface, a digital mapping layer, and sonic output. Different types of DMIs include augmented or hyper-instruments (extending an existing instrument, e.g. Yamaha's *Disklavier*), alternate controllers (which don't try to model or modify an existing instrument, e.g. Laetitia Sonami's *Lady's Glove* and *Spring Spyre*), biosignal and biofeedback interfaces (which use signals (Motion capture, EEG, EMG, etc.) derived in real-time from living things, e.g. Benjamin Knapp's *Biomuse*), as well as 'intelligent' musical instruments (which might emulate life through live natural computing algorithms or adapt their mapping or other parameters in response to a given musical context, e.g. George Lewis's *Voyager* system, discussed below).[6] In addition, software systems initiated in the 1980s such as Max and Pure Data allowed personal computer users (even those with little coding experience) to create their own real-time music processes using a node-based GUI made up of general purpose building blocks (called objects).

Different types of DMIs include augmented or hyper-instruments (extending an existing instrument, e.g. Yamaha's *Disklavier*), alternate controllers (which don't try to model or modify an existing instrument, e.g. Laetitia Sonami's *Lady's Glove* and *Spring Spyre*),

biosignal and biofeedback interfaces (which use signals (Motion capture, EEG, EMG, etc.) derived in real-time from living things, e.g. Benjamin Knapp's *Biomuse*), as well as 'intelligent' musical instruments (which might emulate life through live natural computing algorithms or adapt their mapping or other parameters in response to a given musical context, e.g. George Lewis's *Voyager* system, discussed below).[7] In addition, software systems initiated in the 1980s such as Max and Pure Data allowed personal computer users (even those with little coding experience) to create their own real-time music processes using a node-based GUI made up of general purpose building blocks (called objects).

Refining what was stated at the beginning of this section, the ethos and practice of NSSMs within live electronics is concentrated in 1) biosignal-driven instruments, which position the systemic dynamics of the human body (or other living systems) as data to drive musical processes, within 2) live performance of natural computing systems, which explore how a computational model's development over time can function in/as a musical performance, and 3) 'intelligent' DMIs, which might co-exist and be responsive in collaborative improvisational contexts, questioning the musical implications of human-computer interaction.

In particular, this last class of modern live electronics systems, which engage computational intelligence in the context of collaborative improvisation, is worth further exploration. As defined in the exposition, a complex system is a set of interconnected parts which function together as a complex whole, whose dynamics are not apparent from an analysis of its elements in isolation. A group of improvising musicians with or without live electronics, particularly within 'freer' contexts, fulfills this definition: all of the performers on stage are both individual sound producers but also strive to engage a 'group mind' or become a 'superorganism' during performance, working together towards a common goal of musical creation.[8]

Through the language of neo-cybernetics, Edgard Landgraf describes group improvisation in terms of contingencies, emergent behaviors, and as a reflexive system (one that is aware of its own processes).[9] With this systems-focused lens, then, collaborative improvisation might be viewed as a simulation of an ecosystem, the resultant behavior and dynamics of this ecosystem more or less sonified as the resultant improvised music. The inclusion of live electronics within a collaborative improvisation setting has varied impacts: the electronics might enhance an instrumental performer's sound palette (as is the case with hyper-instruments), they might generate sound on their own through live acousmatic sound generation, or they might act as an independent musical voice, responsive to the current musical contexts, as in the aforementioned *Voyager* project by George Lewis.

Voyager (1987) is an interactive computer music system consisting of 64 asynchronously operating MIDI-driven 'players' that iteratively reorganizes its behavior in response to the current acoustic musical setting, ranging from "complete communion to utter indifference."[10] The encoded dynamics and behavioral shifts of *Voyager* emulate Lewis's own improvisational aesthetics, activating the computer as an included participant within a particular collaborative live music-making practice. Viewing *Voyager* as a NSSM takes collaborative improvisation as a natural system and its computational simulation as natural system model, with inputs from the acoustic domain (through real-time sound analysis) acting as input to the simulation and the output of the simulation being mapped to acousmatic sound generators. This view of improvisation as ecosystem is also exemplified in other, purely acoustic works of Lewis, including *Artificial Life 2007*, which encodes environmental responsiveness and interactivity through explicit directives and sets of situation-specific possibilities that performers navigate during a realization of the work.

Having outlined the history of live electronics and highlighting parts of it that resonate with NSSM concepts along the way, including a view of improvisation as ecosystemic simulation, specific examples of live electronics will now be discussed in more detail.

1.3.1 Performing the Modeling and Mapping of Natural Systems in Real-Time

Real-time performance of natural system sound model interfaces come in all shapes and forms, ranging from purely acousmatic systems to instrument and electronics works. Examples include the first application of swarm intelligence to music, Tim Blackwell's *Swarmusic* (2002), an interactive music improviser in which a swarm algorithm generates musical material by mapping 3-dimensional agent positions onto events (described through loudness, pulse (repeat rate), and pitch) in MIDI space. Real-time interaction with an external musical source (such as an instrument or vocalist) are encoded as attractions of the particle swarm to targets within 3-dimensional space (for example, if a singer sings a loud middle C, the swarm will be attracted to the corresponding location within 3-dimensional space). Swarmusic is a self-organizing live electronics NSSM in which a real-time flocking algorithm, potentially altered by outside acoustic sources, is mapped to musical space.[11]

Al Bile's *GenJam (Genetic Jammer)* (1994) is an interactive GA that learns to improvise jazz. Like *PatchMutator*, this system involves the live performer ('mentor') functioning as a fitness function for the GA.[12] Before a performance, as *GenJam* improvises over the chord progression of a song the mentor can effectively thumbs up or thumbs down certain musical phrases. A new population of phrases is then generated by culling the least-fit phrases and keeping the most-fit phrases, while also mutating (rotating, inverting, transposing, etc.) a small probability of them. After iterating this process many times, the surviving musical phrases are then called upon in real-time during performance. Additionally, *GenJam* has been extended to 'listen' to Biles' own trumpet-playing in real-time and respond: varying the recorded input by taking advantage of certain mutations in tandem with a song's chord progression.

A last example involves the visualization of data being used as a graphic score for performance, a kind of human-mediated live electronics sonification. In Thomas Rex Beverly's work *Telepresent Storm: Rita* (2013), historical meteorological data (including wind speed, temperature, and barometric pressure) from Hurricane Rita (2005) is used to create a real-time graphical score which is then interpreted live by the performer, who triggers and processes electroacoustic sounds of wind, rain, and destruction using 2 iPads as control interfaces. The graphic score interpreted by the performer is projected onto a screen behind them, providing the audience with a key to, if not the mapping itself, the human-mediated mapping source.[13]

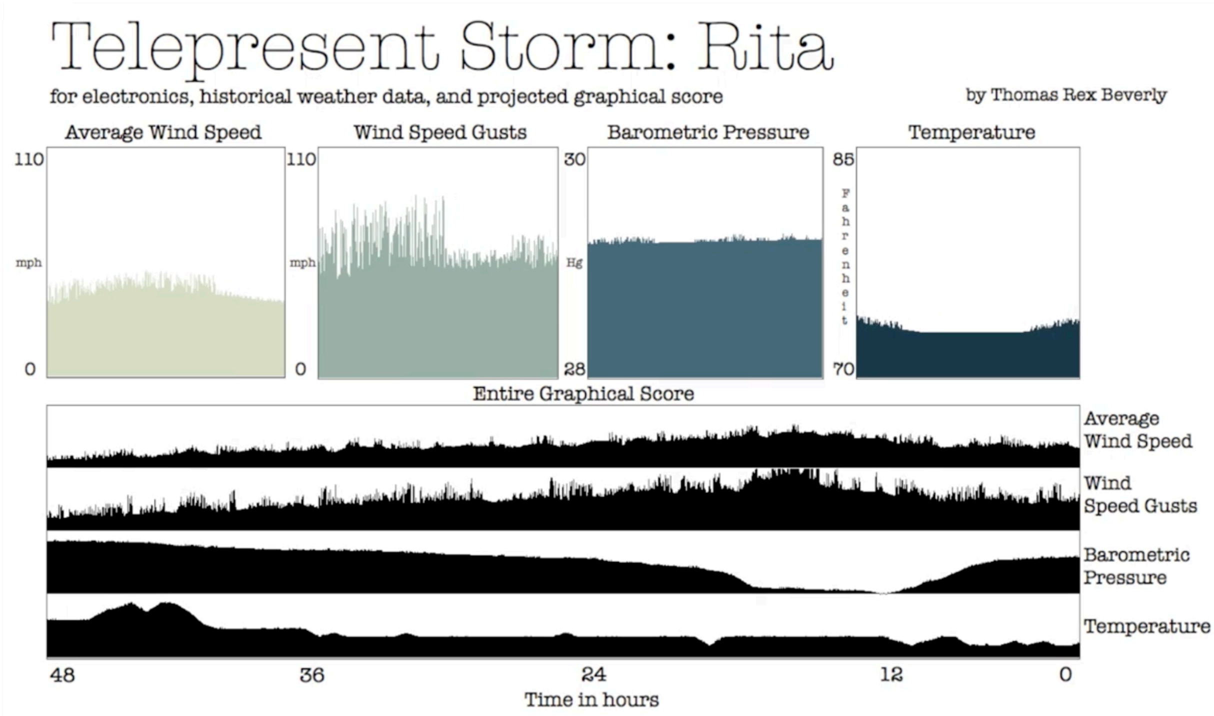

Graphical score of Telepresent Storm: Rita
With kind permission of Thomas Rex Beverly.

1.3.2 Amplifying the (Unheard) Activity of Natural Systems

Other live electronics contexts for NSSMs engage the sonic properties of natural systems present within the performance space: the acoustics of the venue (a la Lucier's *I am sitting in a room* (1969) or Mumma's *Hornpipe* (1967)), the sound mapping of data from performer's bodies, or real-time sonification of other living systems (plants and animals).

These include Agostino Di Scipio's *Audible Eco-Systemic Interfaces (AESI)* project (20030, a unique software and hardware system that positions a performer within a complex digitally-mediated human-machine-environment feedback loop. Tiny gestures, or even just the changing resonance of a performer's instrument, are amplified via microphone and digitally processed before being sent to loudspeakers. In parallel, another microphone (or set of microphones) is used to listen to the effects of this amplification process and modify the parameters of the digital processing being applied accordingly. The result is a NSSM that could be described as a computer-mediated electroacoustic ecosystem that uses sound as an interface to control the system's adaptation and evolution over time. In addition, *AESI* gets at the heart of live electronics' ability to engage with site-specificity and the agency of space: each room this series of works is performed in significantly affects their sound worlds and formal trajectories.[14]

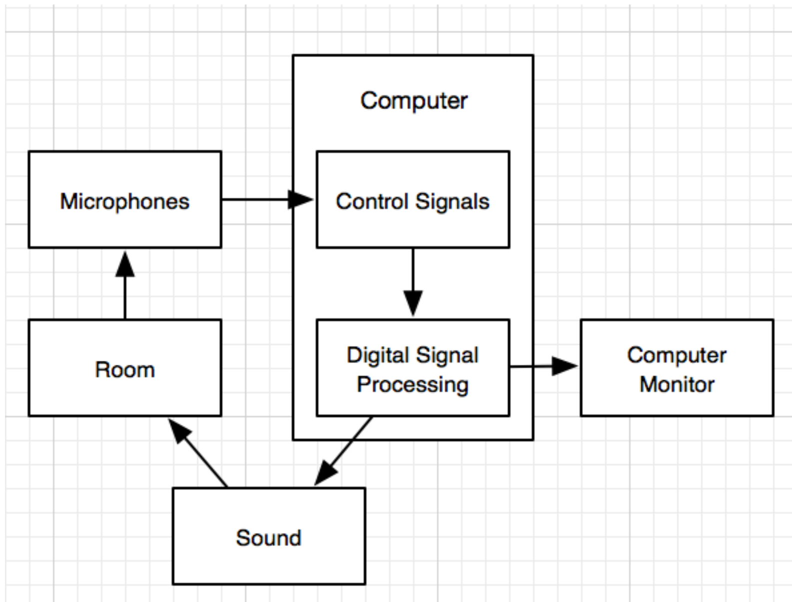

Circular Flowchart of Di Scipio's *Audio Eco-Systemic Interfaces*,
after Meric and Solomon

Biosignal DMIs such as those used in Atau Tanaka's *Myogram* (2015) also fall into this category. *Myogram* uses two Myo sensors (computer control arm-bands that contain multiple electromyographs, accelerometers, gyroscopes, and other sensors) to convert the performer's muscle tension and relaxation and arm movement into data which drives real-time electroacoustic sound processing.[15]

The experiments of Thomas Shannon and John Lifton in the 1960s, with living plants acting as electric pickups, fall into this category.[16] This practice has been updated with digital technologies in the work of Mileece Petre, among others, who maps the electromagnetic current of plants to musical notes using custom software[17], and with *MIDI Sprout*, a portable plant-to-MIDI sound mapper crowdfunded in 2014.[18] The activity of animals is electroacoustically amplified in Céleste Boursier-Mougenot's *from here to ear (v.15)* (2011), an installation which places 14 differently tuned electric guitars within a makeshift aviary containing 70 zebra finches, whose landings and peckings on the instruments are electroacoustically amplified.[19]

The work of Eduardo Miranda with the slime mold *Physarum polycephalum* in the contexts of sound synthesis and live performance, including a live piano duet between sound generated by the mold and a pianist in *Biocomputer Music* (2015), is also an intriguing real-time application of the organic computation sub-discipline of natural computing.[20]

1.3.3 Collaborative (Electroacoustic) Improvisation as Ecosystemic Simulation

Multi-agent live performance, particularly in the context of improvisation, was discussed as modeling the emergent behaviors and complexities of ecosystems. Examples include the aforementioned *Voyager*, and in a (possibly) acoustic context the complex game-based music of John Zorn and the anarchistic control structures within Cornelis Cardew's Scratch Orchestras. [21] [22]

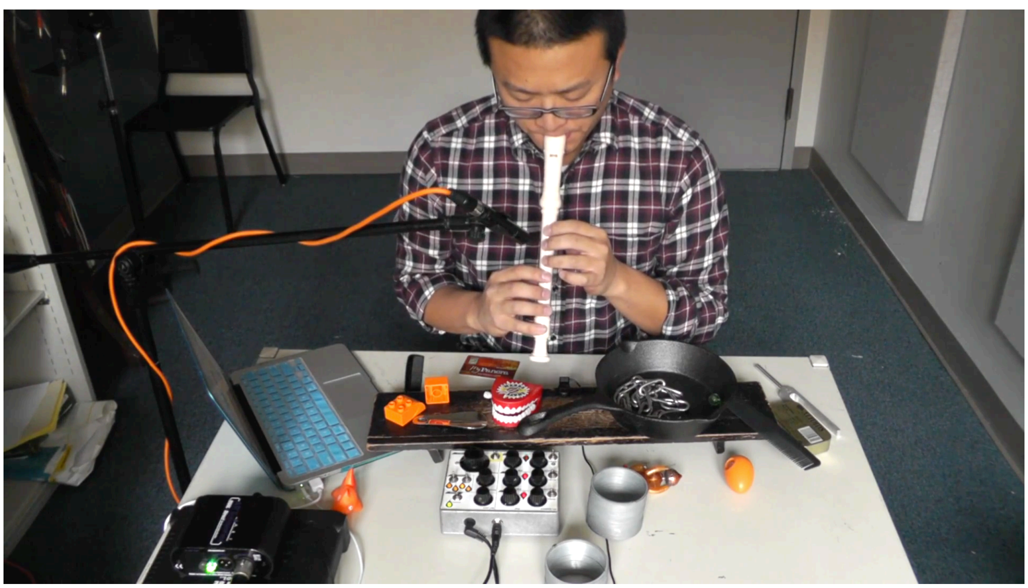

Joo Won Park performing *Touch*
With kind permission of Joo Won Park.

Other works might have more non-human improvisors than human improvisors (perhaps only one human, acting as a a kind of conductor). These works explore the (sonic) agency of objects through a lens of *materialism*.[23] These include the ‘object performances’ of Rie Nakajima, complex electroacoustic ecosystems that Nakajima sets into motion in art galleries, and the playful found object concert works of Joo Won Park, which position various toys, noise-makers, and other found objects as performers within an amplified electroacoustic group improvisation.

These types of works may also exist in an installation context, such as in *failed experiments* (2017) by Anne-F Jacques and Ryoko Akama, a work in which a room is filled with a set of isolated, precarious electroacoustic audiovisual devices, which hum, chirp, scrape, rattle, and resonate to produce an immersive ecosystem of electroacoustic sound.[24]

1.4 Software Art

Software art has come to define two distinct, but deeply related, aesthetic engagements with computers. The first is the design of software to enact art, what I will call *computer art*: the computer artist programs code to produce an aesthetic experience or product. This may be *interactive art* (affording dynamic alteration by humans) and/or *generative art* (autonomously producing a set of materials). The second definition takes as artistic material code, views code itself as aesthetic object, what I will call *code art*. This definition of software art requires a shift of both an audience's and a software artist's focus from visual, acoustic, and tactile digital media to the structures of programming, the digital systems and processes behind their human-comprehensible outputs.[1] It also has the ability, according to software artist Casey Reas, to "comment on our increasingly digital social and political structures and to challenge the underlying formal assumptions of computer code."[2]

These and other philosophies align code artists closely with Free/Libre Open Source Software (FLOSS) communities, as well as various software art offshoots: hacktivism, Net art, live coding, and the demoscene, an artistic subculture that arose from the creation of audio-visual presentations to accompany cracked software ('cracktros').[3] In addition, code art also points away from the digital computer as a necessary element for enacting the structures of programming: deployments of algorithms and the design of algorithmic interactions using purely analog means are also within the purview of this practice. For example, Fluxus artist La Monte Young's *Composition 1961 No. I, January I*, a performance score prescribing the performer to "Draw a straight line and follow it" resonates as algorithmic procedure, along with the chance- and loop-based acoustic music practices in Terry Riley's *In C* and Steven Reich's *It's Gonna Rain*, respectively.[4]

While distinct, as code is both a language in and of itself and a set of pointers to formalized procedures, there is necessarily a bridge between computer art and code art. Both are born from computational recipes devised by humans, the former with the instructions pointing outwards to some externalized human-perceptible art product and the latter pointing inwards to itself, engaging the aesthetics of code (as poetry, as rhetoric, as style, as attitude) and their resultant processes and properties separate from any external experience.[5] The development of the NSSM framework is indebted to tenants of both definitions of software art. Code art's interrogation of computer systems and tools, pointing towards a techno-creative practice that emphasizes customization, openness, and non-linearity, informs the privileging of systemic interconnectivity within the NSSM concept. Computer art's engagement with system models (including natural system models) as artistic agents within a technology-driven art production environment informs the author's own practice of real-time and offline engagement with natural system models in electroacoustic music composition. Rather than exploring how an artist can produce a formal work that encapsulate the beauty or sublimity of our natural environment (as in some traditional media), natural system computer art explores how artists can create processual works (in the form of code) which actively trace the behaviors of natural systems.

Undoubtedly, natural computing is at the technical core of computer art's engagement with natural system modeling, and has a long history with the medium. Some of the earliest instances of software art were born from physical system models, namely ballistic simulations. The 1963 artwork *Splatter Pattern* created by the United States Army Ballistic Research Laboratories was judged to be the best submitted artwork within a competition held by the magazine *Computers and Automation*.[6] Its creators stated that rather than computer art it was "merely an aesthetic by-product" of utilitarian pursuits.[7] Since then, computer art has stayed abreast of natural computing research, with a diverse array of communities continually exploiting natural system model algorithms within interactive and generative art contexts.

Particular focus is made on agent-based systems, cellular automata, and L-systems, algorithms with intuitive maps to visual (and, to a lesser extent, sonic) domains. Communities that engage with these algorithms include, of course, computer musicians and algorithm composers, along with industrial designers and architects, commercial and experimental computer graphics designers and animators (including special interest groups such as ACM's SIGGRAPH), and by programmers in the demoscene and video jockeys (VJs).[8] A complete history of software art and the use of natural system models within software art is neither constructive nor necessary here. For more history of software art see Grant Taylor's *When The Machine Made Art*, Golan Levin's *Audiovisual Software Art: A Partial History*, and Matthew Fuller's *Behind the Blip: Essays on the Culture of Software*. For a history of natural system models within the context of software art, see Mitchell Whitelaw's *Metacreation: Art and Artificial Life* and Daniel Shiffman's *The Nature of Code*.

As was mentioned earlier, code art differs from most traditional media in that code artworks define processes rather than forms, but the technological engagement of code art also diverges from scientific contexts as well. In opposition to traditional natural computing engagements with natural system models, code artists may interact with computational models of natural systems as if they were native to software. For example, a flocking system may not be optimized to be as realistic as possible, to best represent the natural system from whence it came, but rather will be extended, abstracted, and transformed (possibly beyond recognition) as computational material within the medium of software. The result of this process is a systemic caricature, defined in the exposition as an intentionally un-optimized or distorted model of a system to highlight or critique a particular quality of the system. Computer artists may focus on the aesthetics afforded by these systemic caricatures, whereas code artists might use the comprehensible delta between a realistic natural system and its caricature to open up discussions around, perhaps, the changing natural landscape, our place within it, and the role of technology within it (a kind of 'speculative ecological software art').

Mitchell Whitelaw expounds upon the impact of systemic caricatures both within computer art and code art in his article 'System Stories and Model Worlds: A Critical Approach to Generative Art' (2005). This article suggests that artistic software applications that engage system models reveal their implicit 'system stories'. Whitelaw calls for a decoding of the narratives and ontologies inherent in these system stories, and suggests new system models that engage what he calls 'critical generativity'. Critical generative art sketches possible worlds, imaginations of the systems we live in, 'revolutions cast in software'. Whitelaw states that these critical software systems are particularly useful because they are processual rather than formal: they can actually experiment with the emergent outcomes of speculative ontologies, modes of being, and relation. Whitelaw calls for an eschewing of generative art as a practice of reproducing known features (in our case, through optimized sampling or simulation) and feeding these into what he calls 'eye-candy' machines (or in our case, 'ear-candy' machines), but rather to construct systems that generate prospective or utopian ontologies that “might be equally ironic, critical, deconstructive, or fantastic.”[9]

Matthew Fuller, one of the co-founders of the field of software studies, describes these types of software systems simply as ‘critical,’ defining themselves in relation or opposition to pre-existing softwares (in this case, industrial natural computing applications).[10] One modern, electronic music-related example of this is *Argeïphontes Lyre*, a critical software produced by Akira Rabelais that subverts a traditionally explicit, user-friendly, and controllable sound, image, and text processing software into one that is an aesthetic (and potentially frustrating) experience to use.[11]

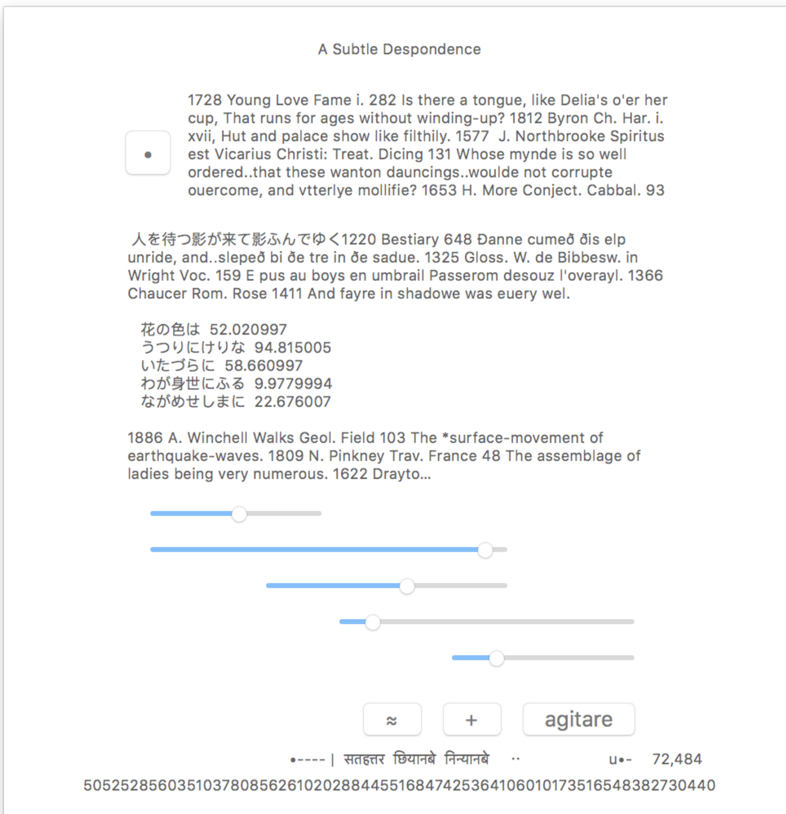

User interface for an audio processing module of *Argeïphontes Lyre*
With kind permission of Akira Rabelais.

Another example of critical software might be a sonification system that, rather than limiting the user to a curated set of mappings through a tightly controlled interface, opens up the mapping floodgates and implements no restrictions to its interface, such as in the author’s own *HabiSpat* (described in the penultimate chapter), where a user can populate an imaginary sonic landscape with any number and distribution of sonic agents rather than just having a fixed, prescribed set of realistic presets. The result is a software whose affordances extend beyond a traditional use case and into exploratory areas of high abstraction and low indexicality. Going past critique, if a software is so radically transformative as to suggest a new type of software or hypothetical computation, it may be defined as ‘speculative’.[12]

The computational modeling of natural systems within an artistic context, realistic or otherwise, also gets to the heart of some traditional questions posed when considering the impacts of computer art. These include questions of originality and creativity: whether or not computers can generate anything novel, or put another way, whether or not the output of artistic software is always a regurgitation or re-mixing of the human programmer’s interactions with the computer.[13] Within the context of natural system models (and other systems with sufficient complexity) the computer acts as a synthetic catalyst for emergent and chaotic dynamics, generating materials which, if not new, are by definition not completely predictable before-hand. This exemplifies one of the unique affordances of software art: the ability of computers to simulate complex systems and to map them to visuals, sound, or other human-perceptible outputs in ways that would be much more difficult (or impossible) if done by other means.

This also speaks to classical ideals of art mimicking nature and, perhaps, a core ethos of this text. Within generative software art nature is mimicked by 1) a necessarily quantitative modeling procedure and 2) as a process rather than form. Quantitative modeling, always being an approximation, inherently points towards the unreachable complexity of natural systems, enhancing their awe-inspiring properties. Computational mimicry via process rather than form mirrors a modern view of natural systems: not as static forms but rather as ongoing processes, constantly changing in response to their environments.[14]

1.4.1 Natural System Computer Art

Many of the software systems discussed earlier to generate acousmatic sound, particularly those that are primarily autonomous, may be considered sound-based generative computer art. What follows are a number of graphical generative computer art systems that engage with NSSM concepts.

Casey Reas is a generative software artist and co-creator of the Processing computer language whose work *The Process Compendium 2004-2010* is a series of computer generated images and documentation videos which have been deployed in a variety of contexts. These include purely digital (web-based) contexts, installations (such *TI*, where animations from the compendium are projected onto sculptural elements within a room), and even live performances (as in *Images for 18 Musicians* (2004), visual accompaniment to Steve Reich's *Music for 18 Musicians* (1976)).[15] On *The Process Compendium,* Rease states "During the last seven years, I have continuously refined a system of Forms, Behaviors, Elements, and Processes. The phenomenon of emergence is the core of the exploration and each artwork builds on previous works and informs the next. The system is idiosyncratic and pseudo-scientific, containing references ranging from the history of mathematics to the generation of artificial life."[16] For example, Element 1 of this series is defined through the following pseudo-code:

```
Form 1: Circle
Behavior 1: Move in a straight line
Behavior 2: Constrain to surface
Behavior 3: Change direction while touching another Element
Behavior 4: Move away from an overlapping Element
```

This pseudo-code is then actualized through real code (in the Processing programming language or another similar graphics-oriented coding environment) to define a work of generative software art. This work might have randomized elements on startup (in this case the size and locations of circles), but after that initial injection of randomness, the explicitly-defined Behaviors take control and the emergent properties of the system unfold over time as frames of animation, resulting in tracings of circle locations in different colors and opacities.

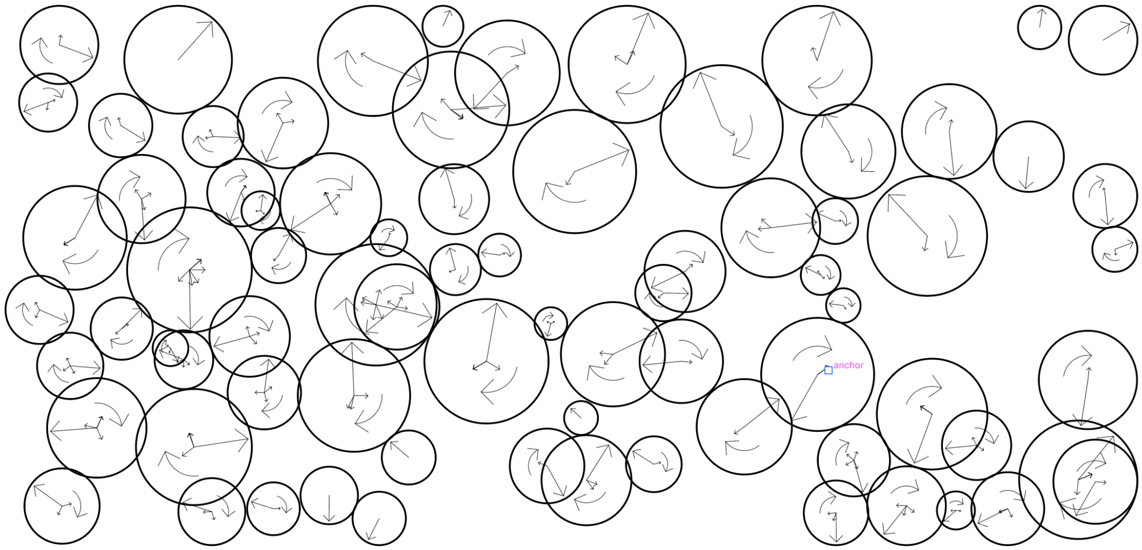

Element 1, Casey Reas
With kind permission of Casey Reas.

Jared Tarbell, another Processing user, computer artist, and co-founder of Etsy, directly engages natural systems in his generative software art. His various code repositories and digital galleries (from the late 2000s to the the early 2010s at complexification.net and currently at infinite.center) showcase a wide array of biologically-inspired, natural computing-derived, and otherwise exploratory generative code art. In addition to providing graphical outputs and interactive web-based versions of these softwares, Tarbell also often provides his code, making much of his software art FLOSS. A representative example of Tarbell's work is *Substrate* (2003), which the artist describes as follows: "Lines likes crystals grow on a computational substrate. A simple perpendicular growth rule creates intricate city-like structures. The simple rule, the complex results, the enormous potential for modification; this has got to be one of my all time favorite self-discovered algorithms."[17] Other works such as *Happy Place* (2004) explore multi-agent pseudo-social structures through the following rules:

```
A. Move close to friends but no closer than some minimum distance.
B. Distance self from non-friends as reasonably as possible.
```

The resultant patterns and shapes have large-scale movement and micro-scale chaos, ultimately telling stories of digital pixel 'friends' and 'non-friends' via computer graphics.[18]The artist applies different color palettes to his works in a variety of ways, when not algorithmically defined: either extracted from famous art (e.g. Jackson Pollock) or derived from pictures of landscapes the artist has taken, a process of mapping from a natural system visual model (picture) to a digital color palette. The way in which these colors often permeate Tarbell's creations, a soft, translucent, grainy shading called 'Sand Stroke', is also derived from a natural, emergent process: that of emulating the movement of sand in the desert (perhaps around Tarbell's locale at the time in Albuquerque, New Mexico).[19]

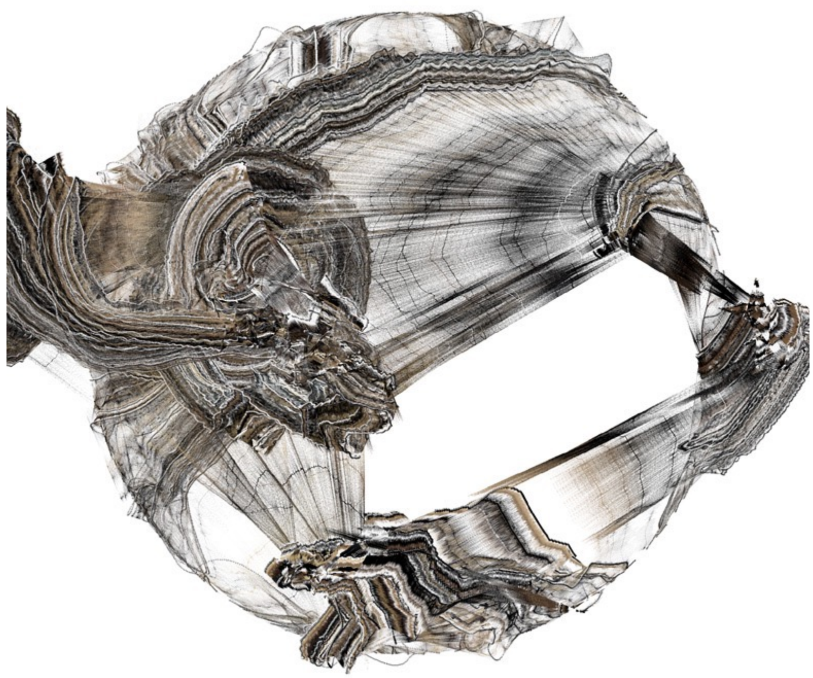

happy place with 500 friends, Jared Tarbell
With kind permission of Jared Tarbell.

Erwin Driessens and Maria Verstappen have been active as collaborators since the 1990s. They attempt an art in which spontaneous phenomena are created systematically, works not entirely determined by the subjective choices of a human being, but instead generated by self-organizing processes. They work with both physical (and organic) media and develop software to simulate artificial growth and evolution.[20] One example of their work that also engages with materialization is

Accretor (2012), a computer program and subsequent construction process that is based on the accretion of particles (in 3D space, 'voxels') to naturally 'grow' sculptures. After an initial random seed of a few voxels is placed within 3D space, rules comparable to CA transition rules decide if a new voxel can be added to a certain place on the surface of the existing form. Many rules are possible, but not all rules will create a coherent structure, a form that can be materialized, and thus Driessens and Verstappen constrain the system to only apply rules that yield materializable forms. These forms are then 3D printed with acrylic resin, the resultant objects being approximately 6 x 6 x 6 inches and consisting of 8 to 12.5 million particles each.[21]

Another example from Driessens and Verstappen is *E-volved Cultures 360°* (2016), an immersive installation created for circular projection screen (installed in Amsterdam's Kunstkapel Zuidas).The warped naturalistic pixel-landscape of this work is generated by virtual agents that leave visual traces of their interaction with their environment. The agents have been bred with the *E-volver* software, custom software developed by Driessens and Verstappen used in a number of their works, that applies artificial GA to the generation of complex 2-dimensional images.[22]

E-volved Cultures 360°, circular projection
The Generative Cinema, Kunstkapel Zuidas, Amsterdam, 2016
With kind permission of Maria Verstappen.

2. Evaluation Framework

Having explored a wide range of domains and examples of works that relate to NSSMs, in what ways can we describe and compare them? This chapter first discusses how the production of NSSMs may be positioned within one or more creative spaces, presents a multi-dimensional taxonomy that defines a NSSM based on what audio technologies and modeling techniques it uses, along with its level of abstraction, and then positions works that engage with the NSSM concept into this framework.

2.1 NSSM Creative Spaces

From a creative (*poietic*) perspective, a number of different modes of engagement, strategies, or procedural domains are at play when developing NSSMs. These range from the aesthetic: making perception-based decisions on the curation and abstraction of natural system model derived musical materials, to the technical: making a quantitative connection between a natural system model and a sound producer, to the practical: necessarily needing to boost the amplitude or change the timescale of musical materials or model handles. These modeling, composition, and mapping procedures take place within three potential *creative spaces*, each of which (barring the first) may or may not be present within the NSSM creation process:

1. *Physical Space*, where the physical natural system exists and the sonic output coexists
2. *Compositional Space*, where the creator makes decisions and selections on the modeling of the natural system and the medium and aesthetics of the sonic output
3. *Mapping Space*, a space within Compositional Space, where handles from the natural system model and the producer of the sonic output are mapped to one another

Natural systems, whether they are living or non-living, are necessarily in Physical Space. Music exists as vibrations across a medium, often air, and must also exist in Physical Space, regardless of its source (instrument, loudspeaker, etc.). Compositional

Space is where the modeler (possibly the composer) simulates or samples a model of the physical natural system, and also where the composer constructs the musical presentation of the natural system model-derived musical materials, through processes of mapping, curation, and/or abstraction. This process of constructing the musical presentation is always in conversation with its perception (both by the composer and its intended/hypothetical audience) within physical (sonic) space.

Entry into Mapping Space is entirely optional, and requires a shared layer, where handles of the modeled natural system and destinations to the musical materials generator coexist at the same logical level. In the case of digital NSSMs, this could be done by converting both to digital data (time series, functions, etc.), harnessing the source agnosticism of computers. Mapping Space's inclusion in Compositional Space is a function of the role of mapping in the composition process: Mapping Space might be 'entered' as part of a compositional decision, a desire to facilitate a direct indexical relationship between model handles and sonic output. As mentioned in the Exposition, mapping may require an iterative sequence of trial-and-error, a loop of execution (choosing a mapping) and evaluation (listening to its results within Physical Space) that may be done many times as part of the compositional process. This mapping may also change dynamically over the course of the NSSM's musical presentation, either through pre-programmed automation, self-adaptation, or performance interface (Figure 19).

Figure 19

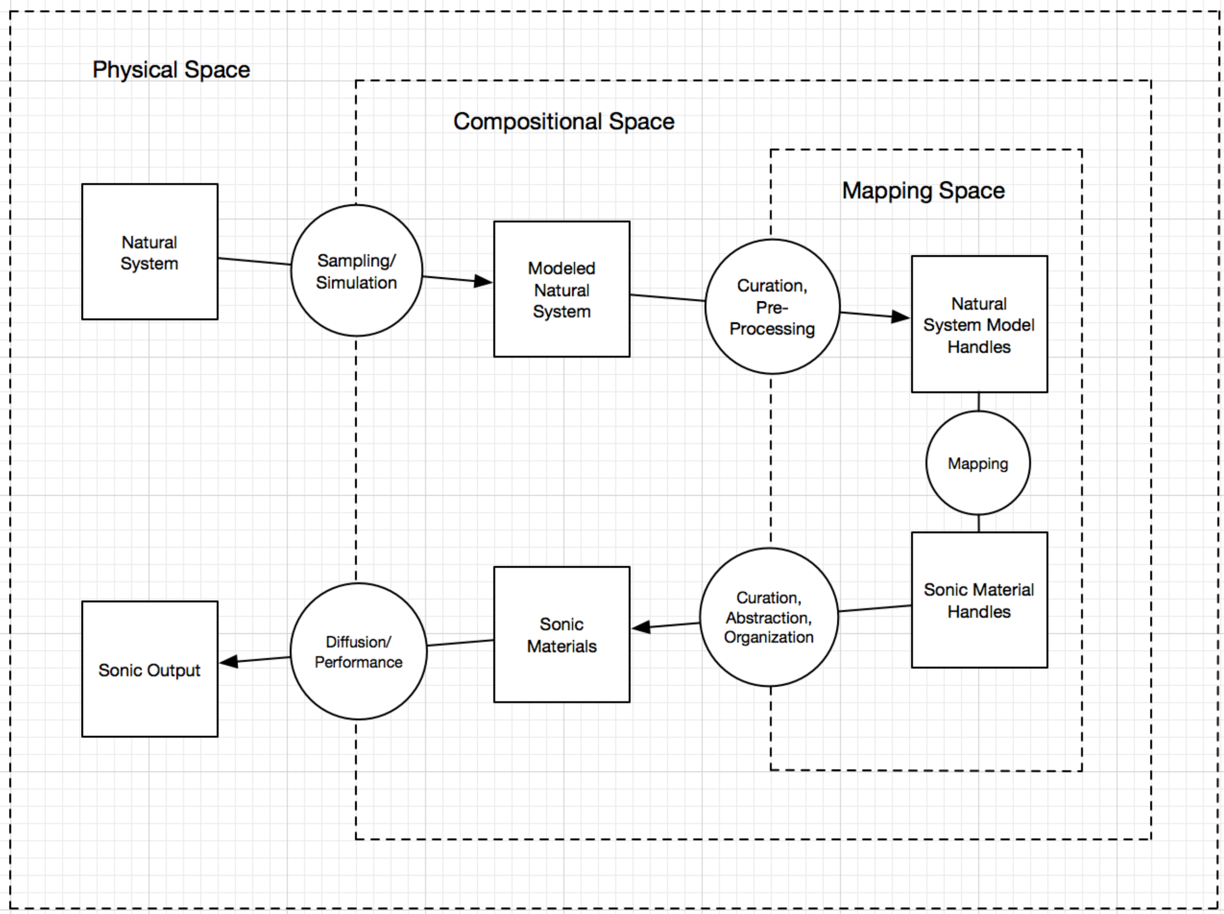

A generalized overview of NSSM creative spaces and the actions available to transition between each

The configuration of these spaces in different NSSMs is highly differentiated. Within the context of a site-specific ecological sound art work creation happens almost entirely within Physical Space, engaging Compositional Space merely to push the natural system to a more 'vocal' position. For example, the Wave Organ is a completely acoustic NSSM created by Peter Richards and George Gonzalez that is installed in the San Francisco Bay. The Wave Organ consists of 25 organ pipes made of PVC and concrete located at various elevations within the Bay which produce sound from the impact of waves against the pipe ends and from the movement of water in and out of the pipes.[1] There is no real quantitative mapping here, but rather an intervention of the site that actualizes a model of ocean waves as musical sound producers.This engagement with NSSM creative spaces contrasts with audification, where the

Mapping Space is both essential but inflexible, a direct, completely indexical mapping from data to digital sound. Within audification, Compositional Space is also minimally used, involving only curation and possible pre- or post-processing. For example, the audification of seismological data presented in the last chapter involved no compositional activity besides time dilation pre-processing, and then perhaps amplification or compression post-processing.

A more complex example may be found in real-time group composition (collaborative improvisation). Here, as was discussed in the last chapter, the interactions of an ensemble may be seen as a model of an ecosystem. The dynamics of this system might change in response to the musical gestures of a member of the ensemble, a particular acoustical feature of the room in which the improvisation is being done, a change in the state of live electronics used within the performance, or a prescribed set of rules set out before performance. In this context, Compositional Space is constantly shifting and reflexively self-organizing, saturated with feedback and feedforward loops. Live electronics NSSMs such as data-driven DMIs might foreground an expressive, highly-explicit mapping between the real-time input data and the instrument's sound generator, a mapping that may dynamically change during a performance based on input (feedforward) or adapt based on the sonic output (feedback).

Each of these different example NSSM creative spaces configurations, by no means an exhaustive list, are visualized below (Figure 20).

Figure 20

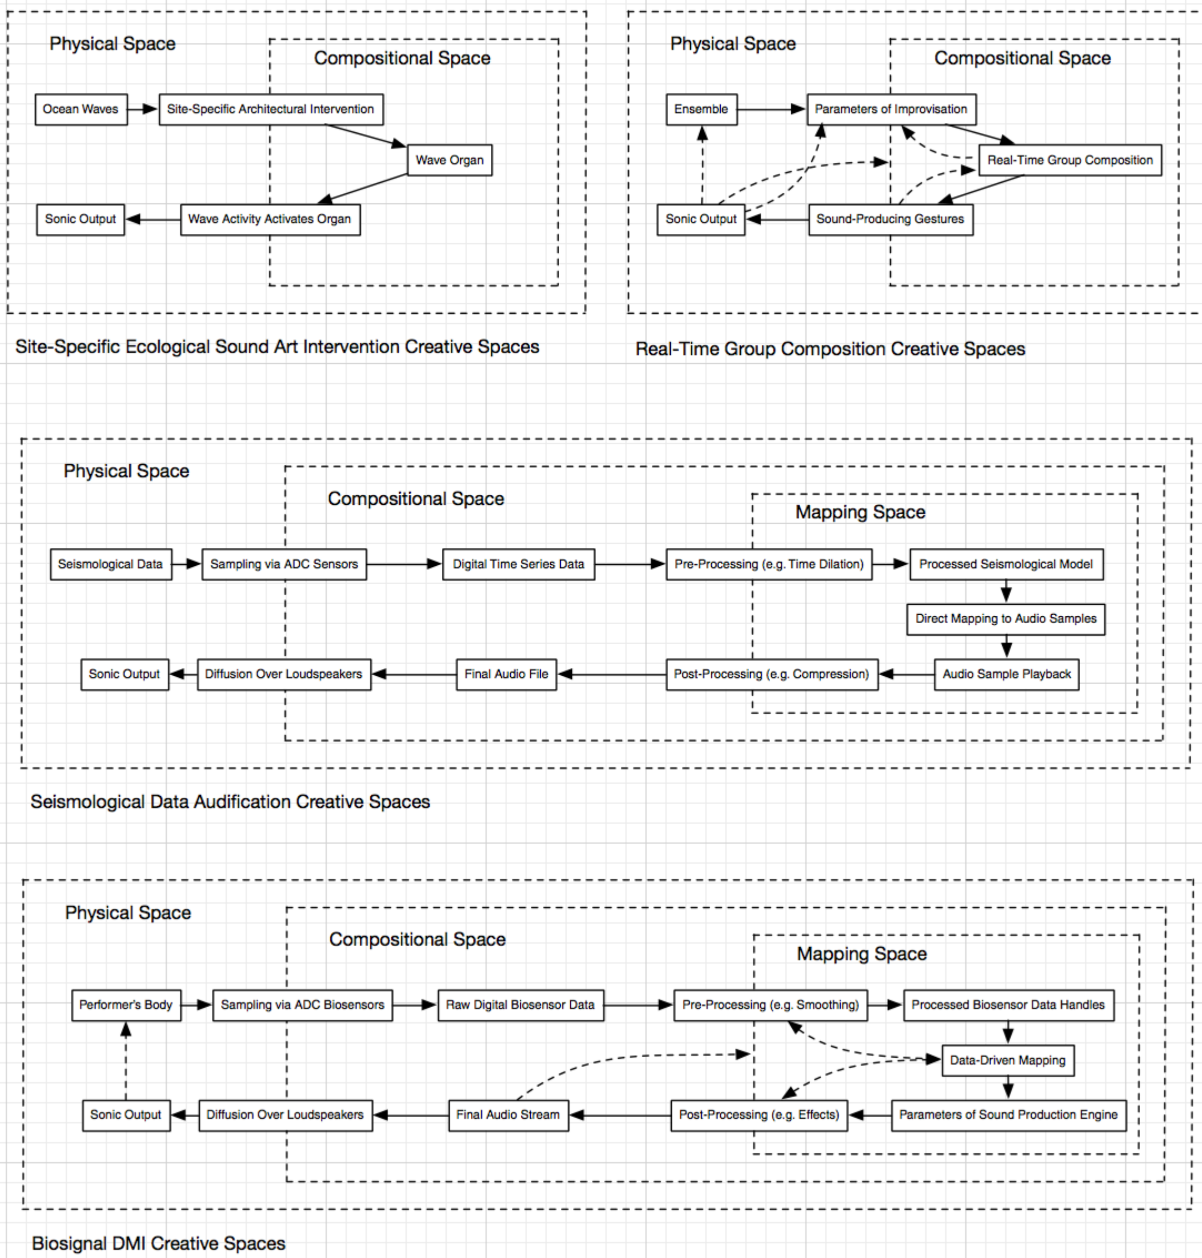

Specific creative spaces for different types of natural system sound models

The different creative spaces presented in this section focus on the internal workings of the NSSMs and, as their name implies, foreground their creative development. Supplementing our analytical method to include listener comprehension and perception and the forces behind NSSMs is a 3-dimensional classification method, discussed next.

2.2 NSSM Classification

All NSSMs have in common several shared components, namely a relationship to a natural system, a model of a natural system, a means for adapting a model of that system to/through sound, and sonic output. Placing these aspects into three dimensions, we may say that a particular NSSM is:

1. *Original*, *Sampled,* or *Simulated*, where the original, physical natural system is harnessed, snapshots of that natural system (samples) are captured and stored, or abstracted principles of the system are used to build a simulation, respectively. This dimension will be called the *modeling methodology* of the NSSM.
2. *Acoustic*, *Live Electronics,* or *Acousmatic*, where only acoustic sound sources are used, where electronics and/or acoustic sound sources are performed in real-time, and where only fixed, electronically-produced sound sources are present, respectively. This dimension will be called the *sound production technology* of the NSSM.
3. *Representative*, *Semi-representative,* or *Abstract*, borrowing terms from art studies, where the creator of the NSSM is emphasizing the literal content of the natural system (representative) or foregrounding aesthetics and stylistically presenting it through various layers of abstraction or curation (abstract). The space in between these two extremes, semi-representative, either uses curation techniques on indexical source material or presents a non-realistic systemic caricature of the natural system. This dimension will be called the *level of abstraction* or *indexicality* of the NSSM, where indexicality is inversely related to level of abstraction.

Taking these dimensions together, NSSMs may be classified based on their modeling methodology, sound production technology, and level of abstraction.

Figure 21

Modeling Methodology	Sound Production Technology	Level of Abstraction
Original	Acoustic	Representative
Sampled	Live Electronics	Semi-Representative
Simulated	Acousmatic	Abstract

NSSM Classification Dimensions

To better classify NSSMs using these dimensions, two important dependencies between dimensions must be mentioned. First, all NSSMs that do not engage a sampling or simulating methodology, acting only upon the natural system itself (‘Original’), are by default acoustic and representative. Second, as was discussed in the exposition, sampling and simulation modeling methodologies are divergent, but not unconnected processes: simulations may produce streams of samples and sampled data may be used to train a simulation. Regardless, NSSMs that use sampled models (data sets, acoustic recordings) will be analyzed separately from NSSMs that use simulations. Taking these dependencies into account, the figure below is a visualization of the different categorizations within the NSSM classification methodology. Below that is a table that describes each category and associates a pertinent NSSM with it, many of which were discussed in the previous chapter.

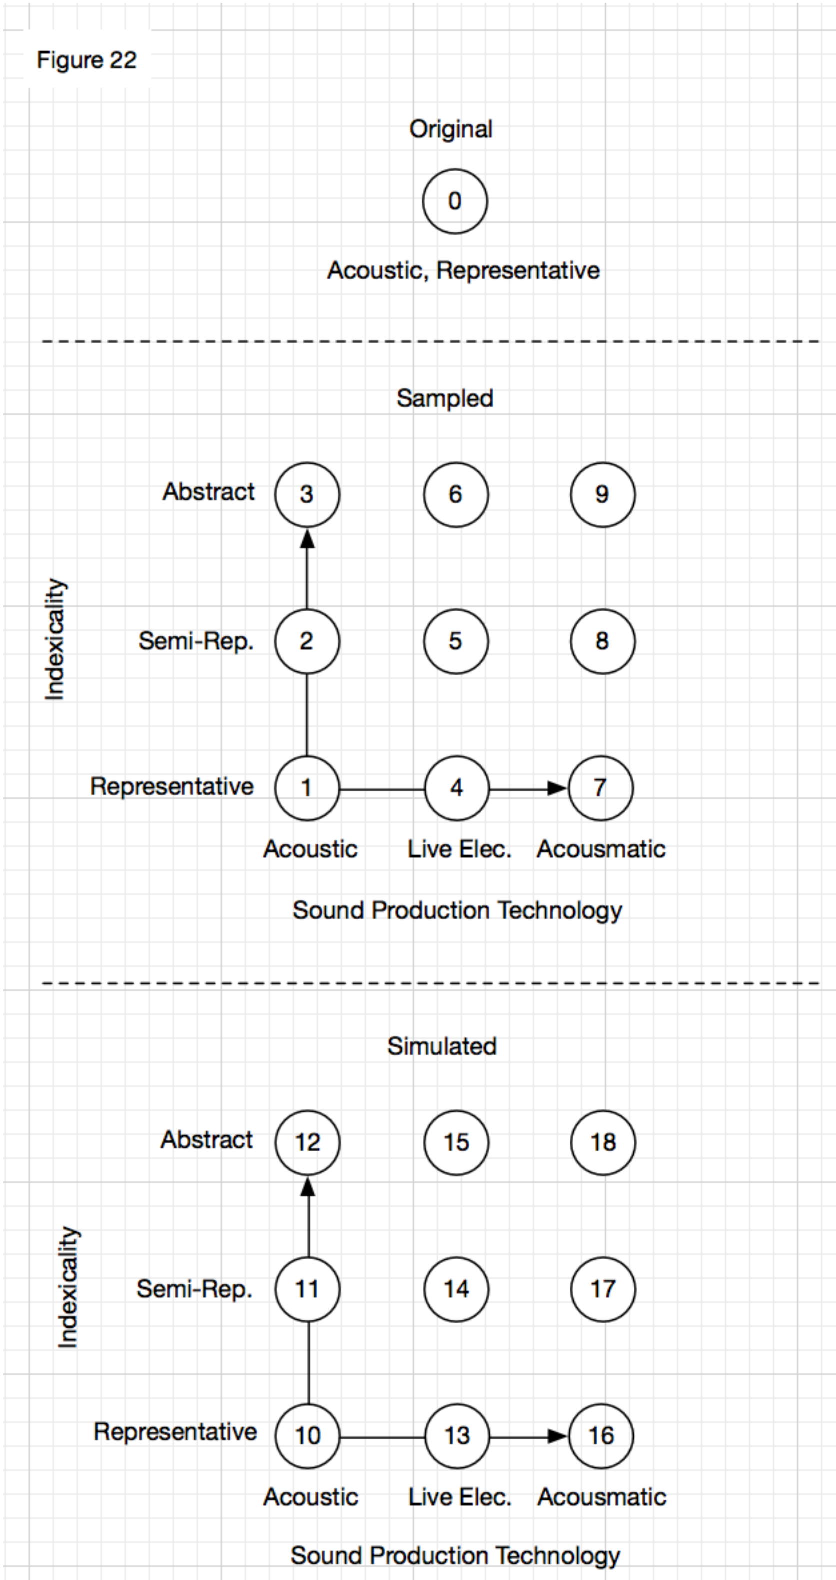

3-Dimensional NSSM Classification Framework

2.3 NSSM Classification Examples

Altogether, this framework contains NSSMs of 19 Types organized within 7 Supertypes. Each Type will now be described and work (many from the previous chapter) of each Type will be given. Figure 23 graphically organizes this information, along with non-exhaustive descriptions of work that might be classified as each Type. As with any taxonomy that realistically attempts to encode many types of creative expression, there are examples that defy classification or may be positioned in multiple Types, some of which are discussed after the outlining of the Types.

NSSM Type 0 contains NSSMs that act only upon the natural system itself ('Original'), and because of this are necessarily acoustic and representative. Instead, the model of the natural system is, perhaps, a coloring, a filtering, a re-organizing. NSSMs of Type 0 may be described as *Interventionist Site-Specific Ecological Sound Art*, typified by Richard and Gonzalez's *Wave Organ* or the aeolian harp.

NSSMs Type 1 through 3 use sampling as their modeling methodology and are entirely acoustic. The natural system sampling in this case is transcription: the mapping from a sound source, famously birds but also other animals and environments, to descriptive or prescriptive notation. Collectively these Types may be described as *Performed Natural System Transcriptions*. These works range from the systematic presentation of representative transcriptions in Messiaen's *Catalogue d'oiseaux* (1958) (Type 1), which presents Messiaen's own transcriptions of bird song within the context of a solo piano work, to more expressive, liberally re-organized sound recording transcriptions in works such as Mâche's *Le printemps du serpent* (2001) (Type 2), a work for large percussion ensemble that incorporates transcriptions of recordings of birdsong, insect sounds, and other natural phenomena such as raindrops, to highly abstracted, computer-mediated transcriptions: Peter Ablinger's *Quadruturen IV* (1998) (Type 3), for example, where field recordings are transcribed by a computer using a decidedly non-representative algorithm and notated for orchestra. [2] [3]

NSSMs Type 4 through 6 use sampling as their modeling methodology and make use of live electronics. Here, the natural system sampling is done through real-time analog-to-digital sensors: often microphones, but also other sensors (motion sensors, electromagnetic sensors, pressure sensors, etc.) which allow real-time, real-world data to be mapped to sound. Collectively these Types may be described as *Amplifications of the (Unheard) Activity of Natural Systems*. These works range from performances that use the microphone, lightly mediated, as sonic microscope, such as in Stockhausen's *Mikrophonie I* (Type 3), to works that abstract live recordings, such as Lucier's *I am sitting in a room* (Type 5), on to works that position the human body or another living system as a source of data within a highly abstracted real-time sound mapping, such as in Atau Tanaka's *Myogram* (Type 6).

NSSMs Type 7 through 9 use sampling as their modeling methodology and are acousmatic. Here, sampled data is mapped to acousmatic sound through mapping processes borrowed from sonification, and collectively these Types may be described as *Natural System Data-Driven Music*. These works range from the highly indexical, barely mediated audification, such as in Hayward's *Listening to the Earth Sing* (Type 7), which audifies seismological data for musical purposes, to semi-representative, mediated and expressively mapped acousmatic sonification, such as in Adam's *The Place Where You Go To Listen* (Type 8), to more abstracted data-driven music, such as Scaletti's *h->gg* (Type 9).

NSSMs Type 10 through 12 use simulation as their modeling methodology and are acoustic. Here, the interactions between performers within a collaborative improvisation context simulate a natural system, and collectively these Types may be described as *Collaborative Improvisation as Natural System Simulation*. These works range from improvisations that have low mediation and are representative of the full range of intra-systemic relationships, such as Cardew's *Scratch Orchestras* (Type 10), to works that mediate interaction through descriptive or prescriptive rules, such as Zorn's *Cobra* (1984) (Type 11), to works that incorporate the modeling of improvisational relationships into their design: abstract meta-modeling that invites the performers to co-create the bounds of their own collaborative improvisation, as in Wolff's *For 1, 2 or 3 People* (1964) (Type 12), an open, graphic score which invites the performers to define their own relationships to the improvisation.[4]

NSSMs Type 13 through 15 use simulation as their modeling methodology and make use of live electronics. Here, sonic ecosystems are created with one or more objects and performers, their mediated interactions constitute the work, and the microphone, and other electronics, are used to amplify the interactions within the ecosystem, and collectively these Types may be described as *Electroacoustic Ecosystems*.These range from works that position the performer as curator, explorer, or investigator of electroacoustic objects, as in Nakajima's Object Performances (Type 13), on to works where the sounds of that simulated electroacoustic world are further abstracted through processing, as in Joo Won Park's *Toccata* (2017) (Type 14), which uses SuperCollider to process the sounds of contact mic'd found objects[5], on to complex, computer-enhanced improvisational ecosystems, such as Lewis' *Voyager* (Type 15). Additionally, some projects in this Supertype might not privilege a direct-to-sound mapping of a simulation. Instead, simulations of learning (via ANNs or GA, for example) might be used to implement an 'intelligent', adaptive mapping strategy, as in Fiebrink and co-author's *Wekinator*.

NSSMs Type 16 through 18 use simulation as their modeling methodology and are acousmatic. Here, micro-level computer sound simulation and macro-level sound organization are models of natural systems in acousmatic sound, and collectively these Types may be described *Digital Acoustical Modeling + Simulation-Driven Acousmatic Sound*. These works range from works that incorporate representative physical modeling synthesis such as Risset's *Sud* (1985) (Type 16), which makes extensive use of standard physical modeling synthesis as musical material, to works that extend physical modeling into a more fantastical domain, such as Cadoz's *Gaea* (2007), a work created with his *GENESIS* system (Type 17), on to works that use computer simulation of natural systems at a much larger temporal scale, informing the structuring of electroacoustic materials, as in Barrett's 'In the Rain' from *Three Fictions: Northern Mix* (Type 18).

Figure 23

#	Description	Example
ORIGINAL		
0	Interventionist Ecological Sound Art	Richard and Gonzalez's Wave Organ
SAMPLED		
Acoustic		
1	Performed Transcriptions of Natural Sounds	Messiaen's Catalogue d'oiseaux
2	Collaged Transcriptions of Natural Sounds	Mâche's Le printemps du serpent
3	Highly-Mediated Natural Sound Transcriptions	Ablinger's Quadraturen IV
Live Electronics		
4	Live Feedback	Stockhausen's Mikrophonie I
5	Sound Out of Time (Looping)	Lucier's I am sitting in a room
6	Natural Data-Driven Performance	Tanaka's Myogram
Acousmatic		
7	Natural Data Audification	Hayward's Listening to the Earth Sing
8	Natural Data Sonification	Adam's The Place Where You Go To Listen
9	Natural Data-Driven Acousmatic Music	Scaletti's h->gg
SIMULATED		
Acoustic		
10	Open Improvisation	Cardew's Scratch Orchestra
11	Bounded Improvisation	Zorn's Cobra
12	Modeling Improvisational Relationships	Wolff's For 1, 2 or 3 People
Live Electronics		
13	Amplification of Sonic Ecosystems	Nakajima's Object Performances
14	Live Processing of Sonic Ecosystems	Joo Won Park's Toccata
15	Complex Electroacoustic Ecosystems	Lewis' Voyager
Acousmatic		
16	Physical Modeling	Risset's Sud
17	Abstracted Physical Modeling	Claude Cadoz's Gaea
18	Organic Systems in Electroacoustic Music	Barrett's In the Rain

Examples of NSSMs Positioned Within the Framework

There are a number of NSSMs that thwart single-Type description. Examples include Di Scipio's *AESI*, which uses sampling (a microphone feed of the performer's instrument) within the context of simulation (the system as a whole functioning as an evolving and adaptive ecosystem), thus existing somewhere in between Types 6 and 15. Alvin Lucier's *I am sitting in a room*, a work that directly engages the properties of acoustical systems through a recursive technological performance, might be better positioned as

Type 16 or even 13. Live electronics performances using natural system data that happen within a studio (for example, Barrett's uses of the *Cheddar* sonification system for multi-channel acousmatic works) might be best described as either of Type 6 or 18.

Combining some concepts from the Exposition (specifically the concept of keys), along with the concept of NSSM creative spaces and the NSSM classification methodology outlined in the previous two sections, we now have the tools to effectively compare and contrast NSSMs. For example, Thomas Beverley's *Telepresent Storm Rita* is a NSSM of Type 6 (abstract, sampled, live electronics) that uses hurricane Rita as natural system, data sets from meteorological sensors measuring hurricane Rita as natural system models, and maps those data to sound through visualizations that are read as a score by the performer. The key to this mapping is visible to the audience through visualizations projected for the duration of the work. *Telepresent Storm Rita* shares level of abstraction and modeling methodology with John Cage's *Atlas Eclipticalis*, a NSSM of Type 3 (abstract, sampled, acoustic music) that takes the stars as natural system, their representation within star maps as natural system model, and the translation of those star maps to five-line staff notation as mapping strategy. The key to this mapping is available within program notes of the work (and within its title).

Park's *Toccata* is a NSSM of Type 14 (semi-representative, simulated, live electronics), that takes a set of found objects as natural system, their interactions within a simulated electroacoustic ecosystem as model, and instead of a direct mapping engages computer processing to abstract that model. The key of the mapping used in this work is somewhat opaque, as the sound processing abstracts and computerizes the natural sound of the contact microphone-amplified objects. *Toccata* shares level of abstraction and modeling methodology with Cadoz's *Gaea*, a NSSM of Type 17 (semi-representative, simulated, acousmatic) which takes physical and acoustic systems as natural system, computational simulations (physical modeling synthesis) as model, and maps the output of fantastical versions of physical models directly to acousmatic sound. The key of the mapping in *Gaea* is presented through a visualization of the physical models, although this can become unclear if many physical models are interacting at once. And so on.

Through this evaluation framework works that engage natural systems and models of natural systems that are in seemingly disjunct aesthetic or disciplinary territories may be compared, connecting the threads of natural system models among a wide variety of musical applications.

3. Project Descriptions

The penultimate chapter of this work outlines three NSSMs created by the author, software systems built in Cycling ’74’s Max programming language that engage models of natural systems in the context of multi-channel electroacoustic music. These projects are NSSM Types 15, 8, and 17, respectively, although their designation might change depending on the context of deployment (performed live or used as sound generation tools for multi-channel fixed media compositions).

The first of these projects,*The Murmurator*, is a multi-channel DMI based around a 3-dimensional bird flocking algorithm. The second project is two-fold, consisting first of a bioacoustics-focused field research project for oyster reef monitoring created in collaboration with an environmental scientist and secondly, the intelligent environmental recording-mapping software *AcousTrans*. The third project is also two-fold, consisting first of a software for sonification of Virginia Barrier Island shorebird habitats again created in collaboration with an environmental scientist and secondly, the habitat-based sound spatialization software *HabiSpat*.

3.1 The Murmurator

Abstract

The *Murmurator* is a natural system-driven multi-channel digital musical instrument. At the core of the Murmurator is a three-dimensional bird flocking simulation that spatializes a corpus of granularized audio files. Heuristics of the flocking simulation affect different sonic characteristics of the audio files: file choice, playback location, speed, and effects. The sonic output is in turn analyzed and used to alter parameters of the flocking simulation: flock cohesion, separation threshold, inertia, and wind resistance. These spatio-sonic and sono-spatial mappings, which together create a trans-modal feedback loop, are managed live by the performer, who pushes this semi-autonomous system into different emergent states and corresponding behaviors. The result is an immersive electroacoustic musical experience that oscillates between dense, ambient textures and fleeting, chaotic gestures.

Introduction

This project began as a quest to seek out a particular electroacoustic aesthetic, one that explored the middle ground between musical gesture and texture in the context of live performance over multi-channel loudspeaker systems. More specifically, the author (and collaborator Kevin Davis) were interested in creating a system that would allow one to perform different colors, timbres, or sound masses over a multi-channel loudspeaker system, to compose in real-time the relationships between multiple textures distributed in space.

The history of performed electroacoustic spatialization can be traced back to 1951 with the development and use of the potentiomètre d'espace by Pierre Schaeffer and Pierre Henry at the Radiodiffusion-Télévision Française.[1] Since then, as more and more systems and spaces dedicated to multi-channel sound diffusion of acousmatic music become available—non-standard multi-loudspeaker systems such as the BEAST or Acousmonium or high density loudspeaker arrays such as the Espace de Projection or The Cube, for example—so do tools (often software tools) designed to control the projection of sound in space. These tools off a variety of methods, algorithmic processes, and interfaces for dynamic control of spatialization.[2] Each also privileges different performance contexts, technologies, and musical aesthetics, ranging in performativity, generalizability, and extra-musical relationships.

Examples include Richard Garrett's Audio Spray Gun[3], Robert Normandeau's Octogris[4], Scott Wilson et al.'s BEASTmulch[5], Jan Schacher's ICST Ambisonics Tools[6], IRCAM's spat~[7], Ico Bukvic's D4[8], and Natasha Barrett's Cheddar[9].

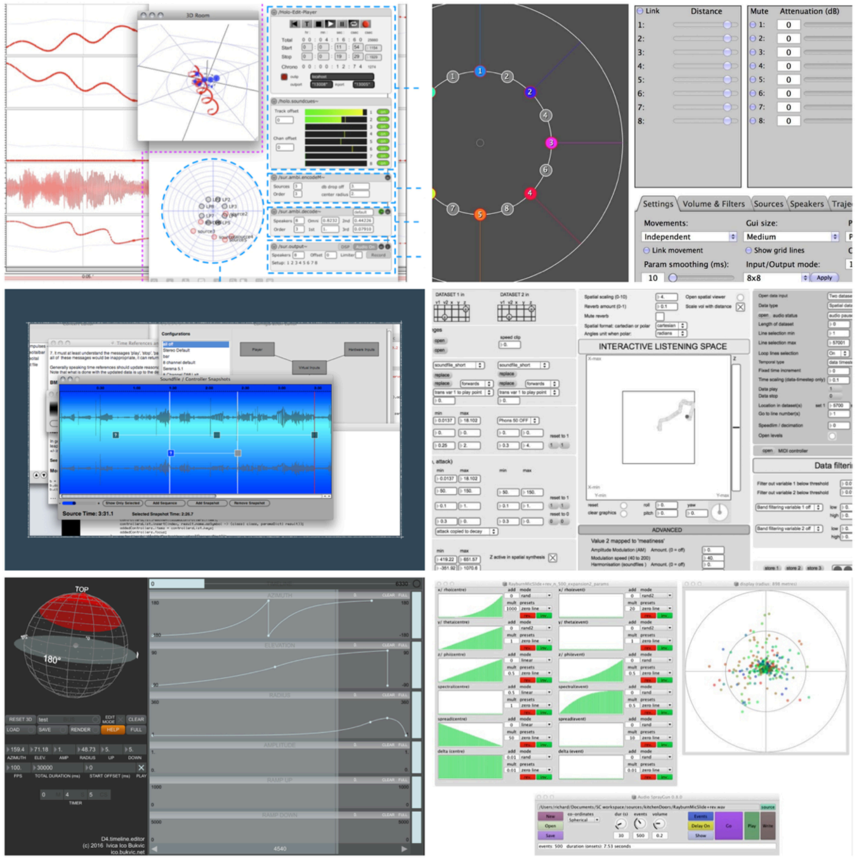

Examples of spatialization software

To reach the authors' sonic goal, it was determined that no off-the-shelf spatialization system was sufficient, so development began on a new system. It was decided that granular synthesis, a flexible synthesis method for coaxing out and sustaining the timbres of portions of audio files, would be used to generate sonic material, and that flocking simulation(s) (specifically Craig Reynold's Boids[10]) would be used to present these timbres not as single points in space, but as tangible electroacoustic masses whose volumes, densities, and positions in space could be controlled. Seeking a spatialization method, the author's experience with ambisonics, a flexible spatial encoding and decoding standard, made it a natural choice for handling the technical details of real-time spatialization in this system.

Taking these different methods and tools in hand, the author and Davis created the *Murmurator*, a flocking simulation-driven multi-channel software instrument, collaboratively over 3 months, with a premiere performance at The Bridge Progressive Arts Initiative on November 9th, 2017 in Charlottesville, Virginia. Since then, the system has been performed in duo or solo form at the 2018 New Interfaces for Musical Expression conference (Blacksburg, VA), the 2018 CubeFest (Blacksburg, VA), at the 2018 International Computer Music Conference (Daegu, South Korea), where a paper outlining its development was also presented[11], at Burning Man 2018 (Black Rock Desert, NV) as a fixed, octophonic piece, at the 2019 (T)echs Machina Festival (Oberlin, OH) and at the 2019 Society for Electroacoustic Music in the United States conference (Boston, MA).

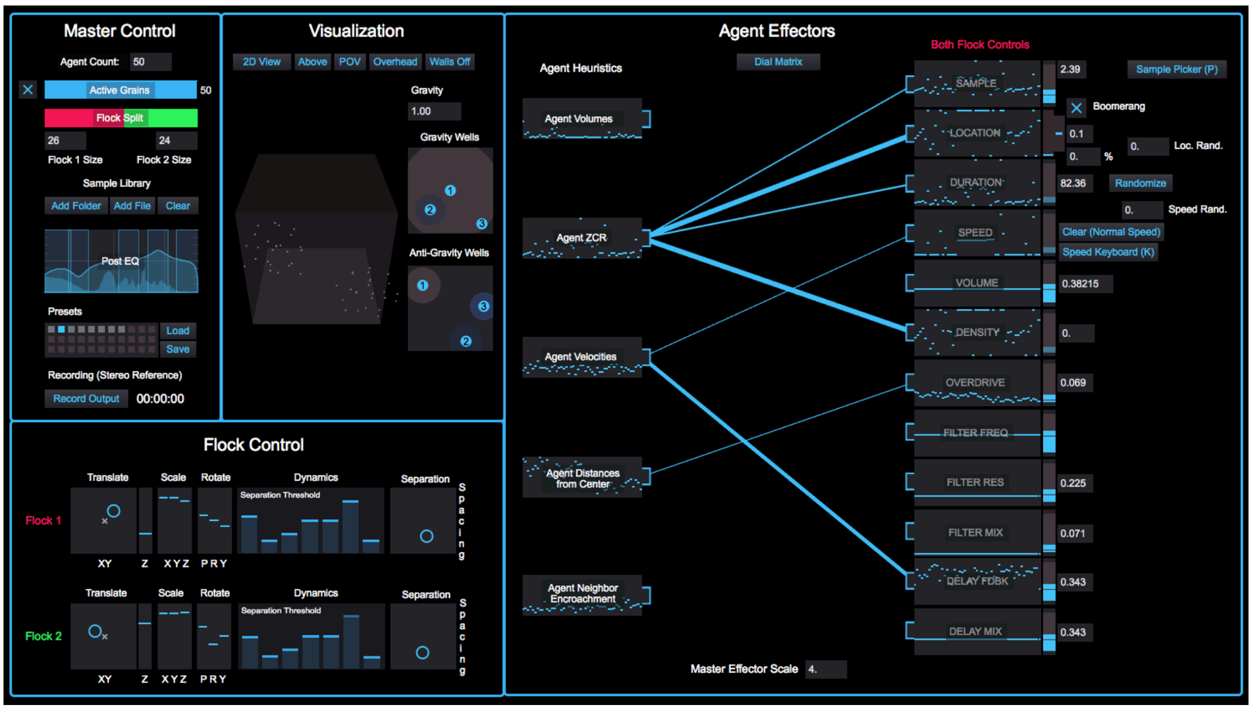

Murmurator v.2 Interface

Design + Implementation

The Murmurator is built around a three- (previously two-) dimensional bird flocking simulation consisting of a set of agents under constraints that result in naturalistic, emergent behavior. Each agent in the simulation corresponds to a stream in a polyphonic granular synthesizer that granulates sound from a corpus of audio files chosen by the performer. The spatialization of each grain stream is controlled in ambisonic space by the corresponding agent's location. The grain streams are further processed using delay, filtering, and distortion, the parameters of which are influenced by the spatial and sonic characteristics of the agents corresponding to each grain stream (a process controlled through meta-interfaces called ‘effectors’).

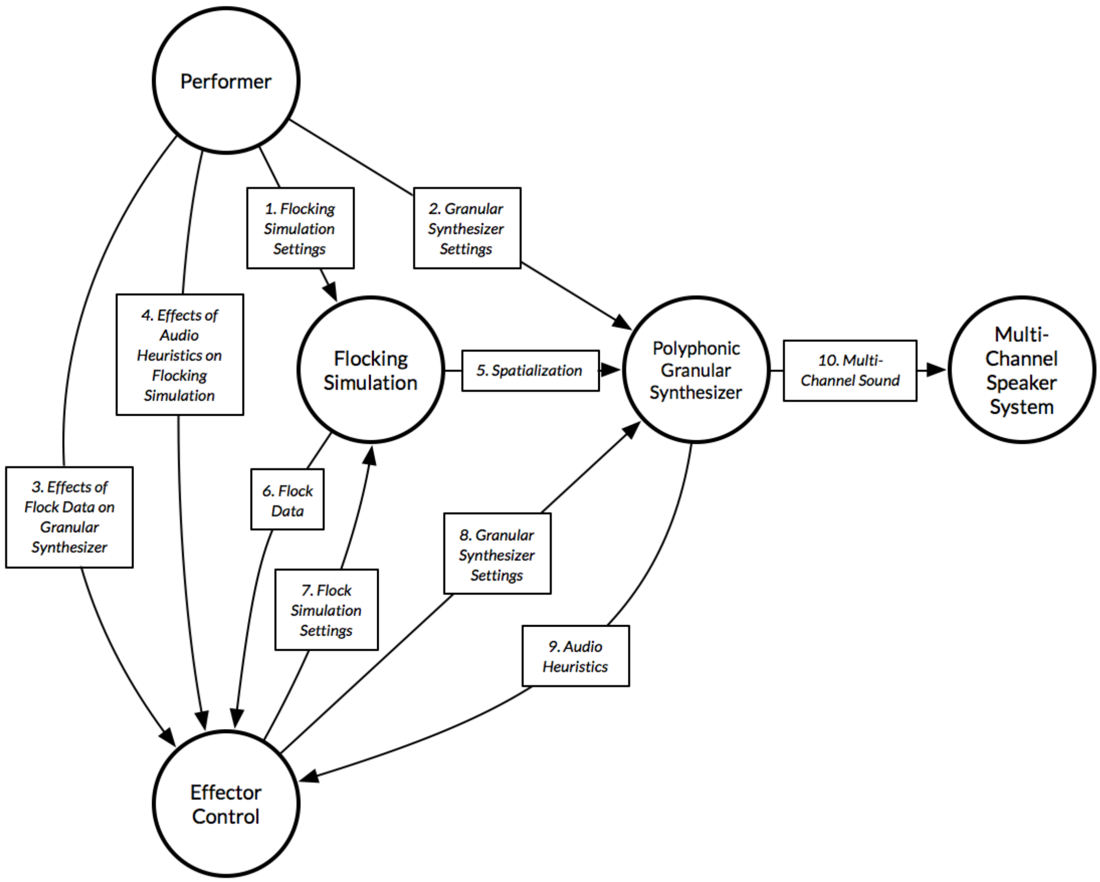

Systemic diagram of the Murmurator v.2 as NSSM

More explicitly, there are a significant number of required or optional connections (some which afford feedback) between the four primary components of the Murmurator: the performer, the flocking simulation, the effector control system (which will be outlined in detail below), and the polyphonic granular synthesizer (see Figure above). Ordered by the component where the connection originates, these connections are as follows:

1. The performer may alter different parameters of the flocking simulation, changing the flock’s behavior
2. The performer may alter the settings of the polyphonic granular synthesizer
3. The performer may alter how the flock data changes the parameters of the granular synthesizer

4. The performer may alter how the audio heuristics change the parameters of the flocking simulation
5. The flocking simulation controls the spatialization of the different streams of the granular synthesizer
6. The flocking simulation's data are sent to the effector control component
7. The effector control may change the flocking simulation settings based on its settings and the audio heuristics
8. The effector control may change the granular synthesizer settings based on its settings and the flock data
9. The granular synthesizer sends audio heuristics to the effector control component
10. The granular synthesizer's output is sent as multi-channel sound to the multi-channel loudspeaker system

In summary, the performer has agency over the quality of the flocking simulation, the sound of the granular synthesizer, and also how much each dynamically changes the other. The data from the flocking simulation is used to spatialize streams of the granular synthesizer and also to control the granular synthesizer. The granular synthesizer generates the sonic material of the system and heuristics extracted from this sonic material can also control the flocking simulation. Lastly, the effector control acts as a meta-interface, a switchboard of sorts that allows the performer to have different components of the system interact with one another. Implicit within this system are also the auditory and visual feedbacks to the performer—the acousmatic sound out of the speakers, the flocking and effector control visualizations— which are constantly informing their performative actions.

Another way to understand this system is by analyzing its graphical user interface, within which the Murmurator's functionality is divided into four primary modules: Master Control, Visualization and Flock Control, Agent Effectors Control, and Granular Synthesizer Control.

Master Control Module

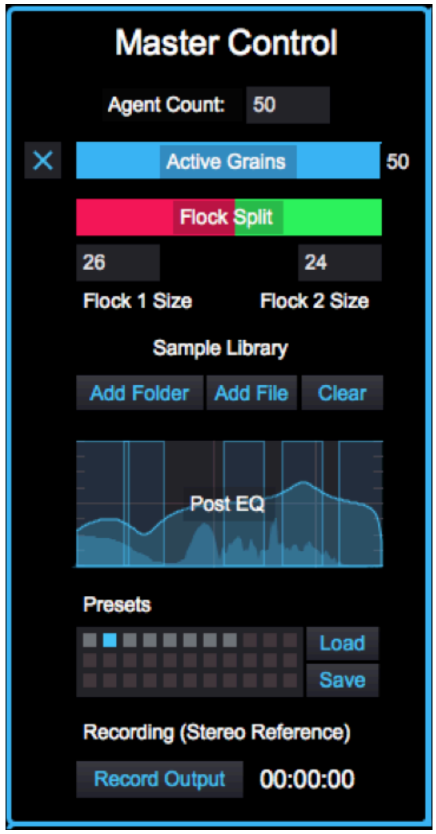

Master Control Module

The Master Control module allows the user to specify how many agents the simulation has (simultaneously controlling the polyphony of the granular synthesizer), how many of those agents are currently producing sound, to assign those agents to one of two flocks that may be controlled independently, and to load in a folder of audio files which act as source sounds for the granular synthesizer. A master equalizer and real-time spectrogram, acting on every channel of the polyphonic synthesizer before it is sent to the speakers, is also included, allowing for global timbral changes as well as hyper-specific equalization tweaking. A preset system allows the performer to save and load all of the parameters of the system. Control-specific presets may also be quickly recalled by hovering over a particular interface and pressing the ‘esc’ key. The ability to record the output of the system to undecoded higher order ambisonics (3rd and higher) is also included, generating recordings that can then be decoded for listening on headphones, over a stereo or octophonic system, etc.

Visualization and Flock Control Module

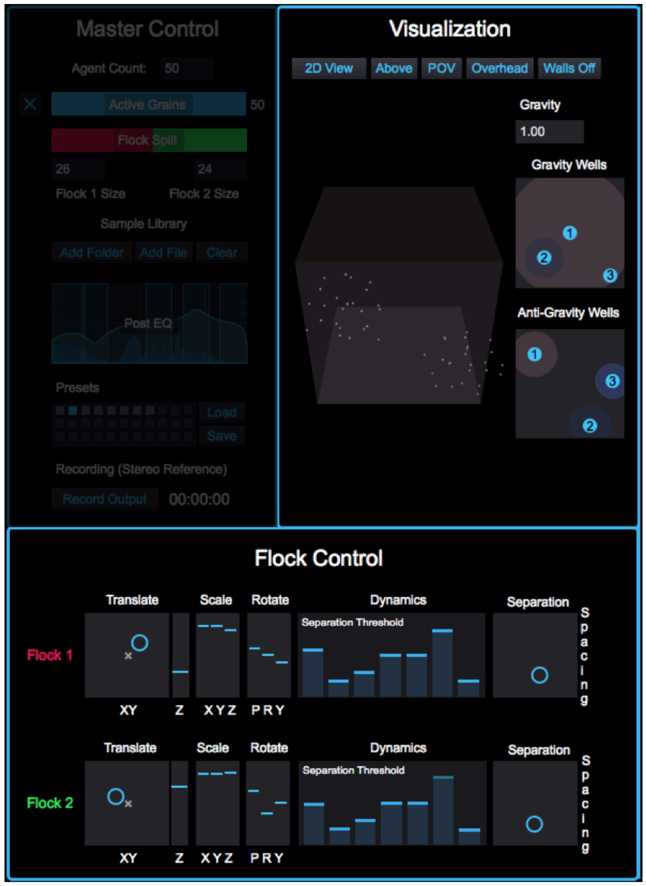

Visualization and Flock Control Module

The Visualization and Flock Control module gives the user high level control over the flock: first, the ability to translate it left-right, forward-back, and up-down (if speakers at multiple heights are available) in multi-channel space, to compress or expand it in three dimensions, defining its volume and shape, and to animate its pitch, yaw, and roll, allowing the sound of the flock to spin and tumble. Low level, precise control over the parameters of the flocking simulation are also included, specifically separation, alignment, inertia, gravity, friction, and coherence. Each of these parameters alters the emergent dynamics of the simulation in a different way, defining its speed, density, and spatial complexity. For example, increasing friction reduces the velocity of the agents, while reducing separation causes them to veer towards one another. Increasing coherence causes the directional movement of the flock to be more unified, whereas increasing inertia causes the agents to slide away from each other, following their own spatial trajectories. Changing these different parameters during performance is one of the primary interfaces to expressive performance within the Murmurator. The location and intensity of gravity and anti-gravity wells, which alter the spatial effects of gravity on the system, may also be tweaked in this module. Split 2-dimensional (XY and XZ displays) and a 3-dimensional visualization of the simulation are provided, with a number of different camera angles and views to assist the performer in evaluating the Murmurator's state.

Agent Effectors Control Module

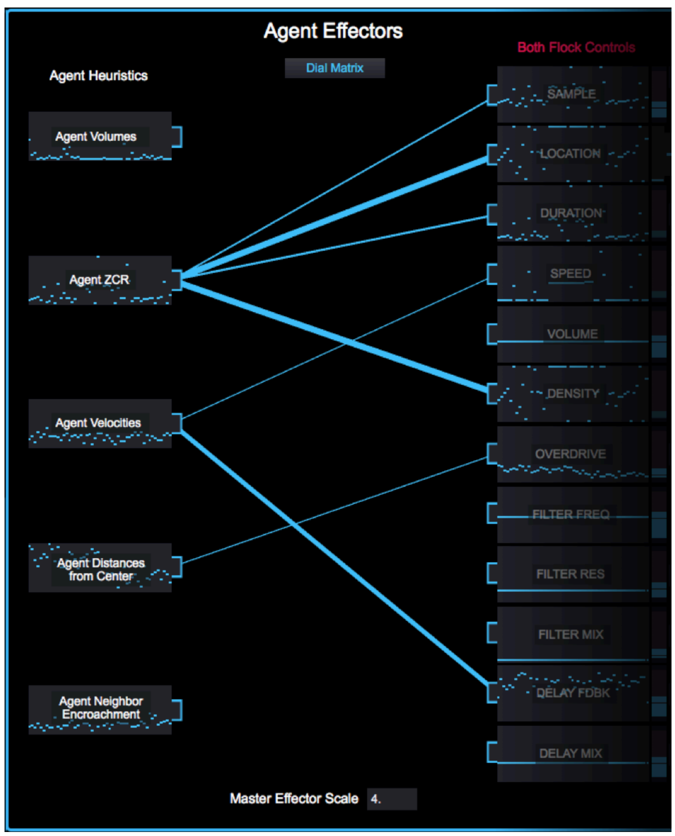

Agent Effectors Control Module

The Agent Effectors Control module gives the user the ability to control effectors, connections that alter the influence of sound and movement-related agent heuristics on granular synthesizer and flocking simulation parameters. These effector connections can be controlled via virtual patch cables, with the patch cable size corresponding to the intensity of affectation or with a dial matrix.

An example of an effector would be the velocity of agents affecting the volume of their respective grain streams: the faster an agent is moving, the louder it is. Alternatively, the overall volume of the granular synthesizer could affect the inertia of the flocking simulation, causing spatialization to be more compact at quieter volumes and more expansive and spread out at higher volumes. These two examples, if used simultaneously, would produce significant positive feedback: the louder the output of the system is, the faster the agents will move, and the faster the agents move, the louder the grain streams will be. While the system does have a master limiter (preventing any catastrophic volume intensity blow-ups) perhaps another, ‘governing’ effector could be used to prevent such a feedback loop and to create dynamic, complex interaction: for example, the position of the agents altering the location within the file that a granular synthesizer stream is playing back from, which would temper the simple feedback of the other two effectors and, given the right settings, create fascinating, complex behavior.

Granular Synthesizer Control Module

Granular Synthesizer Control Module

The Granular Synthesizer Control module gives the user control over all aspects of each channel of the polyphonic granular synthesizer. Within each labelled parameter controller from left to right are the different agents and from top to bottom are the minimum and maximum settings for each of the parameters. At the right of each parameter controller is a master control, which sets all of the grain streams to the same value.

These parameters, for each granular synthesizer stream, are: which sample is being played, the location within that sample that is being granulated, the duration, playback speed, and volume of each grain, the density of the stream (how likely it is at each time step for a grain to be produced), and then the parameters of effects applied to each grain stream. These effects include the intensity of an overdriven tube-like digital circuit, the frequency, resonance, and wet/dry mix of a filter, and the level of feedback and wet/dry mix of a multi-channel chaotic delay line (designed to artificially increase the sense of polyphony in the system).

How files are chosen is worth noting, as it experienced a paradigm shift from version one to version two of the Murmurator. The first version of the Murmurator had control of samples done solely through probability, using a ‘probabilistic sampler’ that the author has used previously in generative DMIs. This system allows the performer to control how often a sound from a corpus is chosen, but not precisely when, shifting control of the audio file selection to be texture- rather than gesture-oriented. A multi-slider interface takes as columns the different audio files, with the height of each slider corresponding to the probability that that audio file will be chosen by all granular synthesis streams. For example, if the corpus contains two files and the height of the first column is very high while the height of the second column is very low, the texture of all of the granular synthesizer streams will be made up almost exclusively of the first audio file, with occasional interjections of the second.

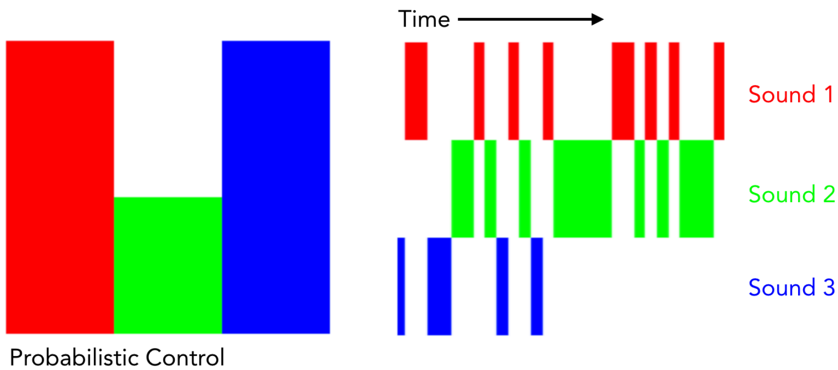

Probabilistic control of samples over time (implemented in Murmurator v. 1)

The second version of the Murmurator drops this functionality in favor of control that allows for each granular synthesizer stream to retain the sample it is granulating. In the current version, the multi-slider interface takes as columns the different granular synthesis streams (voices) and as rows the number of audio files in the corpus. The height of each column, then, corresponds to what sample a particular voice is playing back (if a multi-slider is in between the cracks of the rows a corresponding mixture of two samples is used). The functionality of the first version is also included through a separate interface ('Sample Picker') which populates this multi-slider based on the 'probabilistic sampling' paradigm, allowing for the same fast and intuitive control of the overall texture of the system.

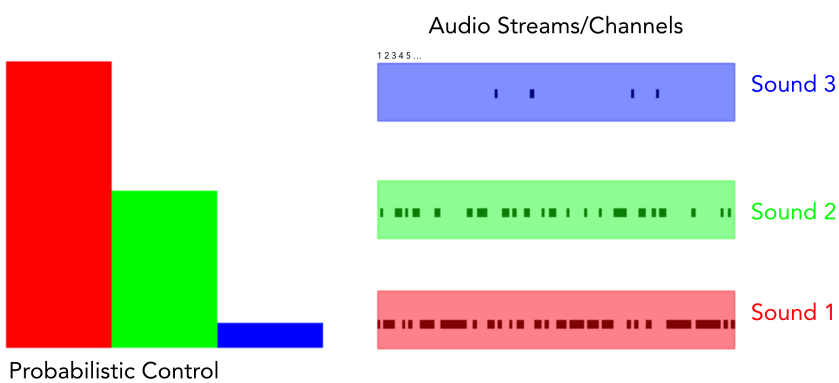

Probabilistic control of samples over all streams (implemented in Murmurator v. 2)

Additionally, the location within the sample can be scrubbed through at a controllable rate (with the scrubbing reversing its direction at the start and end of the file if the 'Boomerang' functionality is turned on) and randomized by a set amount. The playback rate of the samples may be controlled by hand or by using a speed-altering standard piano keyboard interface, and speed may also be randomized by a set amount.

Hardware Interface + Performance Strategies

Each module of the Murmurator presents different affordances of control (some direct, others indirect) and different visualizations that may be used to evaluate the state of the system. The design of these control and evaluation systems was made over many hours of experimentation and performative exploration, determining which controls and visualizations should be privileged and which should be de-privileged. Expanding outwards from just a laptop-based software instrument, a hardware interface and control paradigm was designed to allow for synchronous multi-parameter control of the system using off-the-shelf control surfaces that are intuitively mapped to different parameters.

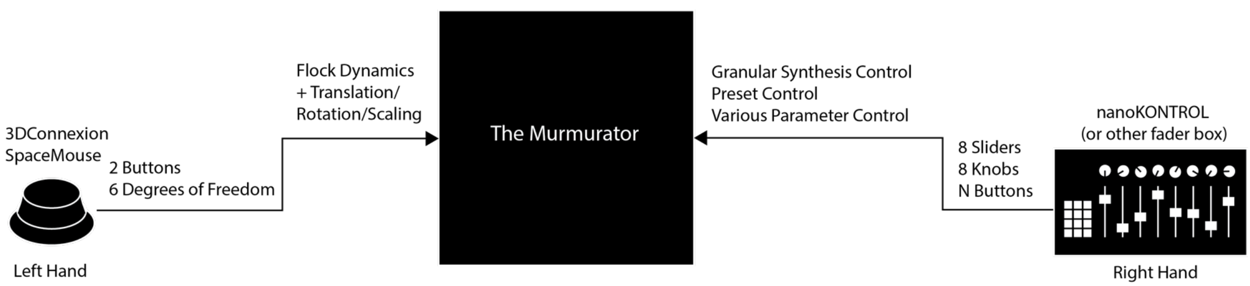

The Murmurator hardware control system

This interface makes use of the 3DConnexion SpaceMouse in the left or non-dominant hand and then a fader box such as the Korg nanoKONTROL or AKAI MIDIMix in the right or non-dominant hand. The SpaceMouse controls flock dynamics, translation, rotation, and scaling with it's unique 6 degrees of freedom 3-dimensional mouse: the left button of this device toggles between 'fine' and 'rough' control modes, allowing for detailed and gross control over parameters, respectively, and the right button on this device switches between modes that afford control over flock dynamics only, translation, rotation, and scaling only, or both.The fader box in the the right hand controls the granular synthesizer parameters and presets, with two different modes: the first where every fader movement is sent to the Murmurator as it happens and the second that sends out the positions of all the sliders when a button is pressed. The effectors control may be controlled via the laptop or alternatively through a knob control surface (not yet implemented, but being experimented with using a Keith McMillan Instrument's QuNeo).

While rehearsing for the premiere performance, the author and Davis co-curated two banks of sounds that each performer would use that complemented the other, in the same way that a duo of instrumentalists decides ahead of time on different materials and correspondingly different musical roles (accompaniment, lead, etc.). A rough form for the piece was then structured via the preset system using these chosen sound banks, with different sections calling for each performer to have more or less control over the system, to give the flocking simulation and its emergent properties more or less of a performative voice. Performance involved immediate or slow transitions to and from each preset (either recalled manually or memorized), with the performative path navigated between each section adapting improvisationally in response to the output of the systems and the actions of the performers.

The author and Davis found that improvisation with the Murmurator often involved oscillating between one of two performance modes: focusing solely on controlling one of the software's modules or listening and fluidly moving between all of the modules, tweaking their individual settings to work towards a desired sonic output. The author and Davis also found that staying out of the spectral space of one another (i.e. if one performer is playing lower frequency sounds avoiding that frequency range) assisted in differentiating the contributions of the two performers and similarly, spatially separating the contributions of each performer was also helpful: e.g. one person placing sounds in the 'middle' of the multi-channel space, the other at its 'sides,' something that is reminiscent of spatial differentiation in laptop or mobile orchestras through local, person-specific speaker systems.

NSSM Analysis

The Murmurator is a NSSM of Type 15 (abstracted, simulated, live electronics). Unpacking the Murmurator as NSSM, engagements with physical space include the natural system of animal flocking behaviors, generally, and the phenomenon of bird murmuration, more specifically. The output of the system may be diffused over high-density loudspeaker arrays (HDLAs) or listened to on headphones.

Engagements with compositional space include the implementation and adaptation of Craig Reynold's Boids flocking simulation, the curation of sounds, the flock control settings, as well as the effectors control, which acts as a meta-interface for the mapping space. Within mapping space, the location of agents in the flock in virtual 3-dimensional space are mapped to the localization of the granular streams in ambisonic space. Additionally, heuristics of the flock movement may be mapped to sonic parameters and conversely acoustic features of the system's output may be mapped to flock movement, affording systemic feedback.

The keys to these mappings might be revealed using projection of the Murmurator's GUI. At the premiere performance, the audience was situated around the performers and was able to view the GUI, with several audience members noting that this enhanced their experience. As of this writing the author feels that the acousmatic potentials of the Murmurator, in particular its ability to situate an audience within a fantastical, constantly-shifting, immersive acousmatic environment, would be lessened by a visual accompaniment.

Future Directions

The trans-modal feedback potentials of the Murmurator are just beginning to be experimented with, producing some amazing results that push this system into a more autonomous role. For example, with the right settings the Murmurator could be used within installation contexts with minimal human involvement, slowly changing its sound over minutes or hours as its chaotic dynamics unfold.

There are also potentials for the Murmurator to be used with live input: instead of granularizing sounds from a corpus of audio files, a live microphone feed or an audio buffer filled with the last ten seconds of live microphone input, for example, could be 'Murmurated'. This extends the system from just a conversation between the Murmurator's operator and the autonomous flocking algorithm to include an improvising instrumentalist or vocalist.

Lastly, rather than limiting the spatialization control to a single system model (Boids), many other simulations and models could be included in the software, each of which would have different input controls, spatial dynamics, and musical potentials.

Conclusion

The Murmurator builds upon previous projects that make use of granular synthesis and natural system models in multi-channel electroacoustic space, but differentiates itself by being designed explicitly to be used in a collaborative improvisation setting. This results in a number of significant, and musically compelling, changes to the system. To reduce the cognitive load required to manage both real-time spatialization of sounds and other musical dimensions and simultaneously to give it an improvisational ‘voice’, distributed, relegated control permeates all levels of the Murmurator’s design, from its probabilistic sampler to the emergent properties of the flocking algorithm. Further, a system to facilitate control of the influence of the spatialization model on other musical dimensions (effectors) establishes a deep connection between the way sounds are spatialized by the system and the processing of the sounds themselves, effectively integrating the parameters of granular synthesis and spatialization in live electronics performance.

3.2a Bioacoustic Monitoring of Oyster Activity on an Intertidal Oyster Reef

Abstract

Marine ecosystems are very loud. In particular, the soundscape of the intertidal oyster reef is a dense polyphonic layering of the ever-present sounds of the tide, the swimming of fish and their throaty calls, shrimps snapping to stun their prey, and, beneath all of this, the quiet, albeit unique sounds of oysters opening and closing. This project includes collaborative marine bioacoustics research conducted during the summer of 2018 at the Anheuser-Busch Coastal Research Center in Cape Charles, Virginia, to develop a methodology for being able to detect those oyster sounds, quiet indicators of an exceedingly important process: the filtering of ocean water, an essential part of the health of coastal environments. Working on this project caused the author to seek the answers to a parallel, highly-related but electroacoustic music-focused question: given recordings of the oyster reef, how can we model the behavior of its many layers, and how can this natural system sound model then be used to create immersive electroacoustic music, music that is gesturally and structurally rich with the behaviors of that ecosystem? What follows is an account of the work undertaken to seek out this answer, including details of its motivation, the iterative design process, the implementation of sound translation and mapping software, and examples of musical deployment.

Introduction + Motivation

The focus of this research project revolves around the Eastern Oyster (*crassotrea virginica*) and its habitat: complex, three-dimensional reef communities that may be found along the Virginia Barrier Islands. The Eastern Oyster is an essential part of the Eastern coast of the United States, filtering the water column and vastly improving water quality.[1] Eastern Oyster populations on the Virginia coast and in the Chesapeake Bay have declined to approximately 1% of pre-1900 levels, and they are currently a major focus of restoration efforts by The Nature Conservancy (TNC) at the Virginia Coast Reserve (VCR).[2]

Bioacoustic monitoring, the recording and analysis of the acoustic emissions of animals, has been used in the study of behaviors and to aid in census counts of many species of animals, including birds, wolves, and marine animals. The marine animals observed include marine mammals, fish, crustaceans, and marine habitats more generally. The excellent propagation properties of sound waves in the water combined with the fact that many marine organisms produce sounds, intentionally for communication or remote sensing, or unintentionally each time they move, makes acoustic monitoring a powerful method for recording animal behavior.[3] A number of different research groups have used acoustic monitoring on the Eastern coast of the US, but no dedicated sound-focused research has been done in the VCR.[4]

Scientific Methodology

Over the course of four weeks of field work during the summer of 2018 the author collaborated with Martin Volaric, a UVA Environmental Science Ph.D. Candidate in the lab of Peter Berg and Matthew Reidenbach, at the Anheuser-Busch Coastal Research Center in Cape Charles, Virginia. The author's roles included sound recordist, sound analyst, and general assistant, and Volaric's roles included marine biologist, non-sound data analyzer, and boat driver. Each work day the author and Volaric would depart the center in accordance with the tide and go out to one of two oyster reefs. Volaric would set out his scientific equipment (recording the depth of the tide, information about oxygen the oyster reef is producing, and other measurements) and the author would set up two hydrophones (underwater microphones) and hook them up to a sound recorder in a small boat. After collecting data over one or two days, the equipment was retrieved and the data looked over (or listened to). At the end of this process 15 tide cycles were recorded, over 180 hours of oyster reef data (including high quality, 96kHz stereo hydrophone audio).

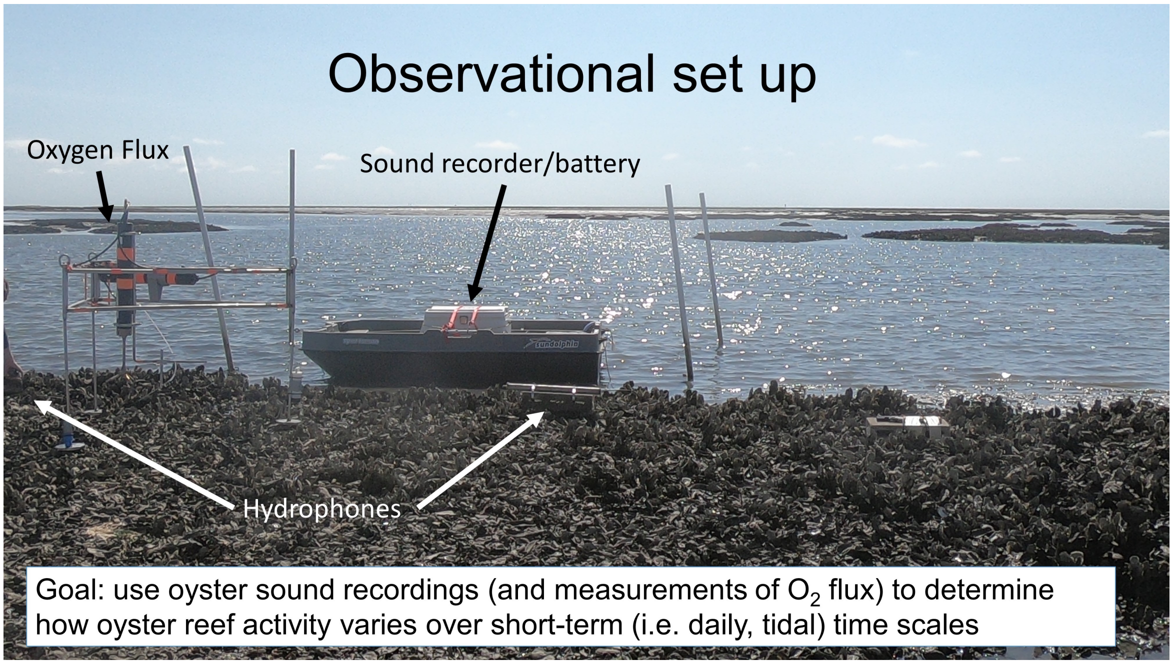

Observational set up for recording oyster reef data

After collecting this data, the author and Volaric worked collaboratively to deploy methodologies for extracting information from the sound data and relating them to the non-sound data. First, the author applied heuristics used in the context of Music Information Retrieval (MIR) and bioacoustics, the results of which included time series data that ranged from the general (zero crossing rate, energy, spectral centroid) to the highly-specific (biodiversity indices developed for rain forest acoustic analysis).[5] While these heuristics were able to be synchronized with the non-sound data and significant correlations were able to be shown, they were not able to indicate information specifically related to the Eastern oyster.

This was primarily because of a single animal: the most dominant feature in all sound recordings made of the reef is of snapping shrimp (*alpheus heterochaelis*), often called pistol shrimp, who use the fast and powerful closing of their large claws (which causes a cavitation bubble) to stun prey. This activity embodies itself in recordings as a dense, cacophonous chorus of wide-band, noisy explosions, exactly the type of chaotic texture that masks other sounds in the soundscape, what could be referred to as 'noise' in this context. A different approach needed to be made, then, one that either filtered out the sound of the snapping shrimp or sidestepped it entirely by honing in on the sounds of the oysters.

Filtering different frequencies of a soundscape to reveal its different components is an incredibly powerful process, one that has been central to acoustic ecology since its birth.[6] As an example, the author made a recording of a dock on the shore which could be bifurcated quite successfully into anthropogenic and non-anthropogenic sounds simply by filtering it into two recordings: a low-pass filtered version containing the anthropogenic sounds of boat motors and human speech, and a high-pass filtered version containing all other sounds within the soundscape. The same is true of the intertidal oyster reef, and research done in this area to define the frequency bands of different biotic and abiotic behavior of marine soundscapes is highly applicable to this research, which will be returned to later.

In order to be able to hone in on the sound of the Eastern oyster, a better understanding of what this sound was had to be made. While bivalve mollusks (scallops, for example) with larger adductor muscles (the muscle that controls the hinge between the two shells) make robust sounds as they close their shells, Eastern oyster muscles are weak in comparison, and in fact no bioacoustics research has been done that includes recordings of them. To isolate the sounds of these animals, the author and Volaric collected clumps of Easter oysters and brought them back to the lab, placing them within tanks of unfiltered ocean water (they were later returned to the location where they were sourced). A video recorder was trained on the oysters and a sound recorder was connected to a hydrophone submerged in the tank. The video and sound recorders were turned on and the oysters were left to filter the water. The sounds of the oysters closing their shells (oyster ‘coughing,’ a byproduct of their filtering process) were then identified by viewing the video and the waveform of the recorded audio. These oyster coughs or clicks were then extracted, denoised, and used within the context of another bioacoustic sound analysis process: acoustic event detection, and more specifically, acoustic template matching.

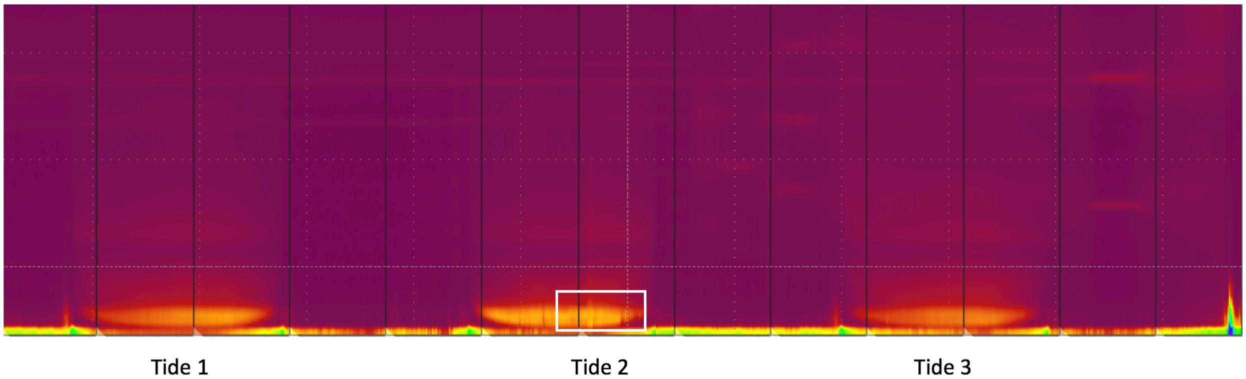

Spectrogram of three tides. Note the “canoe” shape of the spectrogram as the water rises, altering the bandwidth of the snapping shrimp

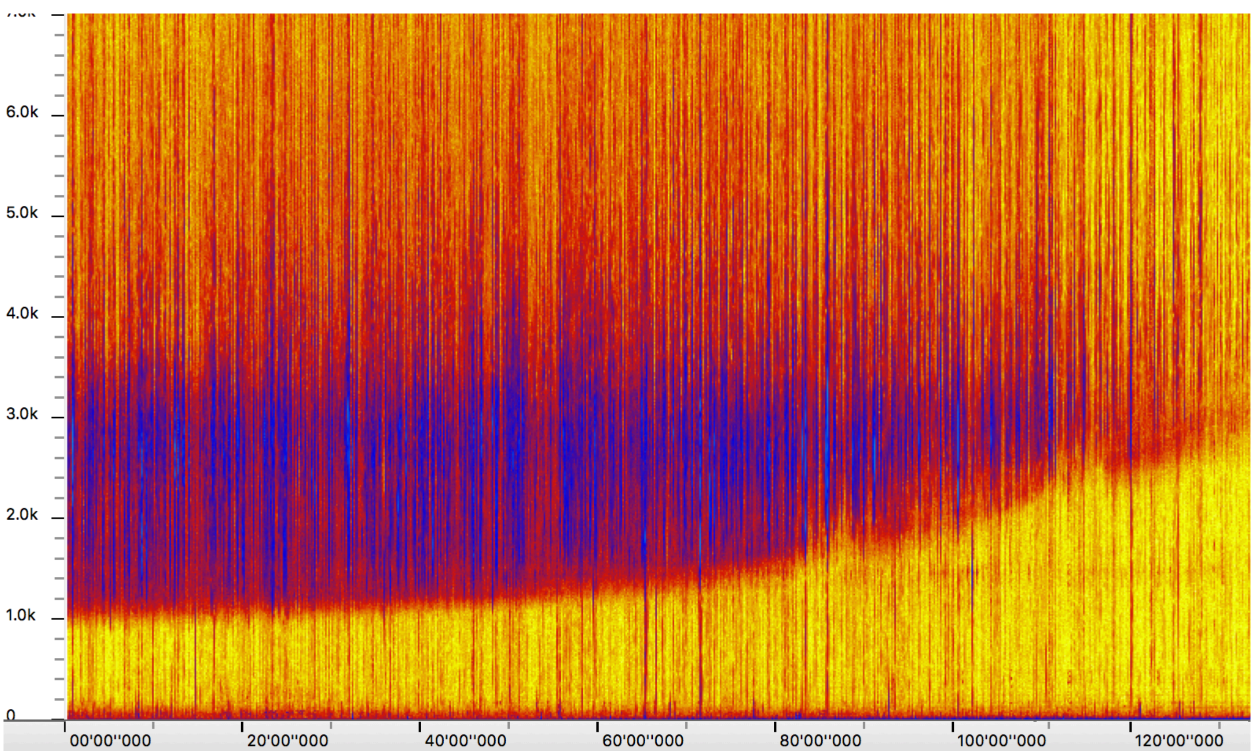

Inset detail of the previous Figure. Note the vertical striations of snapping shrimp activity (center), along with the rising of the lower limit of those striations as the tide recedes

Acoustic template matching (as implemented in the monitoR package, which the author used[7]) involves taking templates (in this case the lab-isolated oyster coughs) and then searching for sounds that match them within a longer recording (in this case the recordings of the reef). The result is time series data of match likelihood: the probability at a given moment that the template sound is present in the recording. This probability data can then be thresholded to generate streams of events: vectors of times when it is significantly likely that the template sound is present.

In the context of the oyster reef recordings this process would be chock full of false positives if the layer of snapping shrimp—noisy, chaotic, and potentially masking the sound of the oysters—was part of this matching process, but fortunately the activity of the snapping shrimps takes place exclusively above 1kHz and the sound of the oyster coughs has most energy below 1kHz. The matching, then, is only done below 1kHz, sidestepping the issue of the dominance of the snapping shrimp within the intertidal reef soundscape. This event data may then be synchronized with the non-sound data.

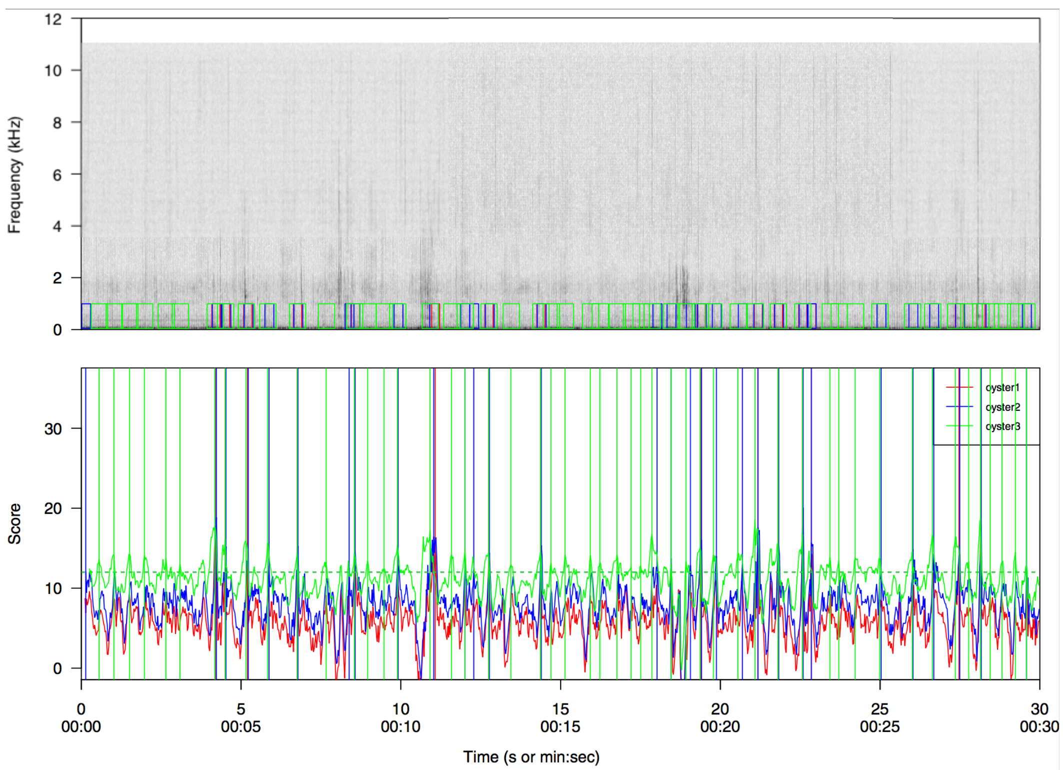

Visualization of the acoustic template matching process. The spectrogram of the reef recording (top) is visually matched with three lab-recorded oyster sounds, generating a time series of likelihood scores (bottom)

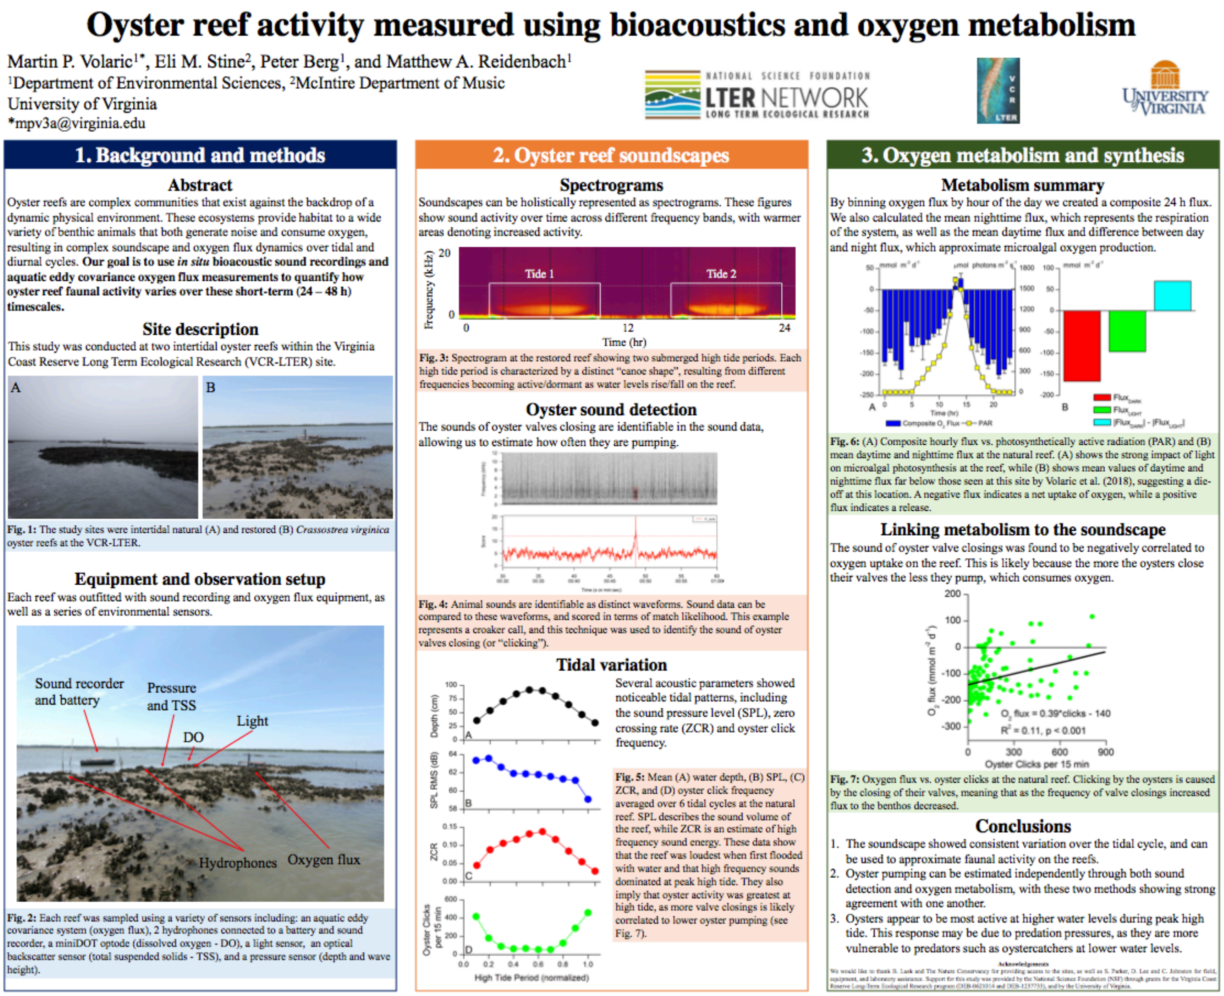

Presentation of scientific research presented at the Long Term Ecological Research (LTER) All Scientist's Meeting

Results

The author and Volaric determined that there was a significant inverse correlation between the oxygen production of the oyster reef and the average cough rate, demonstrating that when the oysters are coughing they are filtering less of the water around the reef and consequently the reef is generating less oxygen. These findings also suggest that this methodology might be deployed on reefs as a surrogate for the oxygen measurement equipment. The results of this work have been presented at University of Virginia as part of the 2018 Coastal Futures Conservatory conference, at the 2018 Long-Term Ecological Research (LTER) All Scientist's Meeting (ASM) in Pacific Grove, California, and at the 2019 Association for the Sciences of Limnology and

Oceanography (ASLO) Aquatic Sciences Meeting in San Juan, Puerto Rico.

While the marine science aspect of this project is essential to the development of its electroacoustic music-focused counterpart, what will be focused on for the rest of this subchapter will be on sound translation and mapping software created in response to, and in concert with, the research project. The following section is adapted from a paper published in the proceedings of the 2nd Workshop on Intelligent Music Interfaces for Listening and Creation (MILC 2019), presented in Los Angeles, California.[8]

3.2b AcousTrans

Abstract

AcousTrans is a software for intelligently mapping a multi-dimensional stream of gestures extracted from one environmental soundscape to an entirely different, multi-channel electroacoustic sound world. After a process of multi-band filtering and event segmentation, events within a stereo environmental recording are intelligently mapped onto multi-channel sound events from another corpus of sounds using a *k*-nearest neighbors search of a *k*-dimensional tree constructed from an analysis of acoustic features of the corpus. The result is an interactive sound generator that injects the organicism of environmental soundscape recordings into the sequencing, processing, and composing of immersive electroacoustic music. This work was created within the context of the bioacoustic analysis of intertidal oyster reefs described in the previous section, but is applicable to any environmental soundscape that may be effectively decomposed using the described method.

Introduction

There exist many different software tools, in a variety of disciplines, which seek to deconstruct and/or map onto sound natural systems.[9] Some approaches engage with natural computing and artificial intelligence: creating musical prescriptions using Cellular Automata, L-Systems, or flocking simulations (such as in the author's own *Murmurator* system).[10] The field of auditory display and sonification, when making use of natural system data, also engages methods for best representing in sound the organic interactions of natural data, with many different sonification softwares available (such as Wilson and Lodkha's Data Sonification Toolkit[11]).

Decoding and transcribing the events within a recording is under the purview of both Music Information Retrieval (MIR) and bioacoustics. Pertinent MIR tasks include automatic transcription, track separation, and speaker diarization, each of which seek to automatically reveal structural decompositions of acousmatic sound. Within acoustic ecology and bioacoustics, techniques have been developed to assist in the decomposition of environmental soundscapes, revealing their underlying sonic components (such as in the work of sound recordist Bernie Krause and composer Hildegard Westerkamp).

The specific method that this work builds from is concatenative sound synthesis, a synthesis technique that may be generally described as granular synthesis driven by audio analysis, and more specifically a process of selecting grains of sound from a file or corpus based on their best fit to some acoustic criteria.[12] There are a wide variety of projects developed over the past few decades that make use of concatenative synthesis, ranging from more or less artistic and scientific applications and from off-line systems to, more recently, real-time implementations.[13] These include cataRT[14], timbreID[15], developments into 'Soundspotting'[16] and 'Audio Mosaicing'[17] techniques, and the work of scientist-musicans Jean-Julien Aucouturier[18] and Aaron Einbond[19], among others. The author's system builds around the MuBu concatenative synthesis engine: connecting it to an environmental soundscape event parser and an expressive electroacoustic sound mapper.

Design + Implementation

The goal of *AcousTrans* (*Acousmatic Translator*) is to allow a user to load in a source stereo audio file (field recording or other environmental recording) and a destination corpus of other audio files and interactively map the events, gestures, and structure of the source onto the destination. What results is a stereo or multi-channel audio file with gestural, rhythmic, and/or structural similarities to the source file, but with entirely different timbral characteristics: those of the destination corpus.

AcousTrans is implemented in Cycling '74's Max 8, taking advantage of ICST's Ambisonics externals to handle multi-channel audio[20] and IRCAM's MuBu for Max[21] and Programming Interface for Processing Objects (PIPO) Max externals to handle acoustic feature analysis.[22]

Event Extraction

AcousTrans operates by first taking in a user-selected stereo audio file of a soundscape (an intertidal oyster reef stereo hydrophone recording, for example) within the segmenter module. This audio file is then played back at a set playback rate. Speeding up of this playback rate contracts time, reducing the time between events, and slowing down of this playback rate expands time, increasing the time between events. Changes in playback rate are done via varispeed (changes in speed also change pitch), which will change the frequencies of the filtering procedure, described next, accordingly.

This stereo audio file is then sent through N pairs of band-pass filters whose frequencies are tuned to the particularities of the soundscape (the specific threshold between the sub-soundscapes of wave movement, oysters, snapping shrimp, etc., within the reef, for example). Within the context of the intertidal oyster reef recordings a value of $N = 4$ was deemed both necessary and sufficient, although the system allows for N to be variable.

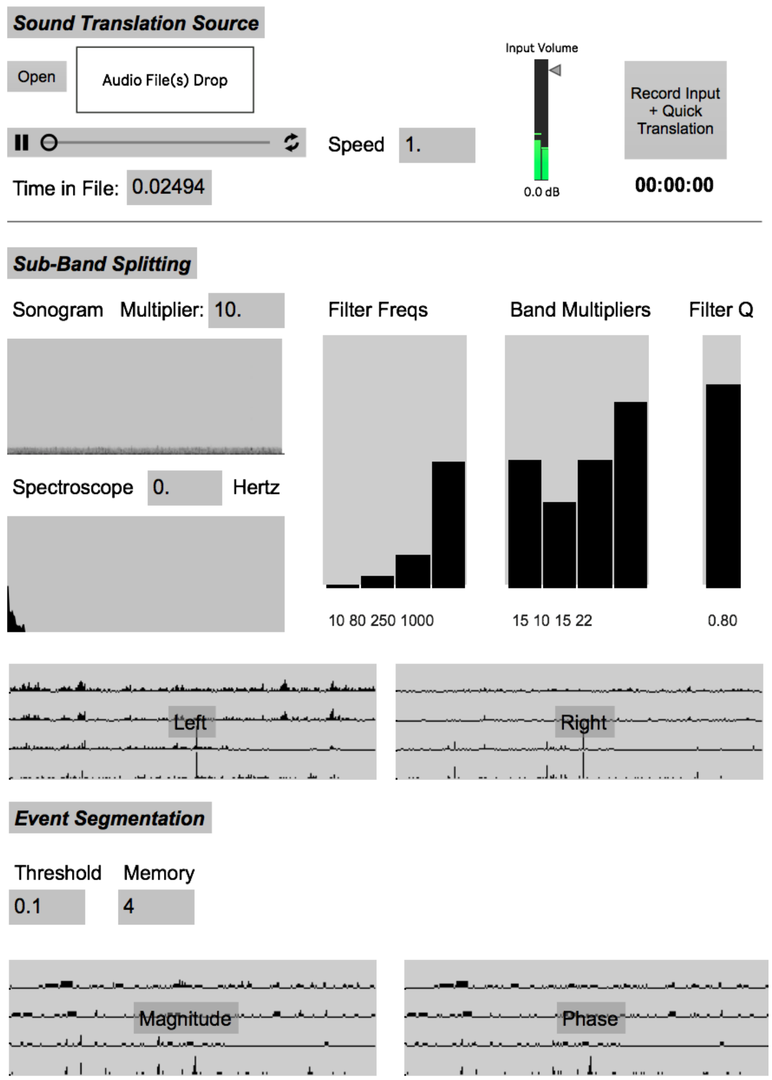

AcousTrans segmenter module, the output of which is fed into the playback module

Each of the 2*N* spectral sub-bands of the source audio file (one for each stereo channel) is then segmented using one of several different segmentation modes: amplitude-based peak detection, spectral flux-derived segmentation, or fixed-size segmentation. The settings of the first two modes may be controlled with two parameters: threshold, which determines the amplitude or spectral flux value at which onset and offset are determined, and memory, which determines the size of an averaging window applied to the most recent amplitude or flux samples received. A lower threshold will result in more events being detected, but increases the potential for a single event to be segmented into multiple events. A higher threshold will result in less events being detected, but has the potential to miss more subtle events. A larger memory helps to smooth out events over time, better representing large activity changes in the bands. A smaller memory results in a finer event granularity but fails to encode large textural changes. Through trial and error the user can tweak these parameters to find settings which effectively represent the activity present in whatever soundscape they have loaded in.

After segmenting each of the $2N$ sub-bands we add them as stereo pairs to get N streams of ‘magnitude’ events, which represents the overall event activity in that stereo sub-band, and calculate their absolute difference as stereo pairs to get N streams of ‘phase’ events, which represent how much each event is shared between the two channels, effectively indicating how spatially localized an event (how much it is panned to the left or right vs. being in the center of the stereo field). Each event is then analyzed in real-time to extract a set of acoustic features.

At the end of this process, each event is encoded as a list including its sub-band, intensity (average volume), duration, stereo localization (position from left to right), and a subvector of acoustic features that describe it, including fundamental frequency (F0) estimation, energy, periodicity, loudness, spectral centroid, spectral spread, spectral skewness, spectral kurtosis. This multi-channel stream of events is then passed into the playback module.

Event Mapping

Two options are available for mapping events extracted from the loaded audio file: Musical Instrument Digital Interface (MIDI) and acousmatic translation. In the MIDI mapping procedure, event intensity is mapped to the pitch of MIDI notes, localization to the corresponding velocity of MIDI notes, and the sub-band tag of each event determines the MIDI channel of those MIDI notes. The acoustic features are discarded (although the author briefly experimented with encoding them as MIDI Polyphonic Expression (MPE) data). The result is a symbolic MIDI encoding of the activity of the soundscape which can be saved, recorded, opened in a DAW, applied to electronic sound or notation, etc. The author started with this event mapping option and then extended it to the much more flexible and expressive acousmatic translation option, described next.

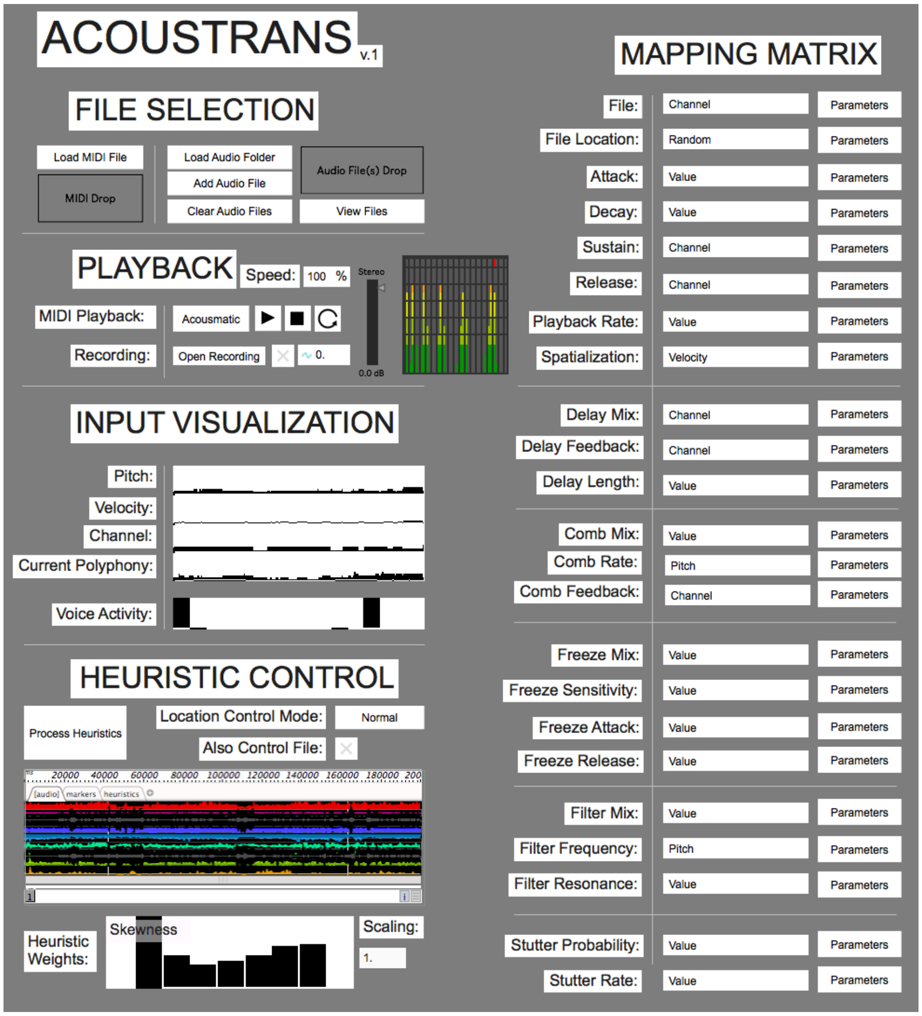

AcousTrans playback module, including file selection, playback, input visualization, heuristic control, and mapping matrix submodules

With acousmatic translation, different dimensions of the events are dynamically mapped to different parameters of sounds generated by *AcousTrans* using a mapping matrix. For example, the intensity of a source event may be used to dictate the volume of a destination event, or the stereo localization value of a source event may be used to dictate the spatialization speed of the destination event. These destination events take as sound material a user-selected corpus of audio files. Further electroacoustic abstractions may also be applied including delay, comb filtering, spectral freezing, filtering, and a probabilistic repetition (stutter) effect, the parameters of each being either set by the user or driven by different event dimensions.

More exhaustively, the musical parameters available are file selection, determining which file in the corpus the event is mapped to, file location, determining where in the file the event is mapped to, the individual settings of an ADSR envelope, playback rate, as well as spatialization. The spatialization parameter operates in three modes: stereo, 360º panning, and dynamic ambisonic mode. In stereo mode, incoming values simply spatialize the destination sound events from left to right through 2 channels. In 360º panning mode, incoming values spatialize the destination sound events in a circle around the listener through undecoded higher order ambisonics. In dynamic ambisonic mode, the mapping is more complex and expressive: if incoming values are high, the sound event's spatial trajectory will be distant and fast, whereas if incoming values are low, the sound event's spatial trajectory will be close and fast. Dynamic ambisonic mode is a useful way to map from stereo space onto higher order ambisonics, representing sounds emanating primarily from the left or right channels as highly localized, fast-moving sound sources.

The acoustic features embedded in each source event may also be used to select a quantitatively similar sound within the user-selected destination corpus via concatenative synthesis (effectively automating control of the 'file' and 'file location' parameters). Using a *k*-nearest neighbors search algorithm on a *k*-dimensional tree constructed from the acoustic features of segments of each audio file in the audio file corpus, the subvector of acoustic features for a source event is mapped to the most similar sound within the destination corpus.[23] At the time of this writing, there is no similarity thresholding procedure (i.e. producing silence when a source event does not have a highly similar destination event), so a match, however distant, will always be made. This is intentional: a forced mapping from one sound corpus to another, possibly significantly different corpus might lead to interesting, unexpected sonic outputs. The user may customize the weighting of the acoustic features used in the search via a multislider interface (bottom left of the Figure above), which can be useful to 'tune' the system depending on the particular source and destination sounds (for example, de-emphasizing F0 estimation if only using sounds with no clear pitch center).

The result of this process is an acoustic feature-driven mapping between the events in the source audio file and those generated by the system from the user-selected audio file corpus. Combined with the electroacoustic abstractions outlined above, this system can generate a diverse array of natural system-derived soundscapes.

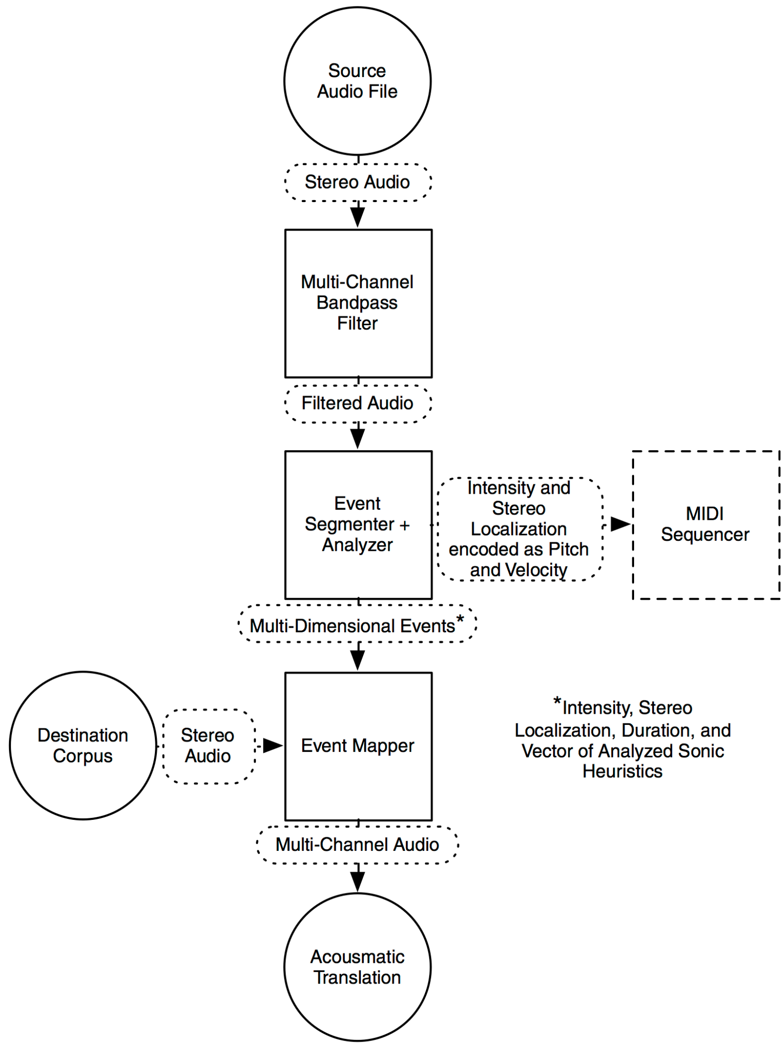

An overview of the functionality of *AcousTrans* as system diagram

NSSM Analysis

AcousTrans is a NSSM of Type 8 (semi-representative, sampled, acousmatic). Unpacking AcousTrans as NSSM, engagements with physical space include the natural systems that acted as impetus for the system: intertidal oyster reefs off the coast of Virginia. The output of the system may be encoded as MIDI files or diffused over high-density loudspeaker arrays (HDLAs) or listened to on headphones.

Engagements with compositional space include modeling Bernie Krause's 'Acoustic Niche Hypothesis'[24] to extract the events within sub-bands of environmental recordings and also curating the destination sound corpus. Mapping space includes the highly customizable soundscape event-to-musical parameter mapping matrix as well as the 'intelligent' mapping of soundscape events onto the most similar sounds within the destination corpus through a criteria of acoustic feature similarity.

The keys to these mappings might be revealed by presenting the source and destination in parallel, that is, playing back the source sound file and its AcousTrans-processed mapping in sync. Each event within the source sound (ideally) will be mirrored in the destination output, creating a chimeric extension of the source to a new timbral world. This key may also be used as part and parcel of electroacoustic composition technique: cross-fading between the source recording and destination recording, for example, to cause the listener to note and then remember the tight relationship between the two.

Future Directions

Future work includes testing out this software on a wide variety of environmental soundscapes, developing the segmentation algorithm to be more sensitive to dense, low dynamic range, 'lo-fi', environmental soundscapes, and expanding the type and parametric control of the acoustic features used.

Conclusion

Ultimately, *AcousTrans* presents a methodology for intelligently mapping a multi-dimensional stream of gestures from one environmental soundscape to an entirely different, multi-channel electroacoustic sound world. It harnesses techniques from both bioacoustics and MIR to facilitate the generation of electroacoustic material derived from the activity of natural systems.

Acknowledgments

Thanks to Environmental Sciences Ph.D. candidate Martin Volaric and the University of Virginia Environmental Sciences Department's Moore Award and the Jefferson Scholars Foundation for providing financial support for this research.

3.3a Virginia Barrier Islands Shorebird Sonification

Abstract

The Virginia Barrier Islands, a chain of islands hugging the coast of Virginia, are one of the world's premiere sites to see (and hear) shorebirds and seabirds. These islands are patchworks of diverse habitats including beaches, mudflats, marshes, and grassy areas. Hundreds of different species of birds, some colonial (local) and others migratory (passing through) call different habitats on these islands their home, either for their entire lives or just as a stopping point along their routes.

This project revolves around a spatially representative habitat-driven sonification that allows listeners to traverse an interactive, immersive model of birds on two of the Barrier Islands, driven by data sources including bird censuses, bird habitat preferences, and geographical information system (GIS) data of the islands. The author then sought a way to deploy this system within an electroacoustic music context, to position it as a habitat model-driven sound spatializer. Transitioning from documentary sonification to electroacoustic music sound spatialization paradigm engenders a number of questions: How can the multi-channel speaker array be used as a lens into a virtual space inhabited by sound materials whose spatial locations and dynamics model population dynamics? Further, what are the effective bounds of the perception of activity through such a sonic lens and what sonic material is best suited to it? The collaborative design and creation of the sonification system, its transformation into an electroacoustic music spatialization tool, and the musical deployment of this spatialization tool is discussed.

Introduction + Motivation

Off the coast of Virginia are fourteen undeveloped islands, the Barrier Islands, which not only protect the Eastern Shore's coastline from storm damage but also contain an incredible diversity of coastal wildlife. Among the wildlife that make these islands home are shorebirds and seabirds, over 100,000 of which visit the Barrier Islands each Spring, either as a pit stop or to breed, as they migrate across the Atlantic coast.[1]

As part of the author's involvement in the Ecological Methods Lab, an interdisciplinary working group formed at the University of Virginia during 2017-2018, the author was able to visit one of these Barrier islands, Cobb island, and to experience first hand the multiplicity of habitats and unique wildlife in these ecosystems.

Virginia's Barrier Islands, with the two islands focused on for this project, Hog and Cobb islands, highlighted

Because of the endangered species status of some of the birds on these islands (the piping plover, Wilson's plover, gull-billed tern, and American Oystercatcher, to name but several) there is a focus among many governmental and academic institutions to monitor the breeding success, population size, and general activity of these coastal birds. Not only does this data help biologists inform best conservation practices for these species, but it also acts as a more generalized environmental health indicator for these coastal ecosystems. These scientists, along with other natural scientists around the world, are also seeking out ways to engage the public in their research, to educate and to get involved the people who call these environments their backyards, the ecosystems of which they are implicitly a part. It is towards this goal that the author first began a project that represented Barrier Island shorebird data using one of the most recognizable features of these animals: their vocalizations.

Sanderlings near the edge of the shore on a Virginia Barrier Island

Design

As a first venture into creating a sonic model of these islands, the author took a seabird census database from the Environmental Data Initiative (EDI) that was created by a research team at University of Virginia.[2] This data set describes the populations of 95 species of seabirds and shorebirds on Cobb and Hog islands over multiple days. This data set was mapped onto a database of the vocalizations of those 95 species, some recorded by the author but most downloaded from the Xeno-Canto wild bird recording database.[3] This mapping was first done with the handle of species size being tied to the destination of bird vocalization volume: the more birds that were seen on a chosen day, the louder that bird's vocalization was in the texture. Metaphorically, it was as if the more members of a species were seen, the higher the volume slider for that species was on a mixer outputting all of the birds' vocalizations.

Seeking a better representation of species multiplicity, species count was then mapped to the number of instances of a granular synthesizer outputting grains from a given bird vocalization recording: the more birds that were seen on a chosen day, the denser their granular 'crowd' was. Upon hearing this sonification, the author's collaborator and fellow member of the Ecological Methods Lab, Environmental Sciences Ph.D. Candidate Alice Besterman, noted that it was unrealistic, because all of the birds were being heard together, as though they could all be in the same habitat at the same time. In reality, certain birds would only be heard with certain other birds they shared habitats with (beaches, mudflats, marshes, grassy areas, etc.).

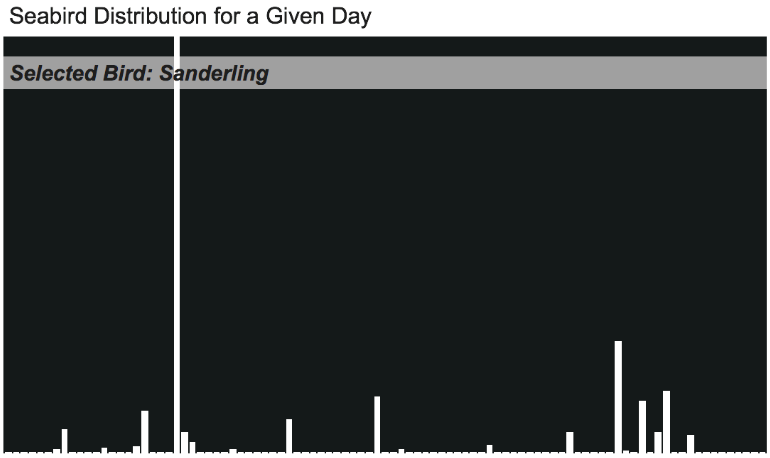

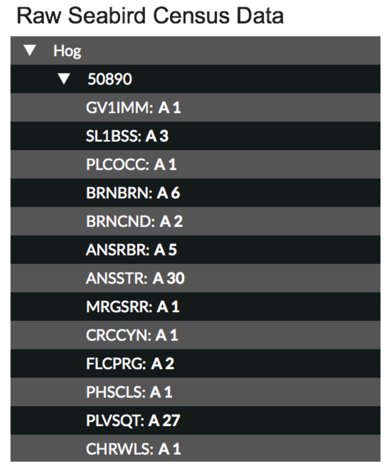

Histogram of bird species populations (left) and corresponding census data (right)

While sonically interesting and likely a very surreal experience for a coastal ornithologist, the author sought a method to make the representation of the birds on the islands more realistic, to work towards distributing the birds in an ecologically informed way. Through discussions with Besterman, the author determined that data that indicated the different habitats on these islands along with the habitat preferences of different species of shorebirds and seabirds were both available.

Linking these data together resulted in a system that assigns each shore bird a location on either island that is true-to-life, that takes into account that bird's habitat preferences and each location on the map. Given a particular census day, then, a 2-dimensional map populated by the number and type of seabirds seen on that day, each located accurately based on their habitat preferences and the habitats on the island, can be made. This 2-dimensional map of locations can then be represented in sound over a multi-channel loudspeaker array, with birds located at the top of the island heard in front of the listener, the birds on the East side of the island heard to the right of the listener, the birds on the West side heart on the left, and at the bottom of the island heard behind the listener, engulfing the listener in a soundscape they would hear if they were standing in the middle of the island and had fantastically exceptional hearing. Zooming into different areas on this virtual sonic island allows a listen to isolate a particular habitat and to hear only the birds that co-habitate it.

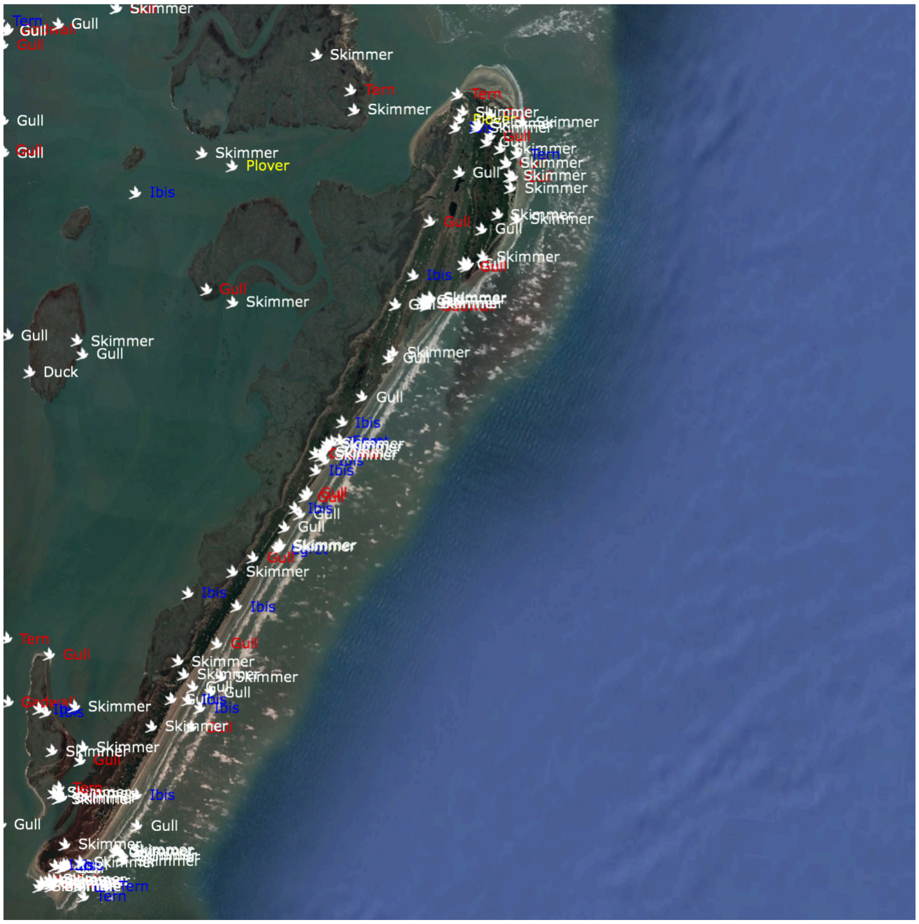

Visualization of birds position on Hog Island based on its different habitats and each bird’s habitat preferences

Implementation

The specifics of implementing this sonification system involved a number of data mapping and analysis tools. First, images of global information services (GIS) data of the two islands were scraped from a GIS database, images in which the color of pixels in the image correspond to different habitat types (what are called ‘land usages’) on the islands.

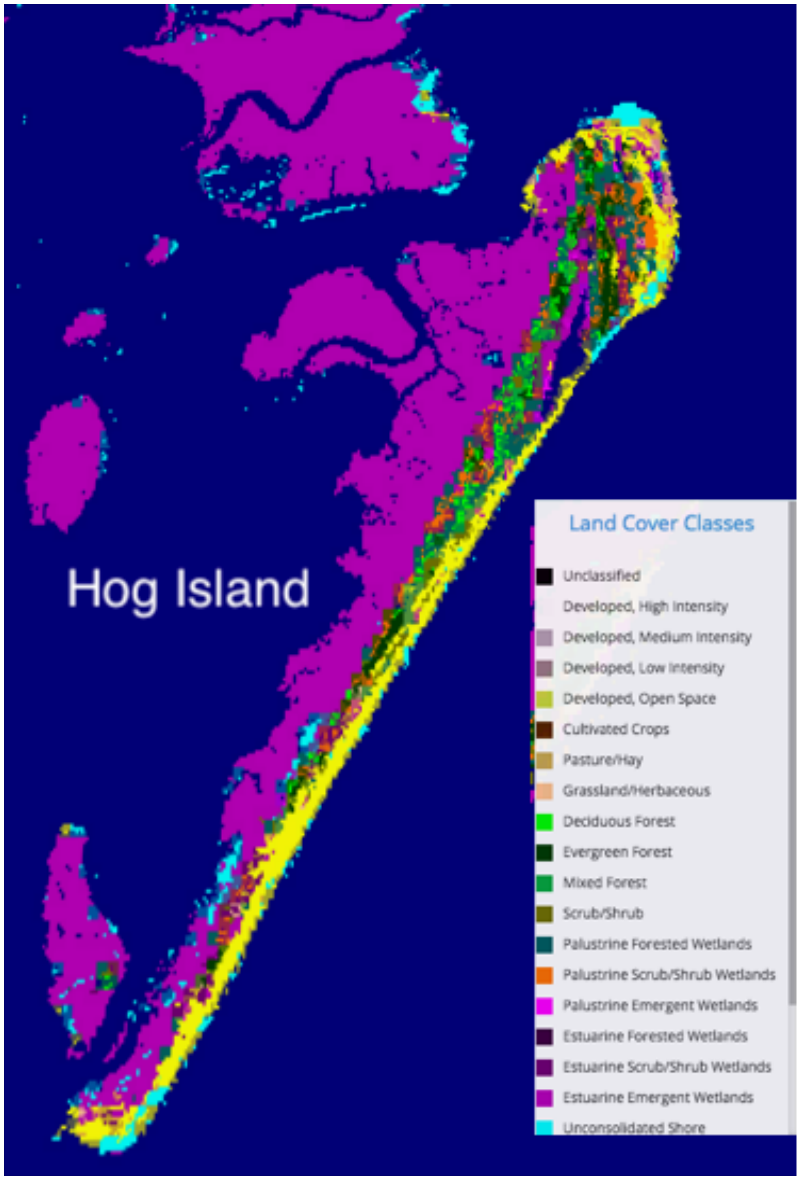

Hog Island GIS data, including a key of Land Cover Classes

Second, bird species habitat preferences data was sourced, consisting of a dictionary of different bird species and the habitats that they have been seen to be active within. In order to get these two datasets into forms that interfaced well with one another, the the author first created a script that scanned a downsampled version of each GIS image and classified and stored all of the pixel coordinates sorted by habitat type. With the help of Besterman, a one-to-one map was then made between the land cover classes and the habitats listed in the habitat preferences data, with some classes (Palustrine Forested Wetlands, Palustrine Emergent Wetlands, etc.) being merged into a single class. All of the different habitats on the islands were now effectively mapped to sets of shorebirds.

Common Name	Upland	Bank	Bay Island	Beach	Mudflat	Rocky Intertidal	Salt Marsh	Brackish Marsh	Tidal-fresh Marsh	Shallow Water	Deep Water
Tricolored Heron	F,L,R		B,F,L,R	F,L	F,L		B,F,L,R	F,L,R		F,L	
Little Blue Heron	F,L,R		B,F,L,R	F,L	F,L		B,F,L,R	F,L,R		F,L	
Cattle Egret	F,L,R		B,F,L,R	F,L	F,L		B,F,L,R	F,L,R		F,L	
Green Heron	B,L,R		B,F,L,R	F,L	F,L	F,L	B,F,L,R	F,L,R	F,L,R	F,L	
Black-crowned Night Heron	B,L,R		B,F,L,R	F,L	F,L	F,L	B,F,L,R	F,L,R	F,L,R	F,L	
Yellow-crowned Night Heron	B,L,R		B,F,L,R	F,L	F,L	F,L	B,F,L,R	F,L,R		F,L	
Sandhill Crane	F,L,R										
King Rail			B,F,L,R		F			B,F,L,R	B,F,L,R		
Clapper Rail			B,F,L,R		F		B,F,L,R	B,F,L,R			
Virginia Rail	B,F,L,R		B,F,L,R		F		B,F,L,R	B,F,L,R			
Sora	F,L,R		F,L,R		F		F,L,R	B,F,L,R	B,F,L,R		
Yellow Rail	F,L,R						F,L,R	F,L,R			
Black Rail			B,F,L,R				B,F,L,R				
Common Moorhen			B,F,L,R				B,F,L,R	B,F,L,R			
American Coot			R		F,L			F,L,R	F,L,R	F,L,R	F,L,R
Wilson's Phalarope								F,L,R		F,L,R	
American Avocet			B,F,L,R								
Black-necked Stilt			B,F,L,R		F,L		F,L,R			F,L	
Wilson's Snipe	F,L,R				F,L,R			F,L,R	F,L,R		
Short-billed Dowitcher	F,L,R				F,L,R		F,L,R	F,L,R			
Long-billed Dowitcher	F,L,R				F,L,R		F,L,R	F,L,R			
Stilt Sandpiper	F,L,R				F,L,R		F,L,R	F,L,R			
Red Knot			F,L,R	F,L,R							
Purple Sandpiper				F		F,L,R					
White-rumped Sandpiper					F,L,R		F,L,R				

Species Habitat Preferences Table, adapted from Bryan D. Watts' "Waterbirds of the Chesapeake Bay: A Monitoring Plan"

Lastly, the shorebird census data was then connected to the parsed GIS data and the species habitat preferences data, to get an ecologically informed placement of each bird on the islands. Outlining this process in its entirety:

1. After selecting a day, the number of each species of bird (if present on that day) is passed to the habitat preferences data
2. For each bird of each species present on that day, the habitat preferences data outputs a vector of potential habitats for that species
3. One of those habitats (land cover classes) is chosen randomly
4. That land cover class is sent to the parsed GIS data, which generates a pixel coordinate
5. Each bird of each species present on that day is then mapped to a pixel coordinate on the map

The result of this process is an ecologically informed, 2-dimensional sonification of bird vocalizations made in Cycling ’74’s Max 8. This sonification design prevents birds that would otherwise not be seen (or heard) near one another from being within the same local soundscape, and presents this soundscape as an island-specific mapping on to 3-dimensional sonic space.

In addition, a dynamic environmental soundscape is automated using computer vision to accompany navigation of the islands: depending on where the user is zoomed into the levels of looped recordings of beach, open water, and marsh are altered. For example, if the user is zoomed into a beach (primarily yellow-ish white) the beach recording will dominate, but if the user moves to open water (primarily blue-ish green) then the open water recording will be turned up, and so forth.

As of this writing, the standalone Max version of this software is being ported to an online interface viewable/listenable through the web as part of a commission from the Coastal Futures Conservatory at University of Virginia.

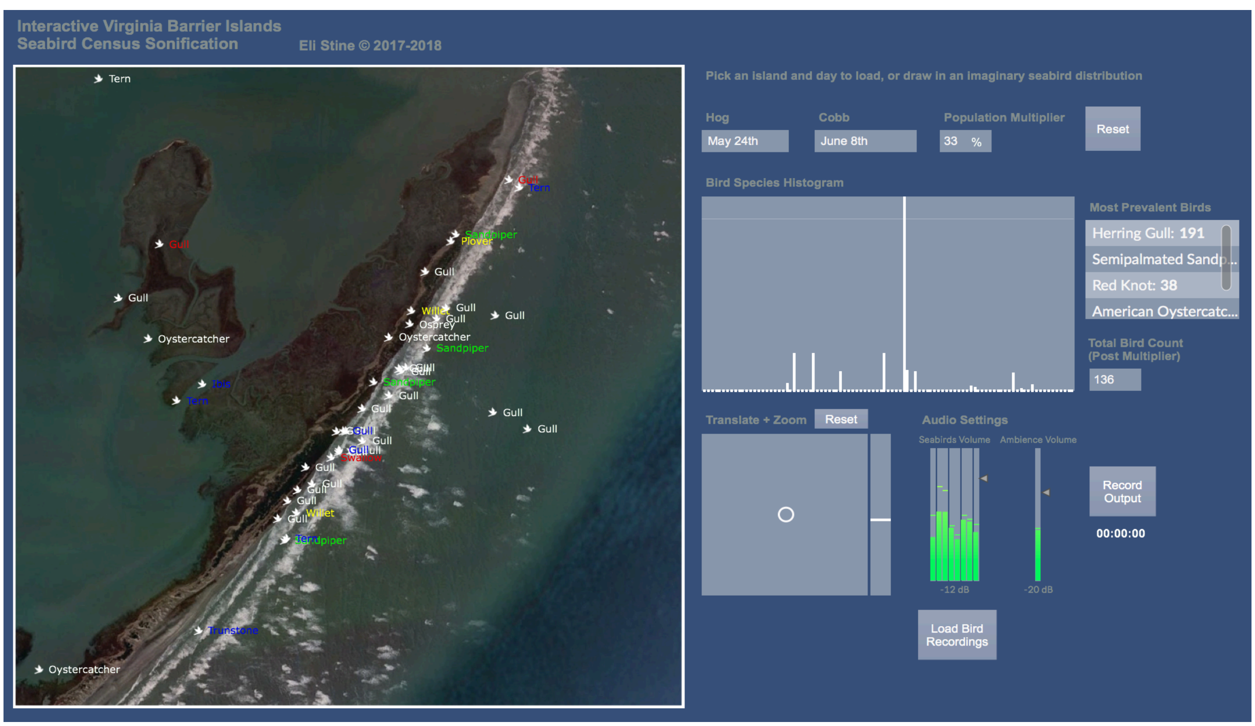

The Interactive Virginia Barrier Islands Shorebird Census Sonification GUI

3.3b HabiSpat

Introduction

During the process of creating the shorebird sonification described in the previous section, the author experimented with loading in different folders of sound files than just the bird vocalizations. This is a simple, playful act, but at the same time is also profound, shifting the system away from a documentary sonification to something much more generalized: a methodology for spatializing sound 'species' in accordance with species localization preferences. Because of the constraints of the system, the author had to load in folders with exactly 95 audio files. Some experiments included 95 bands of a filter bank applied to a drum loop, 95 randomly selected short sounds from the author's hard drive, and collections of room tones (background sounds of spaces recorded for film sound design post-production). In addition, the author also made scripts that were able to simultaneously generate 95 instances—all slightly or, significantly, different—of a particular processed real-world sound or synthesized texture.

Design + Implementation

A new systemic caricature software, *HabiSpat* (*Habitat-based Spatializer*), was then created, which extended the shorebird sonification in a number of significant ways. One first pass, the constraint to 95 species was lifted, the species habitat preferences were made easily customizable (not being scraped from a data set), and the topology of the virtual 'island' (the distribution of habitats within virtual space) could now be designed from scratch. In addition, the day-based method for loading species distributions was extended to be dynamic. Both over both time, through simulations of predator-prey dynamics, and in space, through a simple implementation of topological dynamics (emulating erosion or natural disasters, for example).

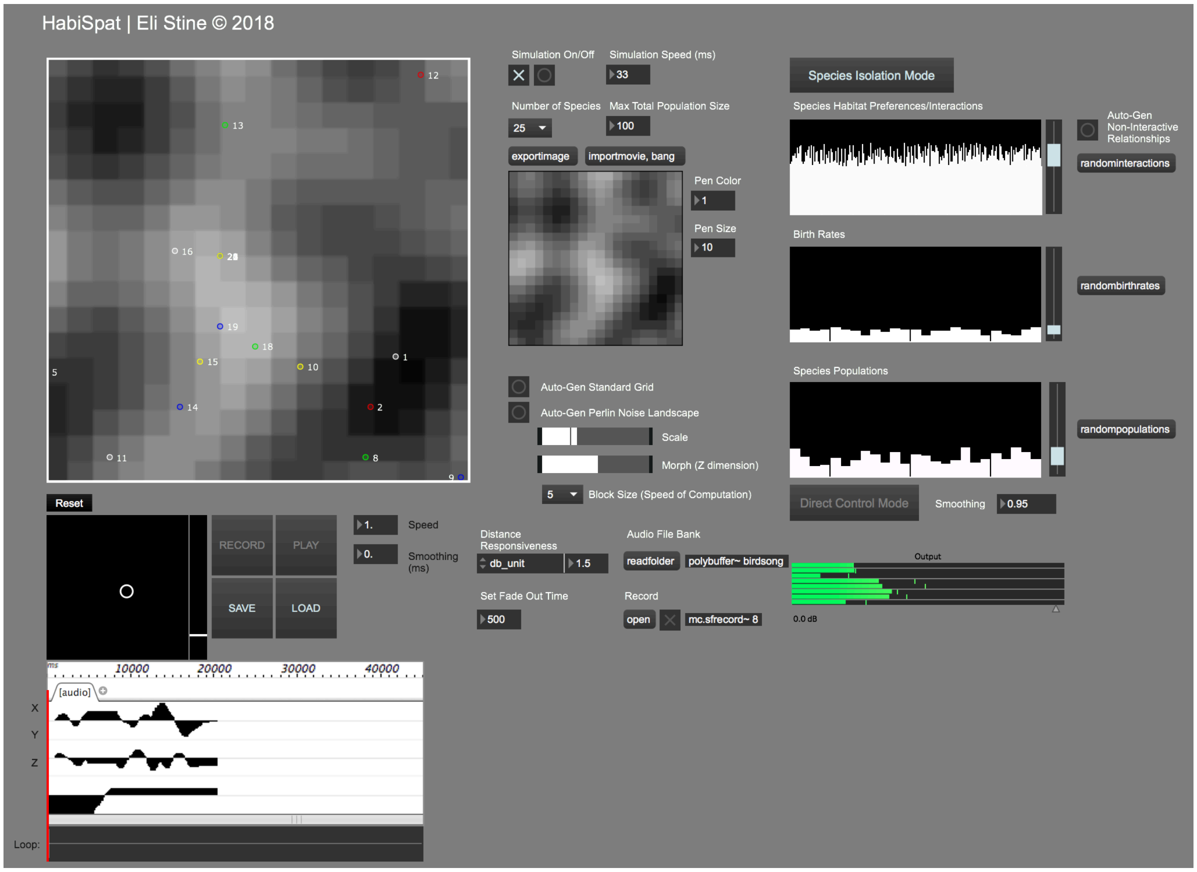

HabiSpat GUI, including Perlin noise-generated landscape (left), control over species interactions, birth rates, and populations (right)

Synthetic Landscape Design

Users of HabiSpat may generate virtual landscapes through three methods. As with the shorebird sonification software, different colors correspond to different habitats. Virtual landscapes may be saved and loaded (as images). First, users may draw in completely customizable habitats using a simple pen-like interface. Second, they may generate a standard grid, akin to the square fields on a farm, which generates a grid of different colored squares that fills the landscape. Alternatively, posterized Perlin noise can be used, creating more naturalistic, ‘map-like’ topologies of habitats which may or may not be contiguous. The scale of this Perlin noise landscape can be altered, a change comparable to subtracting or adding complexity to a topology, or zooming in or out of a map, and its Z dimension can be shifted, with small changes emulating the erosion or general reorganization of a landscape over time.

Perspectival Landscape Navigation

A user can change their perspective in relation to this synthetic landscape through a joystick-like interface, affording X and Y translation and zooming in and out. Traversal of the landscape can be recorded, edited, saved and loaded, and also played back and looped with different start and end points and at different rates with varying levels of movement smoothing. This allows perspectival navigation through the virtual landscape to be able to be treated as an offline process, similar to automating stereo panning in a DAW, but instead of a single sound's position within the sound field being altered, traversal of the landscape changes the listener's spatial relationship to potentially hundreds of sounds.

Species Population and Interaction Simulation

After a virtual landscape has been designed, it may be populated by first choosing the number of species, which must be identical to the number of audio files in the corpus, and then the maximum total population size, which limits the total number of species that can populate the landscape. In addition, the species habitat preferences, that is, how likely they are to populate a particular habitat (color) within the virtual landscape, must be set. This may be done by drawing into a multislider interface (with each bar of the multislider corresponding to the probability that a particular species will be placed within a particular habitat), by pressing a button which assigns each species to their own habitat, or through randomization.

After the species habitat preferences have been set, two modes control how the species populate the virtual landscape: 'Direct Control Mode' and 'Simulation Mode.' In Direct Control Mode, a user may simply draw in the number of each species into a histogram interface: the species (as sounds) are then dynamically distributed across the virtual soundscape in accordance with the landscape's topology and the set habitat preferences, resulting in an immersive soundscape of different sounds in 2-dimensional space.

In Simulation Mode, an author-created implementation of solutions to equations that have been shown to model population dynamics based on specified interactions between species (predatory, symbiotic, etc.) along with birth and death rates, are applied to automate species population dynamics over time. These equations are *N*-species competitive Lotka-Volterra equations, and they, like the natural systems they simulate, are chaotic and complex.[4]

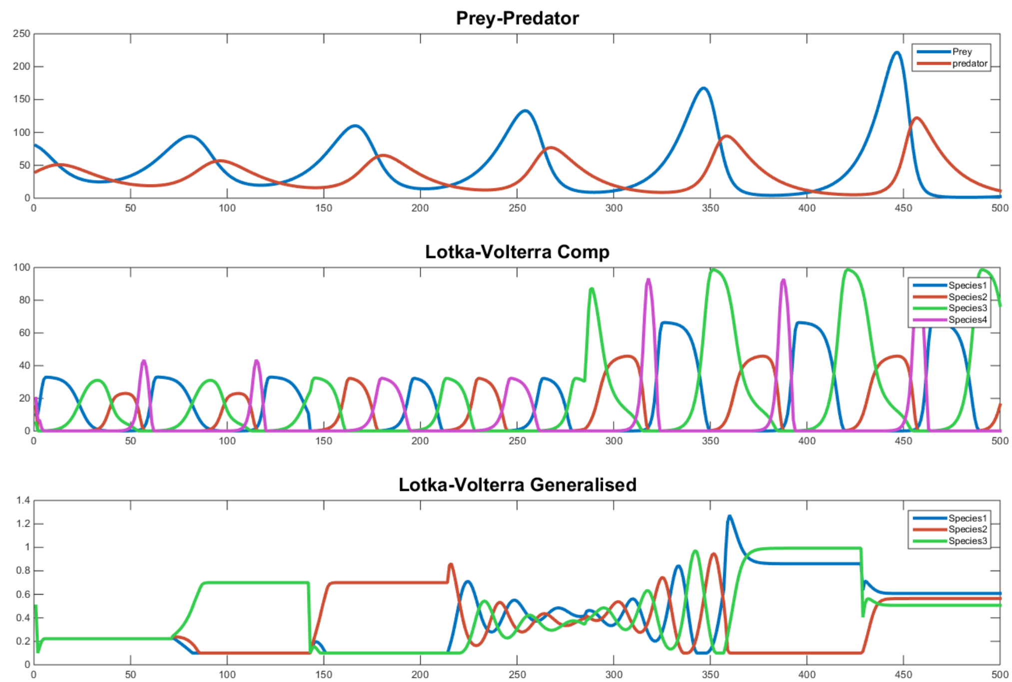

A visualization of the dynamics of different predatory-prey-related simulations over time, after Fiore[5]

In this mode, simulation speed must be set by choosing a time step for each simulation state (in milliseconds). The birth rates of each species must also be set by hand or through randomization. The simulation then takes control, altering the populations of the different species over time according to iterated solutions to the *N*-species competitive Lotka-Volterra equation, formalized as

$$\frac{dx_i}{dt} = r_i x_i \left(1 - \frac{\sum_{j=1}^{N} \alpha_{ij} x_j}{K_i}\right),$$

where x_i is the size of species i's population, r_i is the birth rate of species i, K_i is the maximum population size of species i, and a_{ij} represents the effect species i has on species j. How these effects are calculated depends on whether the software is in 'Habitat Roaming Mode' or 'Species Isolation Mode': in 'Habitat Roaming Mode' there is no effect of species on other species (population dynamics are a function of species carrying capacity). In 'Species Isolation Mode' the effect species i has on species j is dependent on the probability of a species to be found within a habitat outside of its assigned habitat (e.g. how likely species 1 is to be found in species 2's habitat), set through the Species Habitat Preferences/Interactions interface. In this mode, self-interacting terms (e.g. a_{ii}) are set to 1.

Left to run, and depending on the time step used, the spatialized soundscapes that this simulation produces might represent millions of years within a few seconds or a year or two over minutes. Different species and corresponding sounds are foregrounded and pushed to the background of the soundscape over time, in accordance with predator-prey behavior or simply over-population (depending on the mode). Sound localization as a product of species habitat preferences spatially reinforces these dynamics, and, if the Z dimension of a Perlin-noise generated virtual landscape is changed while the predatory-prey simulation runs, the habitats themselves dynamically evolve, resulting in parallel simulations of both species interaction and topological transformation.

NSSM Analysis

HabiSpat is a NSSM of Type 17 (semi-representative, simulated, acousmatic). Unpacking HabiSpat as NSSM, engagements with physical space include the natural system of spatial distribution and population dynamics of species according to habitats and habitat preferences, generally, and the distribution of shorebirds on Virginia Cobb and Hog islands, more specifically. The output of the system may be diffused over high-density loudspeaker arrays (HDLAs) or listened to on headphones.

Engagements with compositional space include the curation of sound species, the adaptation of a sample-based paradigm (data sets in the sonification software) to a simulation-based one in HabiSpat (tools to design virtual landscapes, flexible control over species and species habitat preferences), and the implementation of the predatory-prey modeling Lotka-Volterra equations. Within mapping space, the type and location of members of different species is mapped to sound localization and different audio files and the population dynamics of the Lotka-Volterra equations are mapped to species populations over time.

The transparency of the keys to these mappings is difficult to evaluate without a quantitative perceptual study. The localization of two or more sounds in space, if not timbrally disjunct, may be perceived as a unified ‘environment’ rather than separate, differently localized sounds. This potentially reduces the perspicuity of the mappings within this system, but again, no quantitative study has been done to determine how well listeners can hear changing population dynamics through sound. This system also has a visualization which, particularly in the case of the shorebird sonification, is helpful for listeners unfamiliar with auditory display contexts to be able to hear the distribution of the shorebird vocalizations.

This system also involved a *parallel perceptual model*: that of the author’s collaborator, who perceived the real-world incongruity in a non-spatially distributed sonification of the shorebird population data sets, suggesting the ecologically informed spatial methodology that is at the core of this software system.

Future Directions

Future directions for this software include updating the interface to be more streamlined and less cluttered, adding on more sound generation possibilities (synthesized sound, granular synthesis, etc.), and also potentially conducting a study (perhaps in conjunction with a cognitive scientist) on how well, if at all, people can hear population dynamics over time using these systems.

Conclusion

Created within the context of an interdisciplinary ecological research group, this project began first as a quest to represent ornithological data in a way that was both intuitive and accessible and that also privileged sound, primarily the vocalizations of shorebirds, as a way to get people to connect with their populations and ecosystems. Through the implementation of an ecologically informed 2-dimensional mapping of these birds onto their island habitats the author recognized a unique paradigm for the spatialization of sound, and in turn explored ways in which such a system could be applied to the non-scientific context of multi-channel acousmatic music. The result is a system that spatializes sounds using a simulation of species habitat preferences, either to create fixed distributions of sounds in 2-dimensional space or in conjunction with predatory-prey simulations and simple emulations of topological dynamics to generate dynamically changing immersive soundscapes.

Acknowledgements

Thanks to Environmental Sciences Ph.D. candidate Alice Besterman and the University of Virginia Institute of the Humanities & Global Culture (IHGC) for providing financial support for this research.

3.4 Musical Applications

In addition to the *Murmurator*'s live performance over high-density loudspeaker arrays, the author has composed several fixed multi-channel electroacoustic compositions using material generated by The *Murmurator*, *AcousTrans*, and *HabiSpat*. These include *No Where* (2018) for octophonic fixed media, *Vestigial Wings* (2018) for video and higher-order ambisonics (full 3D 5th order HOA), and *Artificial Reef* (in progress) for higher-order ambisonics (full 3D 7th order HOA). What follows are analyses of *No Where* and an outline of the work-in-progress *Artificial Reef*, demonstrating how these NSSM softwares may be used within an electroacoustic composition process.

To listen to/watch these works, to see systemic demonstrations, and to download the code of these softwares please visit www.elistine.com/diss.

3.4.1 No Where

No Where is a 7 minute octophonic fixed media work composed from January to May of 2018 in the Virginia Center for Computer Music. This work shifts the listener between different places and spaces, some authentic, others synthetic, most only inhabited for moments before transporting to another. Ambience tropes (for example, filmic tropes of what archetypal spaces—restaurants, carnivals, offices—sound like) and impossible deformations of recorded and virtual spaces (pushing the ceiling beneath the floor, for example), are juxtaposed and interposed to dis- and un-place the listener. This work makes substantial use of both *The Murmurator* and *HabiSpat*.

No Where consists of three layers of sound spatialization encodings: first, an octophonic layer that directly addresses each speaker, second, a layer of first-order ambisonics (encoding 360° sound through 4 channels), and third, a layer of third-order ambisonics (encoding 360° sound through 16 channels, offering more precise sound localization than the first-order layer). When the piece is performed these layers are played back simultaneously. Each layer contains different sound materials and has different roles in the story-telling and spatialization of the work.

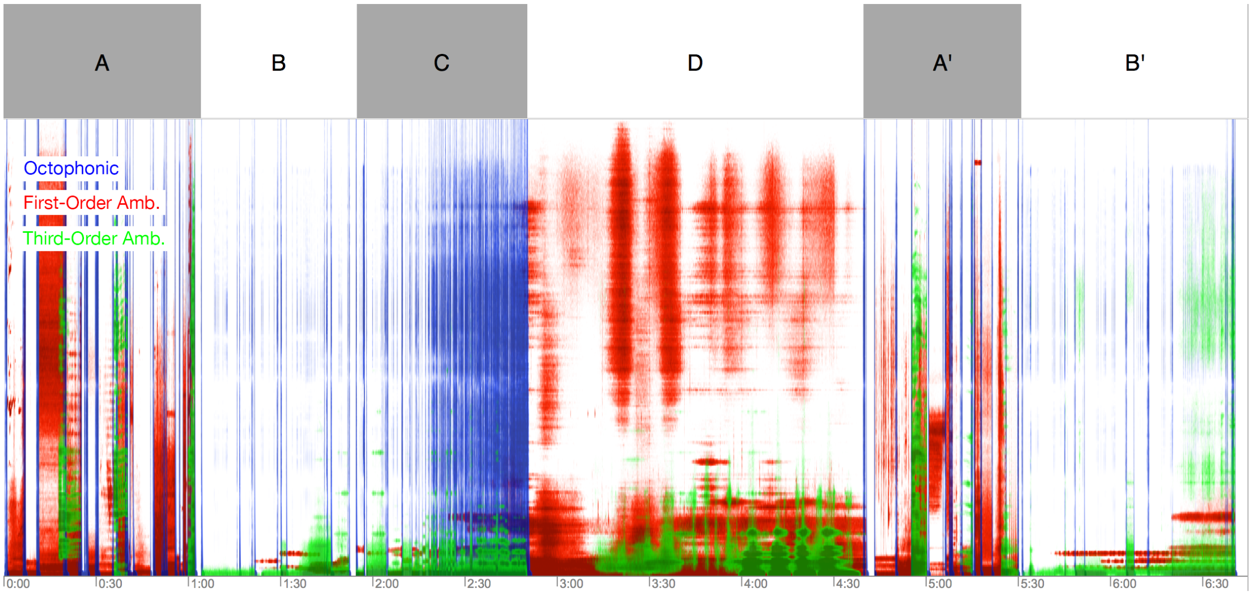

Formal analysis of *No Where*, including its three spatial-encoding layers

A

The work begins by introducing a 'space shifter' device—the sounds of a close-mic'd tape deck being opened, closed, the author pressing play, record, etc.—which is diffused in mono through the octophonic layer (the vertical blue lines in section A). This device 'shifts' through a wide variety of different 'places' represented through first-order ambisonic recordings (made with the Sennheiser *Ambeo* VR mic), or stereo recordings encoded in first-order ambisonics (red in section A). These include recordings of city streets, the interior of cars, and parks. Fragments of these recordings are put into the third-order layer and further abstracted using looping, tremolo, and exaggerated spatialization (green in section A).

B

Section B retains the 'space shifter' clicks but introduces another sound, a 'space tuner,' created from the sound of zip ties being closed (the intermittent vertical blue lines in the middle of sections B and C). The 'tuning' that the 'space tuner' does is of a studio performance of *HabiSpat*, in particular a *HabiSpat*-facilitated navigation of a spatially distributed corpus of room tones (ambient recordings of different places) (green in section B). As this distribution of room tones is navigated, resultant sonic changes—room tones getting closer, further away, spreading out, sudden switches of location—are mirrored with the 'space tuner.' In addition, a layer of harmonic material, composed freely in response to the resonances of the room tones, is added (red in section B).

C

Section C is denoted by a shift in the harmonic material—from consonant to more dissonant—which is mirrored by a speeding up of the ‘space shifter’ sounds. The clicking sounds of the ‘space shifter,’ first mono, then diverge into a cacophonous octophonic texture (blue at start of section C). This noisy texture begins to resonate, almost as though feeding back, an effect which was created by reproducing the iterative record-and-playback process of Alvin Lucier’s *I am sitting in a room* using an octophonic loudspeaker array and 360° microphone. Specifically, the sound material of section C was played back through the octophonic array and recorded using the 360° microphone, that recording was then played back through the array and recorded again, that recording was then played back and recorded, and so on. The results of this process are multiple versions of the material of section C, ranging from completely dry to versions overwhelmed by resonant acoustics of the room housing the octophonic array. These different copies were then faded one into the other—from completely dry, inside the box, to resonant and nearly unrecognizable—to create section C (blue at end of section C).

D

Section D is a decided break from the fast, collaged material heard up until this point in the work. Instead, the last sounds of section C (the resonant acoustics of a room represented through 8 channels) are put through a first-order ambisonics reverberation plug-in (Bruce Wiggin’s *AmbiFreeVerb 2*) on its ‘freeze’ setting, resulting in a smooth, engulfing wash of sound (red in section D). Over the course of section C this wash, and other washes created by putting different sound materials into the infinite reverb, are modified through composer-automated filtering and ambisonic spatialization. In addition, this section introduces the sounds of the *Murmurator*, encoded in third-order ambisonics (green in section D). Vocal and church bell samples that meld with the harmonic content of the washes, along with samples of the washes themselves, were loaded into the *Murmurator* and performed by the author to accompany the washes. This section ends with four repeated overhead, front-to-back sweeping gestures produced by the *Murmurator*, which lead the listener into the last two sections of the work (green in bottom right of section D).

A’

This section has the same gestural vocabulary and texture—the ‘space shifter’ moving between different recordings—as section A, but with different sound recordings. These include recordings of a cave, a cafe, and rain, along with the only comprehensible speech within the work, a recording made by the author on a plane to Brazil in 2016. This section is also slightly faster-paced and shorter than section A, contrasting with the listener’s memory and creating momentum towards the final section of the work.

B’

Section B’ returns to the sounds of the ‘space shifter’ and ‘space tuner’ in combination with the navigation of the room tone landscape using *HabiSpat* heard earlier in the work (section B). The harmonic material presented earlier in section B (red in section B’) is developed and also used as material for the *Murmurator*, the output of which contributes to an immersive, constantly fluctuating texture cycling around the listener (green at end of section B’). At the end of this section, the volume of the room tone landscape and the *Murmurator* output is raised, encapsulating the listener within multiple real (room tone recordings) and synthetic (*Mumurator* output) sonic locales before the composition ends with a loud, final click from the ‘space shifter’ (vertical blue bar at end of section B’).

Conclusion

This composition demonstrates how *HabiSpat* in particular, and the *Murmurator* in a more supporting role, may be used in the context of electroacoustic story-telling. Recordings of the author performing the *HabiSpat* interface are curated and accompanied through other layers—in this case, the ‘space tuner’ and harmonic materials—that mirror and complement, respectively, the gestural content of the recorded *HabiSpat* performance. The result is a spatial environment that is rich, consisting of upwards of 100 room tone recordings at given points, but that was algorithmically constructed and expressively performed in real-time. The *Murmurator*’s ability to create both immersive granular textures (as in section B’) and compelling large-scale spatial gestures (as in section D) is harnessed in this work, the former enhancing a pre-existing texture and the latter taking the foreground and indicating the end of a section.

3.4.2 Vestigial Wings

Vestigial Wings is a 5 minute work consisting of a short poem, a video art setting of that poem, and higher-order ambisonics electronic sound (full 3D 5th order HOA). It was created from October to December of 2018 in the Virginia Center for Computer Music and the Thomas Jefferson Starship (a small studio at the University of Virginia). *Vestigial Wings* first started as a work of sonic art, which then inspired the poem, which then inspired the visual component. The poetry written for the work is inspired by the work of poets whose writing the author has set in the past, predominantly feminist: Sylvia Plath, Muriel Rukeyser, Edna St. Vincent Millay, and William T Barry. The poem:

```
At the boundary of the desert
Beneath the telescopic sky
I stopped to take the world in
As it went on rushing by

I thought ten hundred futures
Of what could and would become
As the dark of night got closer
Slipping disk of orange sun

I thought of all I'd loved and lost:
Of dropped, forgotten things
Of books with unread pages
Broken roots, vestigial wings

I thought of names gone unremembered,
And of places never seen,
Of the last of every species,
Silent forests, noiseless seas

And as dusk made way to nightfall
Black sky pricked with yellow light
I had not moved a single muscle
And so doing lost my life

Because in thinking and not doing
All I did was just compare
What could and would become of
Rather than what was really there
```

This poem consists of both visual imagery and a general message: that over analysis of the past and/or the future will result in neglect of the present. The visual imagery traces an experience of time passing while looking over a desert: morning slipping to afternoon, afternoon, to evening, and evening to night, a metaphor of the passage of a lifetime, while the speaker remains a thinking onlooker rather than an active participant.

The sonic component of *Vestigial Wings* focuses on the unintended sounds at the ends of sample bank recordings (for example, the *Kontakt* piano sample packs, *Logic* sample packs, etc.), the unavoidable room sounds and other sounds that the makers of sample libraries try to erase but are always present. In addition, recordings of bird's wings, room tones, toy pianos, and glockenspiels that the author made, in addition to sounds generated by Vocaloid software (synthesizing spoken/singing voices), are also incorporated. These sounds are processed using Bhob Rainey's *BoomerangPointer* and Iain McCurdy's *Grain3* scripts, both programmed in CSound, Paul Nasca's *PaulStretch* Algorithm (as implemented in Audacity), Nuno Fonseca's *Sound Particles* software, and the author's *Murmurator* software, described earlier in the Project Descriptions section. These processed sound materials were then arranged and spatialized in Cockos' *REAPER*. The visual component of *Vestigial Wings* was created using *Cinema4D Studio R19*, a 3D motion graphics software.

The visual component has six sections which roughly follow the stanzas of the poem, depicting the journey of an unsettlingly black square (perhaps referencing Vantablack, one of the darkest substances created, developed in 2016) within a set of virtual landscapes with evolving lighting configurations. *Vestigial Wings* has been programmed on the 2019 International Computer Music Conference/New York City Electroacoustic Music Festival in New York, New York and on Digitalis 2019 at the University of Virginia in Charlottesville, Virginia.

Section I: Desert Sunrise

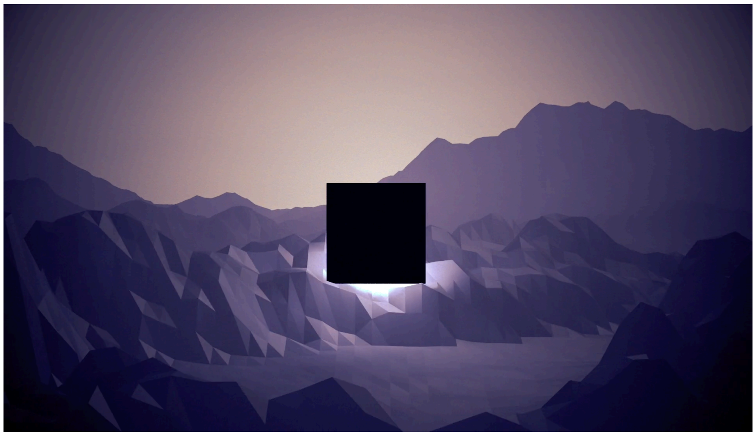

Vestigial Wings begins in darkness with recordings of piano resonance (from the *Alicia's Keys* Kontakt sample library). As the sun rises a black square enters stage left and a gentle swarm of granulated piano recordings (created using *Sound Particles*) surrounds the listener. At 0:45, as the sun can be seen behind the synthetic mountains, a breathy processed Vocaloid-generated sound fades in and out. As the black square moves across the screen and exits stage right the screen cuts to black and the swarm of granulated piano recordings follows suit, gently fading to nothing.

Section II: Desert Sunset

The piano resonance recordings return, accompanied by low, bass-y sounds, as the same landscape and black square, now at sunset, appears. Processed recordings of the resonance and tails of glockenspiel notes enter the texture, along with a return of the granulated piano recordings at 1:28. As the sun falls behind the synthetic mountains and the light changes from orange to dark purple the breathy, Vocaloid-generated sound returns, this time louder and more human. At 1:55 an immersive granular texture enters, akin to a thick torrent of sand, the output of the Murmurator processing a recording of a bird's wings. This sound builds to a small crescendo and suddenly, with a delicate, metallic ping, the viewer is thrust forward *into* the black square and engulfed in darkness.

Section III: Nightland

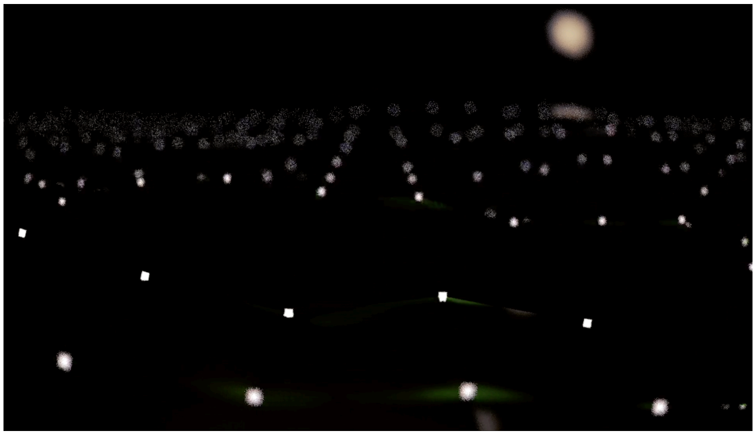

Fading up from black the viewer finds themselves a small distance above (and sometimes passing through) an expansive ocean carpeted with small lanterns, which gently rise and fall beneath a synthetic moon. The delicate metallic sounds, derived from the same bird wing recording processed using Rainey's *BoomerangPointer* CSound script, are overlaid with another processed version of that same sound, a recorded performance of the author using McCurdy's *Grain3* CSound script. Gestures made by arranging the glitching, warbling output of these two scripts interact as a faster-paced presentation of the piano resonance recordings enter, all the while the ocean of small lanterns pulsates beneath the light of the moon. The sonic texture begins to attenuate as the viewer gets closer to the digital ocean, then suddenly plunges beneath the waves into darkness.

Section IV: The Wormhole

The viewer rises to find themselves within a fast-moving cave or tunnel, the color of terra cotta and filled with light. A dense version of the granulated piano resonance recordings (again created with *Sound Particles*) fades in, layered on top of the sounds generated by processing the bird wings recording using the CSound scripts and the *Murmurator*. The Vocaloid-generated sound returns, this time un-processed and distinctly vocalistic, accompanied by a sequence of low bass tones, pitched-down reversed low piano notes, as the tunnel pulsates and shakes. The viewer gets closer and closer to the end of the tunnel, a white light, and is then enveloped within it.

Section V: Return to Nightland

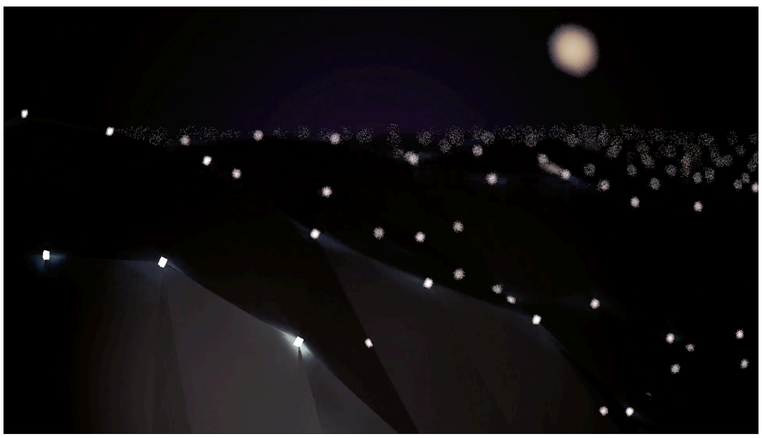

The soundscape and virtual landscape of Section III ('Nightland') returns, this time accompanied by less-processed recordings of toy pianos and glockenspiel along with a delicate, ethereal Murmuration of the piano resonance recordings. The viewer continues to float over the ocean as at 3:43 a flash of white light coincides with the sonic texture thinning, leaving only the sound of the Murmurated piano resonance recordings. As the sky lightens and a dark blue light begins peaking over the horizon, several piano chords (processed to be more synthetically resonant using multiple iterations of convolution reverb) rise and fall. The breathy processed sounds of the Vocaloid-generated material heard in Section II enter the texture and a final piano resonance, overlaid with reversed toy piano and glockenspiel recordings, builds to a hard cut.

Section VI: Winter Plains

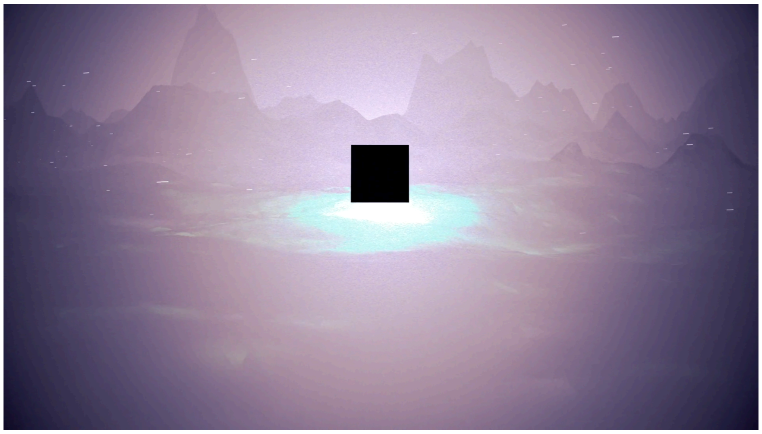

The black square returns, now within a landscape of high mountains, windswept and in the process of being covered in snow. As the processed Vocaloid-generated sounds fade out, a sequence of room tones, distributed around the listener, enters, shifting the listener into and out of different spaces. At 4:39 a distant beeping sound enters and repeats, fading to nothing, as a last iteration of the Vocaloid-generated sound (heard in Section I) fades in and out. The landscape is presented in silence for a few seconds before cutting to black to end the work.

Conclusion

Vestigial Wings is a multi-faceted work, engaging the same artistic object through poetry, sonic art, and video art. Both the sonic and visual portions of this work engage synthetic models of natural systems: within sound, the immersive textures created using the *Murmurator* and *Sound Particles* put the listener within virtual 3-dimensional environments, within visuals, simulations of landscape topologies and oceans, virtual sun- and moon-based dynamic lighting, as well as digital snow create a naturalistic (but still eerily synthetic) virtual world. More than being just an exploration of these digital natural system models, however, the poem at the core of *Vestigial Wings* gives this work a tangible story in addition to a (potentially moralistic) message. This causes this work to not only be multimedia from a purely 'medium-based' standpoint, but to also have multiple layers of interpretation and meaning as a function of the relationships between its poetic, sonic, and visual components.

3.4.3 Artificial Reef (work-in-progress)

Artificial Reef is a work-in-progress fixed media composition composed in higher-order ambisonics (full 3D 7th order HOA). This work directly explores the composing of indexicality of NSSMs over time using recordings of intertidal oyster reefs—those made within the context of the bioacoustics research described earlier in this text— as base musical material. Over the approximately 8 minute duration of this work, the recording of the reef starts completely abstracted (using *AcousTrans* to map its events onto a synthesized sound world), transitions to semi-representative (a mixture of the original and the *AcousTrans* mapping), and finally to a purely representative (playing back the reef recordings as is) presentation. During the second half of the work this process backtracks, but through a different type of abstraction: one that uses *HabiSpat* and the *Murmurator* to develop the recordings of the reef into spatially rich, densely immersive textures, which build and then fade to nothingness.

This direct engagement with level of abstraction as dictator of both macro level form and lower level development (especially through *AcousTrans*) resonates with two NSSM concepts. First, the narrative of this electroacoustic work is a direct byproduct of an exploration of the gestural (first half) and sonic (second half) abstractive potentials of the underlying real-world recording (as natural system). Second, this work suggests an electroacoustic composition practice that is deeply tied to recorded sound materials as models of natural systems (akin to soundscape composition). The gestures and internal movement inherent to the recordings curated by the composer are enhanced and abstracted through computational algorithms, but are ultimately retained, striking a balance between the hand of the composer (or algorithmicist) and the music implicit in the sound recording.

The engagement of this work with full 3D higher-order ambisonics also allows for it to be diffused over a wide variety of loudspeaker arrays. Over a 3-dimensional high-density loudspeaker array, for example, the listener is immersed from head to toe in the sound of the reef, positioning the listener, through sound, into an artificial, electroacoustically enhanced intertidal oyster reef.

3.5a Future Directions

There are many future directions for these projects, some of which are ongoing. These include extended engagement with the expressive, trans-modal feedback potentials of the *Murmurator* in live performance, as well as continued use of *AcousTrans* and *HabiSpat* within multi-channel acousmatic works, several of which are in the composition or pre-composition stages as of this writing (including *Artificial Reef*, outlined in the previous section). These softwares may also be expanded, refined, and possibly ported to different music software contexts (in particular highly-accessible, Web-based audio applications) to enhance their usability and applicability to different artistic and/or scientific contexts.

In addition, the author hopes to continue exploring more 'intelligent' methods for engaging natural system models in electroacoustic music contexts, such as the use of machine learning within the context of concatenate synthesis in *AcousTrans* and the implementation of competitive Lotka-Volterra equations in the context of *HabiSpat*. Further engagement with these kinds of complex algorithms within an artistic context point to continuing collaboration with scientists in a wide array of fields, including environmental sciences as well computer science, computational biology, neuroscience, and sociology, to name but a few. For example, using high-dimensional astronomical data to inform musical structure or using the features of particle simulations to drive sound synthesis. The potentials for applying natural systems data and simulations to electroacoustic sound-making are limitless.

Lastly, the author plans to release these softwares as open source projects for other electroacoustic composers, bioacoustics researcher, and sound artists to use and modify as they see fit. The communities around the software platform used (Max) and scientific contexts of some of these works (bioacoustics research, natural computing), along with institutions with access to high-density loudspeaker arrays (HDLAs) or just multi-channel loudspeaker setups, will be targeted.

3.5b Final Thoughts

These three projects explore many aspects of the modeling of natural systems in sound. They engage with both simulation and sampling of natural systems, sometimes using both in combination. They explore a gamut of modeling methodologies (some borrowed and adapted, others novel). Different mapping strategies are used, although each focuses on the potentials of natural system models to enact immersive electroacoustic sound environments. The interfaces and meta-interfaces to these mappings afford musical expressivity: both within the context of live electronics performance and as sound generation tools for multi-channel fixed media.

The latter two of these systems were spawned out of both an experience of natural systems and a scientific engagement with them. This speaks to the importance of NSSM development that engages the qualitative, personal perception of the natural system at play (if possible) in addition to a quantitative, scientific understanding of them. For the author, the contrast or reconciliation of these two views of the same object is a platform for sonic art-making: one that involves learning about natural systems, building digital tools to express them in sound, and creating musical works that make use of those computer-generated sounds.

These musical works are not merely inspired by natural systems nor are they sonic demonstrations of scientific results. Rather, the hope of the author is that the art and software produced as part of this dissertation suggest a different model of art-making, one that explores the space in between mysticism and science, inspiration and simulation, and actively celebrates the aesthetic potentials of models of our natural world.

4. Appendix

4.1 Project Documentation

All project documentation, including musical applications, software demonstrations, and code downloads, may be found at www.elistine.com/diss.

4.2 Glossary of Terms

0. Exposition

abstraction - a process of reordering, layering, time-stretching or compressing, processing with effects, or otherwise altering musical materials

adaptation - quality of an open system to endure changes in its environment

chaos - the quality of a system exhibiting extreme sensitivity to the initial state of its elements, resulting in seemingly erratic behavior

closed system - system with non-permeable boundaries

complex system - a system whose dynamics are not apparent from an analysis of its elements in isolation

comprehensibility - the perception of how much a musical material is a direct index to its system model handle origins

computational intractability - the inability of a computer to efficiently (within a reasonable amount of time and/or resources) compute the next state of a particular simulation

curation - the extraction of certain portions of musical materials deemed more interesting than others

dynamics - how a system changes over time as a function of its elemental relationships and, if open, its input and/or output

element - indivisible components of a system

elemental relationships - interactions between elements of a system

emergence - the quality of a system as a whole exhibiting properties which are meaningful only when attributed to its whole, not its parts

environment - the world outside of a system's boundaries

feedback loop - formed within an open system if the output of the system informs its input

handle - system model outputs that afford re-presentation of the behavior of the model

hard systems viewpoint - views systems through quantitative metrics, seeks to optimize

homeostasis/systemic equilibrium - quality of a system to self-regulate in order to maintain its behavior regardless of changes in its environment

indexicality - how much a musical material is a direct index to its system model handle origins

interface - a tool for dynamically altering a mapping

key - an understanding of the mapping(s) from system model to musical materials

mapping - the procedure by which an original entity is adapted to a map

meta-interface - an interface which alters the mapping of another interface

model - an imitation of some entity that allows for analysis, experimentation, or other procedures which, when applied to the original, would be expected to produce identical results

modeling decisions - context-dependent choices of how to construct a model that best represents desired properties of the original

natural system sound models - an electroacoustic composition framework that involves natural systems, models of those natural systems, mappings from those models to create musical materials, and composing with those musical materials to create a musical presentation

open system - system with free exchange of materials across its boundaries

sample-based model - imitation of an entity through records of that entity's behavior

scale - a function of at what point elements are considered indivisible and at what point the effects of the interactions of those elements are unregarded *simple system* - a system whose dynamics are apparent from an analysis of its elements in isolation

simulation-based model - imitation of an entity through simulation

soft systems viewpoint - views system qualitatively, not seeking to optimize or even clearly define

state - a description of a system at a discrete moment in time

system - a set of interconnected parts which function together as a complex whole

systemic caricature - intentionally un-optimized or distorted model of a system to highlight or critique a particular quality of the system

1. Historical Threads

algorithmic music - music that uses formalized procedures to generate or manipulate musical material

algorithmicist - one who choses algorithm procedures, how they are formalized, and which materials are generated or manipulated

artificial neural network - a set of artificial neurons connected together in a neural net, used to simulate learning on computers

audification - auditory display which uses a sample-based model in conjunction with a completely indexical mapping

auditory display - the transformation of data relations into perceived relations in an acoustic signal for the purposes of facilitating communication or interpretation

auditory icons - real-world sound recordings used to indicate events within a data stream in the context of auditory display

cellular automaton - an algorithm that models complex spatial interactions through an array of neighboring cells whose values are determined by a set of transition rules

code art - software art that takes as artistic material code, views code itself as aesthetic object

computer art - software art whose primary outcome is an aesthetic experience or product

data-driven music - music that uses sonification methodologies within an artistic context

earcons - synthesized sounds used to indicate events within a data stream in the context of auditory display

generative art - computer art that autonomously produces a set of materials

genetic/evolutionary algorithms - an algorithm that models natural selection through an iterated process of death, reproduction, and mutation using the genotypes of a population in conjunction with a fitness criteria

interactive art - computer art that affords dynamic alteration by humans

Lindenmayer systems (L-systems) - an iterative rewriting process used for simulating the fractal patterns of plants and trees

materialism (musical practice) - an engagement with the sonic affordances and agency of non-conventional instruments

musique concrète - type of acousmatic music which engages sound recordings directly as musical materials through reduced listening

natural computing - the process of extracting ideas from nature to develop computational systems, or using natural materials to perform computation

parallel perceptual model - a person-specific model of a natural system that acts in parallel with its quantitative model

physical modeling synthesis - synthesis method which uses computational models such as simulations of masses and springs to create expressive and nuanced emulations of the sounds of physical systems (often instruments)

software art - art that takes as central compositional material code, either to produce an aesthetic experience or product or that views code itself as aesthetic object

sonification - auditory display that involves higher levels of abstraction, deeper data-to-sound metaphors, and flexible mapping schemes, often for data exploration or entertainment and arts purposes

sound diffusion - a live electronics practice which involves manipulating interfaces in real-time to control the spatialization of (often) fixed media over loudspeaker arrays

spectralism - type of music which takes as musical material and organizing principle structures derived from sound recordings or acoustical phenomena, more generally

zoomusicological composition - type of music, often acoustic, which incorporates the sounds of animals and other natural phenomenon (via transcription or emulation)

2. Evaluation Framework

abstract (indexicality) - a musical presentation of a natural system model that foregrounds aesthetics and stylistically presents the model through various layers of abstraction and/or curation

acousmatic (sound production technology) - NSSMs which only use fixed, electronically-produced sound sources are present

acoustic (sound production technology) - NSSMs which only use acoustic sound sources

compositional space (creative space) - space within a NSSM where the creator makes decisions and selections on the modeling of the natural system and the medium and aesthetics of the sonic output

creative spaces - nested spaces within which NSSMs are designed, namely physical, compositional, and mapping

indexicality - how much a musical presentation is a direct index to its natural system model handle origins

level of abstraction - inversely related to indexicality, how abstract a musical presentation is with respect to its natural system model origins

live electronics (sound production technology) - NSSMs in which electronics and/or acoustic sound sources are performed in real-time

mapping space (creative space) - space within a NSSM (within compositional space) where handles from the natural system model and the means of producing sonic output are mapped to one another

modeling methodology - how the natural system is modeled, either sampled or simulated *original (modeling methodology)* - engagement with a natural system that extends its sonic affordances without quantitative modeling

physical space (creative space) - space within a NSSM where the physical natural system exists and the sonic output coexists

representative (indexicality) - a musical presentation which emphasizes the literal content of a natural system model

sampling (modeling methodology) - model of an entity via (possibly trans-modal) records of that entity

semi-representative (indexicality) - a musical presentation which uses curation techniques on indexical source material or presents a non-realistic systemic caricature as natural system model

simulation (modeling methodology) - model of an entity via (possibly trans-modal) a systemic emulation constructed from observed behaviors, measured elemental relationships, and/or known environmental catalysts

sound production technology - the forces that the NSSM uses to produce sound, either acoustic, live electronics, or acousmatic

4.3 References

Introduction

1. Roginska, Agnieszka, et al. "Immersive sonification for displaying brain scan data." International Conference on Health Informatics, HEALTHINF 2013. 2013.

Exposition

1. Hitchins, Derek K. Putting systems to work. Vol. 325. Chichester: Wiley, 1992: 3.
2. Smithson, Peter, Ken Addison, and Ken Atkinson. Fundamentals of the physical environment. Routledge, 2013: 9.
3. Daley, Ben and Rebecca Kent. Environmental Science and Management. SOAS, University of London. 2014. Accessible: https://www.soas.ac.uk/cedep-demos/000_P500_ESM_K3736-Demo/module/topindex.htm# Unit 1 Section 2.1
4. Shishko, Robert, and Robert Aster. "NASA systems engineering handbook." NASA Special Publication 6105 1995: 3.
5. Hitchins, Derek K. Putting systems to work. Vol. 325. Chichester: Wiley, 1992: 7.
6. Daley, Ben and Rebecca Kent. Environmental Science and Management. SOAS, University of London. 2014. Accessible: https://www.soas.ac.uk/cedep-demos/000_P500_ESM_K3736-Demo/module/topindex.htm# Unit 1 Section 2.1
7. Moere, Andrew Vande. "Beyond the tyranny of the pixel: Exploring the physicality of information visualization." Information Visualisation, 2008. IV'08. 12th International Conference. IEEE, 2008.
8. Jones, W. P., and Brian Edward Launder. "The prediction of laminarization with a two-equation model of turbulence." International journal of heat and mass transfer 15.2 1972: 301-314.
9. Fee, Michale S., et al. "The role of nonlinear dynamics of the syrinx in the vocalizations of a songbird." Nature 395.6697 1998: 67.
10. Trevisan, M. A., and G. B. Mindlin. "New perspectives on the physics of birdsong." Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences 367.1901 2009: 3239-3254.
11. Roads, Curtis. Microsound. MIT press, 2004: 3-4.
12. Emmerson, Simon. "Aural landscape: musical space." Organised Sound 3.2 (1998): 135-140.
13. Smalley, Denis. "Space-form and the acousmatic image." Organised sound 12.1 (2007): 35-58.

14. Barrett, Natasha. "Trends in electroacoustic music." The Cambridge companion to electronic music 232 (2007): 234.
15. Emmerson, Simon. "The relation of language to materials." The language of electroacoustic music. Palgrave Macmillan, London, 1986. 17-39.
16. Huron, D. “A six-component theory of auditory-evoked emotion.” In Proceedings of ICMPC7, Sydney, Australia, 2002. 673–676.
17. Norman, Katharine. "Conkers (listening out for organised experience)." Organised Sound 15.2 (2010): 117.
18. Barrett, Natasha. "Little animals: compositional structuring processes." Computer Music Journal 23.2 (1999): 11-18.
19. Andean, James. "Narrative Modes in Acousmatic Music." Organised Sound 21.3 (2016): 193-195.
20. The author has found that complexity may often be better engaged with or described using a qualitative viewpoint rather than a quantitative one.

Algorithmic Acousmatic Music

1. Allen, Aaron S and Kevin Dawe. Current Directions in Ecomusicology: Music, Culture, Nature. Routledge, 2015. Chapter 1.
2. Nierhaus, Gerhard. Algorithmic composition: paradigms of automated music generation. Springer Science & Business Media, 2009: 21-27.
3. Essl, Karlheinz. "Algorithmic composition." Cambridge Companion to Electronic Music. Cambridge University Press, Cambridge (2007):109.
4. Ariza, Christopher. "Two pioneering projects from the early history of computer-aided algorithmic composition." Computer Music Journal 35.3 (2011): 40.
5. de Castro, Leandro Nunes. "Fundamentals of natural computing: an overview." Physics of Life Reviews 4.1 2007: 3.
6. Jain, Anil K., Jianchang Mao, and K. M. Mohiuddin. "Artificial neural networks: A tutorial." Computer 3 (1996): 33.
7. Adamatzky, Andrew, and Theresa Schubert. "Slime mold microfluidic logical gates." Materials Today 17.2 2014: 86-91.
8. Wang, Ge. "A history of programming and music." The Cambridge Companion to Electronic Music (2007): 58-63.
9. Doornbusch, Paul. "Early Hardware and Early Ideas in Computer Music: Their Development and Their Current Forms." The Oxford Handbook of Computer Music. 2012: 54.

10. Smith, Julius O. "Physical modeling synthesis update." Computer Music Journal 20.2 (1996): 45.
11. Wang, Jui-Hsien, et al. "Toward wave-based sound synthesis for computer animation." ACM Transactions on Graphics TOG 37.4 2018: 109.
12. Prusinkiewicz, Przemyslaw, and Aristid Lindenmayer. The algorithmic beauty of plants. Springer Science & Business Media, 2012.
13. Hingston, Philip F., Luigi C. Barone, and Zbigniew Michalewicz, eds. Design by evolution: advances in evolutionary design. Springer Science & Business Media, 2008: 172.
14. Nierhaus, Gerhard. Algorithmic composition: paradigms of automated music generation. Springer Science & Business Media, 2009: 149-153.
15. Manousakis, Stelios. "Do Digital Monkeys Inhabit Virtual Trees?" Modular Brains, 2007, modularbrains.net/portfolio/do-digital-monkeys-inhabit-virtual-trees/.
16. Reynolds, Craig W. "Flocks, herds and schools: A distributed behavioral model." ACM SIGGRAPH computer graphics. Vol. 21. No. 4. ACM, 1987.
17. Bates, Enda, and Dermot Furlong. "Score File Generators for Boids-Based Granular Synthesis in Csound."Audio Engineering Society Convention 126. Audio Engineering Society, 2009.
18. Schacher, J. C., D. Bisig, and M. Neukom. "Composing with swarm algorithms—creating interactive audio-visual pieces using flocking behaviour." Proceedings of the International Computer Music Conference. Vol. 2011. 2011: 106.
19. Wolfram, Stephen. A new kind of science. Vol. 5. Champaign, IL: Wolfram media, 2002: 231-233.
20. Conway, John. "The game of life." Scientific American 223.4 1970: 4.
21. Miranda, Eduardo Reck. "Granular synthesis of sounds by means of a cellular automaton." Leonardo 28.4 1995: 297-300.
22. Burraston, Dave, et al. "Cellular automata in midi based computer music." Proceedings of the 2004 International Computer Music Conference. International Computer Music Association, 2004: 5.
23. Burraston, Dave, and Ernest Edmonds. "Cellular automata in generative electronic music and sonic art: a historical and technical review." Digital Creativity 16.3 (2005): 172.
24. Miranda, Eduardo Reck, and Marcelo M. Wanderley. New digital musical instruments: control and interaction beyond the keyboard. Vol. 21. AR Editions, Inc., 2006: 241-243.

25. Magnus, Cristyn. "Evolutionary musique concrete." Workshops on Applications of Evolutionary Computation. Springer, Berlin, Heidelberg, 2006: 689.
26. Lidström, Magnus. “Synplant.” Sonic Charge, 2009, soniccharge.com/synplant.
27. Eldridge, Alice, and Oliver Bown. "Biologically inspired and agent-based algorithms for music." The Oxford Handbook of Algorithmic Music (2018): 209.
28. Steels, Luc. "The artificial life roots of artificial intelligence." Artificial life 1.1_2 (1993): 25.
29. Lee, Michael, Adrian Freed, and David Wessel. "Neural networks for simultaneous classification and parameter estimation in musical instrument control." Adaptive and Learning Systems. Vol. 1706. International Society for Optics and Photonics, 1992.
30. Fiebrink, Rebecca, and Perry R. Cook. "The Wekinator: a system for real-time, interactive machine learning in music." Proceedings of The Eleventh International Society for Music Information Retrieval Conference (ISMIR 2010)(Utrecht). 2010.
31. Schneider, David. "Build a neural synthesizer." IEEE Spectrum 55.7 (2018): 16.
32. Oord, Aaron van den, et al. "Wavenet: A generative model for raw audio." arXiv preprint arXiv:1609.03499 (2016).
33. Snyder, Jeff, Aatish Bhatia, and Mike Mulshine. "Neuron-modeled Audio Synthesis." In Proceedings of New Interfaces for Musical Expressions (NIME). 2018.
34. Solis, Gabriel, and Bruno Nettl, eds. Musical improvisation: Art, education, and society. University of Illinois Press, 2009: 43.
35. Fineberg, Joshua. *Spectral music: history and techniques*. Vol. 19. Routledge, 2000: 116.
36. Barrett, Natasha. Three Fictions, 2000, www.natashabarrett.org/fictions.html.
37. Barrett, Natasha. Involuntary Expression, 2017, http://www.natashabarrett.org/Involuntary_Expression.html.
38. Truax, Barry. "Soundscape composition as global music: electroacoustic music as soundscape." Organised Sound 13.2 (2008): 104.
39. Westerkamp, Hildegard. "Linking soundscape composition and acoustic ecology." Organised Sound 7.1 2002: 51-56.
40. Gilmurray, Jonathan. "Ecological Sound Art: Steps towards a new field." Organised Sound 22.1 2017: 34.

Sonification + Data-Driven Music

1. Kramer, Gregory, et al. "Sonification report: Status of the field and research agenda." 1997: 4.
2. Bly, Sara Ann. “Sound and Computer Information Presentation.” University of California, Davis, 1982.
3. Dean, Roger T., ed. The Oxford handbook of computer music. OUP USA, 2009: 321.
4. Ballora, Mark. Data analysis through auditory display: applications in heart rate variability. Diss. McGill University, 2000: 29.
5. Dombois, Florian. "Using audification in planetary seismology." Georgia Institute of Technology, 2001: 227.
6. Edwards, Alistair, et al. "Sonification strategies for examination of biological cells." Georgia Institute of Technology, 2010.
7. Vickers, Paul, and Bennett Hogg. "Sonification Abstraite/Sonification Concr\ete: An'Aesthetic Perspective Space'for Classifying Auditory Displays in the Ars Musica Domain." arXiv preprint arXiv:1311.5426 (2013): 215
8. Sturm, Bob L. "Surf music: Sonification of ocean buoy spectral data." Georgia Institute of Technology, 2002.
9. Schmele, Timothy, and Imanol Gomez. "Exploring 3D audio for brain sonification." Georgia Institute of Technology, 2012.
10. Taylor, Stephen. “From Program Music to Sonification: Representation and the Evolution of Music and Language.” ICAD 2017.
11. Adams, John Luther. The place where you go to listen: In search of an ecology of music. Wesleyan University Press, 2010.
12. Personal communication with Jem Altieri, sound designer of *The Place*, October 25th, 2018.
13. Barrett, Natasha, and Karen Mair. "Sonification for geoscience: Listening to faults from the inside." EGU General Assembly Conference Abstracts. Vol. 16. 2014.
14. Barrett, Natasha, and Karen Mair. "Aftershock: A science–art collaboration through sonification." Organised Sound 19.1 (2014): 4.
15. Barrett, Natasha. "A compositional methodology based on data extracted from natural phenomena." In Proceedings of International Computer Music Conference (ICMC). 2000.
16. Scaletti, Carla. “h->gg.” Mu Psi, 2017, carlascaletti.com/sounds/sound-art/h-gg/.

Live Electronic Music

1. Barrett, Natasha. "Trends in electroacoustic music." The Cambridge companion to electronic music 232 (2007): 241.
2. Collins, Nick, and Julio d'Escriván, eds. The Cambridge companion to electronic music. Cambridge University Press, 2017: 39
3. Hodgson, Jay. "The Time-Lag Accumulator As A Technical Basis For Brian Eno's Early Large-Scale Ambient Repertoire, 1973-75." The Art of Record Production, 2008. http://www.artofrecordproduction.com/content/view/167/109/.
4. Maconie, Robin. "Stockhausen's Mikrophonie I: Perception in Action." Perspectives of New Music (1972): 92-101.
5. Licht, Alan. "Sound Art: Origins, development and ambiguities." Organised Sound 14.1 (2009): 5.
6. Miranda, Eduardo Reck, and Marcelo M. Wanderley. New digital musical instruments: control and interaction beyond the keyboard. Vol. 21. AR Editions, Inc., 2006: 224.
7. Miranda, Eduardo Reck, and Marcelo M. Wanderley. New digital musical instruments: control and interaction beyond the keyboard. Vol. 21. AR Editions, Inc., 2006: 224.
8. Borgo, David. "Sync or Swarm: Musical improvisation and the complex dynamics of group creativity." Algebra, meaning, and computation. Springer, Berlin, Heidelberg, 2006: 2.
9. Landgraf, Edgar. "Improvisation: form and event–a spencer-Brownian calculation." Emergence and embodiment: New essays on second-order systems theory (2009): 179-204.
10. Lewis, George E. "Too many notes: Computers, complexity and culture in voyager." Leonardo Music Journal (2000): 35.
11. Blackwell, Tim M., and Peter Bentley. "Improvised music with swarms." Proceedings of the 2002 Congress on Evolutionary Computation. CEC'02 Cat. No. 02TH8600. Vol. 2. IEEE, 2002.
12. Biles, John A. "GenJam: A genetic algorithm for generating jazz solos." ICMC. Vol. 94. 1994.
13. Beverly, Thomas Rex. "Telepresent Storm: Rita for IPads, Historical Weather Data, and Projected Real-Time Graphical Score (2013) - 7'." Thomas Rex Beverly Audio, 2013, thomasrexbeverly.com/products/telepresent-storm-rita-for-ipads-historical-weather-data-and-projected-real-time-graphical-score.

14. Di Scipio, Agostino. "'Sound is the interface': from interactive to ecosystemic signal processing." Organised Sound 8.3 2003: 269-277.
15. Tanaka, A. and Knapp, R.B. "Multimodal Interaction in Music Using the Electromyogram and Relative Position Sensing." In Proceedings of International Conference on New Interfaces for Musical Expression NIME02, May 24-26, Dublin, University of Limerick, 2002: 43-48.
16. Whitelaw, Mitchell. *Metacreation: art and artificial life*. Mit Press, 2004: 17.
17. Kurutz, Steven. A Haven for Making Music With Plants. The New York Times, 19 Oct. 2018, www.nytimes.com/2013/08/29/garden/a-haven-for-making-music-with-plants.html.
18. Patitucci, Joe, and Jon Shapiro. MIDI Sprout. 2016, www.midisprout.com/.
19. Timms, Peter. Going Off, Virally [online]. Island, No. 125, Winter 2011: 89-96.
20. Miranda, Eduardo R., Andrew Adamatzky, and Jeff Jones. "Sounds synthesis with slime mould of Physarum polycephalum." Journal of Bionic Engineering 8.2 2011: 107-113
21. Brackett, John. John Zorn: Tradition and Transgression. Indiana University Press, 2008.
22. Cardew, Cornelius. "A Scratch Orchestra: draft constitution." The Musical Times 110.1516 1969: 617-619.
23. Sergeant, Mathew. "Towards a Nonhuman Affordance through Performative Contemporary Music." In Proceedings of 2018 Material Cultures of Music Notation Conference. Forthcoming.
24. Akama, Ryoko. "Failed Experiments." Ryoko Akama Site, 2017, ryokoakama.com/archives/works/failed-experiments.

Software Art

1. Cramer, Florian, and Ulrike Gabriel. "Software Art." American Book Review, issue "Codeworks"(Alan Sondheim, ed.) (2001): 2.
2. Casey Reas, Engeli, Maia. "Ars Electronica 2003: CODE--The Language of Our Time." Leonardo 37.3 2004: 174.
3. Nash, Chris, and Alan F. Blackwell. "Tracking virtuosity and flow in computer music." In Proceedings of International Computer Music Conference (ICMC). 2011: 577
4. Malina, Roger F., and Sean Cubitt. Software studies: A lexicon. Mit Press, 2008: 230.

5. Cramer, Florian, and Ulrike Gabriel. "Software Art." American Book Review, issue “Codeworks”(Alan Sondheim, ed.) (2001): 3.
6. Taylor, Grant D. When the machine made art: the troubled history of computer art. Bloomsbury Publishing USA, 2014: 27.
7. Rockman, Arnold, and Leslie Mezei. "The electronic computer as an artist." Canadian Art 21.6 (1964): 365.
8. Galanter, Philip. "What is generative art? Complexity theory as a context for art theory." In GA2003–6th Generative Art Conference. 2003: 3.
9. Whitelaw, Mitchell. "System stories and model worlds: A critical approach to generative art." Readme 100 2005: 135-154.
10. Fuller, Matthew. Behind the blip: Essays on the culture of software. Autonomedia, 2003: 22.
11. Rabelais, Akira. Argeïphontes Lyre. 2014, akirarabelais.com/o/software/al.html.
12. Fuller, Matthew. Behind the blip: Essays on the culture of software. Autonomedia, 2003: 29.
13. McCormack, Jon, et al. "Ten questions concerning generative computer art." Leonardo 47.2 (2014): 135-136.
14. Watkins, Holly. "Toward a post-humanist organicism." Nineteenth-Century Music Review 14.1 (2017): 103.
15. Reas, Casey. “REAS.com Is a Database for Casey REAS.” REAS.com Is a Database for Casey REAS, 2019, reas.com/.
16. Reas, Casey. "PROCESS COMPENDIUM 2004-2010." Paradigms in Computing: Making, Machines, and Models for Design Agency in Architecture 2015: 44.
17. Tarbell, Jared. “Substrate.” Complexification l Gallery of Computation, 2003, www.complexification.net/gallery/machines/substrate/.
18. Tarbell, Jared. “Happy Place.” Complexification l Gallery of Computation, 2004, www.complexification.net/gallery/machines/happyPlace/index.php.
19. Tarbell, Jared. “Sand Stroke.” Complexification l Gallery of Computation, 2004, http://www.complexification.net/gallery/machines/sandstroke/.
20. Driessens, Erwin, and Maria Verstappen. "Natural processes and artificial procedures." Design by Evolution. Springer, Berlin, Heidelberg, 2008. 101.
21. Driessens, Erwin, and Maria Verstappen. “Accretor.” Driessens & Verstappen, 2007, notnot.home.xs4all.nl/accretor/Accretor.html.
22. Driessens, Erwin, and Maria Verstappen. “E-volved Cultures 360°.” Driessens & Verstappen, 2016, https://notnot.home.xs4all.nl/evolvedimages/evolvedcultures/

evolved-360.html

Evaluation Framework

1. Richards, Peter, and George Gonzalez. “The Wave Organ.” Exploratorium: the Museum of Science, Art and Human Perception, 14 Feb. 2018, www.exploratorium.edu/visit/wave-organ.
2. Doolittle, Emily. "Crickets in the concert hall: A history of animals in western music." TRANS. Revista Transcultural de Música 12 (2008): 11.
3. Wulfson, Harris, G. Douglas Barrett, and Michael Winter. "Automatic notation generators." Proceedings of the 7th international conference on New interfaces for musical expression (NIME). ACM, 2007: 346.
4. Riveire, Janine. "Using improvisation as a teaching strategy." Music Educators Journal 92.3 2006: 43.
5. Park, Joo Won. “Toccata.” Joo Won Park, 15 Dec. 2018, joowonpark.net/works/.

The Murmurator

1. Zvonar, Richard. "A history of spatial music." Montreal: CEC 1999.
2. Marshall, Mark T., Joseph Malloch, and Marcelo M. Wanderley. "Gesture control of sound spatialization for live musical performance." International Gesture Workshop. Springer, Berlin, Heidelberg, 2007.
3. Garrett, Richard. Audio Spray Gun 0.8–the Generation of Large Sound-Groups and Their Use in Three-Dimensional Spatialisation. Ann Arbor, MI: Michigan Publishing, University of Michigan Library, 2015.
4. Normandeau, Robert. "Octogris2 et ZirkOSC2: outils logiciels pour une spatialisation sonore intégrée au travail de composition." Proc. of Journées d’Informatique Musicale JIM, Montréal, Canada 2015.
5. Wilson, S., J. Harrison, and S. L. Ancona. "Beastmulch.”, 2008, https://www.birmingham.ac.uk/facilities/ea-studios/research/mulch.aspx
6. Schacher, Jan C. "Seven years of ICST Ambisonics tools for maxmsp–a brief report." Proc. of the 2nd International Symposium on Ambisonics and Spherical Acoustics. 2010.
7. Carpentier, Thibaut, Markus Noisternig, and Olivier Warusfel. "Twenty years of Ircam Spat: looking back, looking forward." 41st International Computer Music Conference ICMC. 2015.
8. Bukvic, Ivica Ico. "D4: an Interactive 3D Audio Rapid Prototyping and

Transportable Rendering Environment Using High Density Loudspeaker Arrays." 2016.
9. Barrett, Natasha. "Interactive spatial sonification of multidimensional data for composition and auditory display." Computer Music Journal 40.2 2016: 47-69.
10. Reynolds, Craig W. "Flocks, herds and schools: A distributed behavioral model." ACM SIGGRAPH computer graphics. Vol. 21. No. 4. ACM, 1987.
11. Stine, Eli and Davis, Kevin. “The Murmurator: A Flocking Simulation-Driven Multi-Channel Software Instrument for Collaborative Improvisation.” In International Computer Music Conference ICMC. 81-86.

AcousTrans

1. Coen, Loren D., et al. "Ecosystem services related to oyster restoration." Marine Ecology Progress Series 341 (2007): 303-307.
2. Kemp, W. Michael, et al. "Eutrophication of Chesapeake Bay: historical trends and ecological interactions." Marine Ecology Progress Series 303 (2005): 1-29
3. Di Iorio, Lucia, et al. "Hydrophone detects cracking sounds: Non-intrusive monitoring of bivalve movement." Journal of Experimental Marine Biology and Ecology 432 (2012): 9-16.
4. Lillis, Ashlee, David B. Eggleston, and DelWayne R. Bohnenstiehl. "Estuarine soundscapes: distinct acoustic characteristics of oyster reefs compared to soft-bottom habitats." Marine Ecology Progress Series 505 (2014): 1-17.
5. Sueur, Jérôme, et al. "Acoustic indices for biodiversity assessment and landscape investigation." *Acta Acustica united with Acustica* 100.4 (2014): 772-781.
6. Krause, Bernard L. "The niche hypothesis: a virtual symphony of animal sounds, the origins of musical expression and the health of habitats." *The Soundscape Newsletter* 6 (1993): 6-10.
7. Katz, Jonathan, Sasha D. Hafner, and Therese Donovan. "Tools for automated acoustic monitoring within the R package monitoR." Bioacoustics 25.2 (2016): 197-210.
8. Stine, Eli. “Creating Immersive Electronic Music from the Sonic Activity of Environmental Soundscapes.” In Joint Proceedings of the ACM IUI 2019 Workshops: forthcoming.
9. Castro, Leandro Nunes. 2006. Fundamentals of natural computing: basic concepts, algorithms, and applications. Chapman and Hall/CRC, Chapter 1.
10. Stine, Eli and Davis, Kevin. “The Murmurator: A Flocking Simulation-Driven Multi-

Channel Software Instrument for Collaborative Improvisation.” In 2018 International Computer Music Conference ICMC. 81-86.

11. Wilson, Catherine M., and Suresh K. Lodha. "Listen: A data sonification toolkit." Georgia Institute of Technology, 1996.
12. Schwarz, Diemo. Corpus-based concatenative synthesis. IEEE signal processing magazine 24, 2 2007, 92–104.
13. Schwarz, Diemo. 2006. Concatenative sound synthesis: The early years. Journal of New Music Research 35, 1 2006, 3–22.
14. Schwarz, Diemo, Grégory Beller, Bruno Verbrugghe, and Sam Britton. Real-time corpus-based concatenative synthesis with catart. In 9th International Conference on Digital Audio Effects DAFx. 2006. 279–282.
15. Brent, William. A timbre analysis and classi cation toolkit for pure data. In International Computer Music Conference (ICMC). 2010. 224–229.
16. Casey, William. Soundspotting: A new kind of process? In The Oxford Handbook of Computer Music. 2012.
17. Driedger, Jonathan, Thomas Prätzlich, and Meinard Müller. 2015. Let it Bee- Towards NMF-Inspired Audio Mosaicing. In International Symposium on Music Information Retrieval ISMIR. 350–356.
18. Aucouturier, Jean-Julien, and François Pachet. "Jamming with plunderphonics: Interactive concatenative synthesis of music." Journal of New Music Research 35.1 (2006): 35-50.
19. Einbond, Aaron, et al. "Fine-tuned control of concatenative synthesis with catart using the bach library for Max." International Computer Music Conference (ICMC). 2014: 1037–1042.
20. Schacher, Jan C. Seven years of ICST Ambisonics tools for maxmsp–a brief report. In Proc. of the 2nd International Symposium on Ambisonics and Spherical Acoustics. 2010.
21. Schnell, Norbert, et al. "MuBu and friends–assembling tools for content based real-time interactive audio processing in Max/MSP." In International Computer Music Conference ICMC. 423–426.
22. Schnell, Norbert, et al. "PiPo, A Plugin Interface for Afferent Data Stream Processing Modules." In International Symposium on Music Information Retrieval (ISMIR). 2017.
23. Schwarz, Diemo, Norbert Schnell, and Sebastien Gulluni. Scalability in content-based navigation of sound databases. In International Computer Music Conference

ICMC. 2009. 13–16.
24. Krause, Bernard L. "The niche hypothesis: a virtual symphony of animal sounds, the origins of musical expression and the health of habitats." *The Soundscape Newsletter* 6 (1993): 6-10.

HabiSpat

1. Ruthenberg, Jessica, and Ruth Boettcher. “Virginia’s Barrier Islands: A Vital Refuge for Shorebirds and Seabirds.” Virginia Department of Game & Inland Fisheries, 2018, www.dgif.virginia.gov/blog/virginias-barrier-islands-a-vital-refuge-for-shorebirds-and-seabirds/.
2. Beck R. 1996. Nesting seabird census of Hog Island and Cobb Island of the Virginia Coast Reserve 1990. Environmental Data Initiative. Dataset accessed 11/05/2018.
3. Planqué, Bob, and Willem-Pier Vellinga. Xeno-Canto: Sharing Bird Sounds from around the World, www.xeno-canto.org/.
4. Malcai, Ofer, et al. "Theoretical analysis and simulations of the generalized Lotka-Volterra model." Physical Review E 66.3 (2002): 031102.
5. Fiore, Vincenzo. Lotka–Volterra Equations: Basic Dynamics and Input-Controlled Parameter Modulation. 12 Dec. 2017, www.vincenzofiore.it/algorithmic-models/lotka-volterra-equations-basic-dynamics-and-input-controlled-parameter-modulation/.