Building Birds

Kyle McCrosson

Abstract

Physical modeling synthesis is the process of using wave synthesis to emulate a naturally occurring, real-world sound. This process often involves mathematical formulas and studies of biological structures to inform the design of the synthesis process in order to most accurately achieve the desired timbres. The goal of physical modeling is usually to create a fabricated product that is indistinguishable from the target sound in the natural world. In this paper, I explore the historical relevance of birdcall in composition, evaluate different methods taken to model birdcall accurately, and analyze the approach that I have chosen to take using Max/MSP, a visual programming language and development environment that is specialized to the needs of music and multimedia projects.

Introduction

Birdcall is a source of musical inspiration for composers that is so engrained in human history, many cultures hold beliefs that humans learned their capacity for music through the mimicry of birdcall. As a result of this fascination, capturing and expressing birdsong as a part of, or along with musical accompaniment, is a pervasive theme of music composition (Head, 1997). Perhaps birdcall captivates human interest because of its similarity in logic, and use of pitch, to music that humans create with their own voices and instruments. Once audio recording became a possibility, birds have been incorporated into everything from orchestral works to soundscape compositions.

As a composer and music technologist, I am interested in the way that birdcall can be incorporated into electronic music, and how this allows composers to explore the intersection of nature and technology, a theme that is increasingly prevalent in our modern world. This interest propels my exploration of synthesis methods that are capable of creating a wide array of natural and synthetic-sounding birdcall, with sound design, compositional, and musical motivations.

Although a complicated process, in this paper we review a few approaches that previous researchers have used to approach birdcall synthesis, followed by my own synthesis model.

Composition with Birdcall

Composers can have a variety of relationships with birdcall, different goals and methods of its incorporation, and varying desires to accurately represent birdcall. Before sound recording techniques made it possible to playback real birdsong within a performance group, composers took melodic inspiration from it. For example, Handel's "Sweet Bird" aria from *L'Allegro, il Penseroso ed il Moderato*, composed in 1740, has the flute playing the part of a nightingale (Judd, 2019). The part is meant to be reminiscent of a bird's song, but not to actually be a realistic depiction of what the nightingale sounds like. Handel took inspiration from the nightingale's melody to imbue the tone of nature into this piece.

With the advent of audio recording technology, it became possible for composers to use actual recordings of birds in their orchestral, electroacoustic, and soundscape pieces. Ottorino Respighi's *Pines of Rome* is a symphonic poem concerning the pine trees "that so characteristically dominate the Roman landscape" (Ferguson, 1954). The third movement, "The Pines of the Janiculum," uses a gramophone recording of a nightingale at the end. This the first example of incorporating pre-recorded material to be played alongside a group of performers, and in his score Respighi notated exactly what disc recording of the nightingale he wanted to be used (Symes, 2004). As opposed to previous examples, where either a composer used birdsong simply as melodic inspiration, or used instrumentation to attempt to recreate the call, Respighi used actual birdcall to represent birdcall. Now, a composer could not only imbue the atmosphere or the idea of birds into their pieces, but they could literally incorporate birdcall right into it.

Although it was composed after the conception of recording technology, Olivier Messiaen chose to take an accurate approach to the musical representation of birdsong—without recording it—in one of his most influential series of pieces. Messiaen's *Catalogue d'oiseaux* is a series of piano solos composed between 1956 and 1958, each modeled after a bird that is native to a different region of France (Chiat, 2005). Trilling dissonances and high flutters capture the essence of these birds along with their environments. Messiaen approached this composition rather scientifically, attempting to actually create somewhat accurate portrayals of the birds. His process for the composition involved "actually [notating] the transcriptions in a specific time and place, presenting it as if it were a precise documentary work" (Chiat, 2005). In other words, he transcribed the piano parts that would represent each bird as he was physically listening to that bird in its natural environment. On defending his use of the piano as the instrument to portray this birdsong, Messiaen noted:

> The rendering of the timbre of the birdsong is particularly difficult. As the timbre is determined by a greater or lesser number of harmonics, I had to look for compositions of unexpected, reinvented sounds for almost every note of each birdsong … The piano alone is able to render the raucous or grating percussions of the raven and the great reed warbler, the scraping of the corncrake, the howl of the water rail, the barking of the herring gull, the dry and imperious struck-tone-like timbre of the black-eared wheatear, the sunny charm of the rock thrush or the black wheatear (Chiat, 2005).

Clearly, Messiaen took painstaking effort to map out the way in which he constructed his birdcall. His methodical construction of timbre with the individual piano notes can almost be seen as an early attempt of additive synthesis, an electronic composition technique through which complex sounds and spectra are built through the layering of simpler waveforms.

This concept of recorded birdcall became even more relevant with the genesis of soundscape composition led by R. Murray Schafer—a Canadian writer, composer, music educator, and environmentalist—and his organization, the World Soundscape Project (WSP). Beginning in the 1970's, the WSP sought to create compositions solely using recorded material from a geographic location, usually a city, in order to bring attention to issues surrounding noise pollution and the clash of city environments with the natural world. They wanted to "document acoustic environments, both functional and dysfunctional, and to increase public awareness of the importance of the soundscape, particularly through individual listening sensitivity" (Truax, 2008). Soundscape projects from the WSP ranged from accurate representations of the sounds of a particular location, to intentionally structured pieces that used the recorded sounds as compositional material. Hildegard Westerkamp, another composer from the WSP, created a variety of pieces that would fall into the latter category. In her series of pieces *Beneath the Forest Floor*, she uses birdcall, running water, wind, and other natural sounds to create an electronic composition of the natural world. Westerkamp describes the construction of the composition as follows:

> Most of the sounds for this composition were recorded in one specific location, the Carmanah Valley on Vancouver Island. This old-growth rainforest contains some of the tallest known Sitka spruce in the world and cedar trees that are well over one thousand years old. Its' stillness is enormous, punctuated only occasionally by the sounds of small songbirds, ravens and jays, squirrels, flies and mosquitoes. Although the Carmanah Creek is a constant acoustic presence it never disturbs the peace. Its' sound moves in and out of the forest silence as the trail meanders in and out of clearings near the creek (Westerkamp, 1992).

Westerkamp uses this "punctuation" of birdsong as melodic and structural material throughout the piece.

With the increased commonality of recorded material and fixed media pieces, being able to synthesize birdcall opens more possibilities for performers to incorporate environment and ambience into their pieces. This further incorporation of environment, ambience, and birdcall makes it useful to have more control over the exact sounds that composers use, allowing them to exert their own compositional intention rather than being subject to what is already present in a recording. A solution for this desire for control comes with sound synthesis, and the ability to synthetically produce a seemingly biological sound from oscillators and filters. The composer could directly control the pitch, tone, and movement of the birdcall that is used in their piece.

Björk's album *Utopia,* produced by Arca (an experimental electronic musician), is a source of inspiration for me because of its incorporation of recorded and synthesized birdcall; some of the synthesized material sounds realistic, and others sound more abstract and fantastical. Together, Björk and Arca created "a collection of birdsongs culled both from Iceland and Arca's homeland of Venezuela" (Larson, 2017) to craft the sonic environment of the album. The birds that they chose to highlight were unusual and uncommon, like the Montezuma Oropendola that "sounds like a Moog synth in a microwave" (Larson, 2017). The intent of this compositional approach was to blend these specifically chosen birdcalls that sounded electronic with synthesized material, in such a way as to create a cohesive natural environment. In addition, Björk had a higher order message to share with the incorporation of both real and synthetic birdcall throughout the album, a "proposal of our possible future when we learn to get technology and nature to collaborate" (Yalcinkaya, 2018). In *Utopia*, the birdcalls are not only used as ambience, but as compositional material. This shows the listener that birdcalls, whether realistic or distorted, can provide a meaningful musical experience; thereby bringing the relevance of sound design beyond classical and electroacoustic music, and into experimental forms of pop music. This artistic fusion of the natural and the synthetic is what propelled me to create a method of birdcall synthesis that has a range of adaptability to be realistic, synthetic, and musical, with flexibility and ease of use.

Bird Anatomy

In order to understand how to synthesize birdcall in an accurate and effective manner, it is important to look at the way that birds physically produce sound. In his book *Designing Sound*, computer scientist and sound designer Andy Farnell discusses how bird vocalization is unique when compared to humans or other mammals. The source of bird vocalization is rooted in the syrinx, an organ situated "at the base [of the throat] where the bronchi split off" (Farnell, 2010). Although similar in its sound-producing functionality, the syrinx is different from a human larynx. The human larynx is located further up in the throat than the syrinx, and only has one air input as opposed to the syrinx's two. Birds have the ability to control each lung and each half of the syrinx individually, allowing them to create continuous song for an extended period of time, whereas mammals must pause to breathe when they have depleted the air in their lungs. This ability can also give rise to a phenomenon called *biphonation* where a bird can sing two notes at the same time (Cornell Lab of Ornithology, 2014). The last distinguishing feature of birdcall physiology is that birds "resonate their whole breast and throat area" (Farnell, 2010) when they vocalize, rather than just resonating through their beaks. Most mammals only resonate sound through their throat and mouth when they vocalize. This means that birds are able to create much more powerful vocalizations proportionally to their size. Although biologists are still not certain of the exact way in which birds use the physical properties of the syrinx to create their song, they

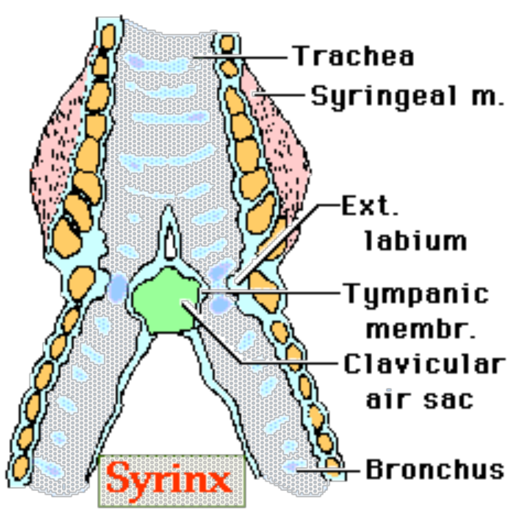

Figure 1 (Ritchison, n.d.)

know that it is “a combination of resonances from the membranes and interactions of the labia, cartilage, and trachea entrance.” All of these factors are important to take into account when discussing the modeling of birdcall, and the programming of a synthesizer capable of replicating birdcall. The mechanics of the syrinx and trachea lead Farnell to believe that a combination of frequency modulation and amplitude modulation of a carrier signal would be an effective method of synthesizing birdcall.

The anatomy of an avian vocal tract is not particularly clear-cut, as its configuration depends on the species of bird in question (Kahrs & Avanzini, 2001). The syrinx is generally considered to be “the source of oscillation of the vocal tract,” but while oscine species (songbirds) have two syrinx entrances, psittacines (birds in the parrot family) have only one. Birds’ beaks “act in a similar matter to horns: they exhibit a cutoff frequency that acts as a high pass filter.” In short, the most accurate way to fully and completely create a physical model of a specific birdcall would require mathematical formulas to compute the effects of each part of the vocal tract (syrinx, trachea, mouth, and beak) on the overall timbre of the call. These models would also have to be specific to the type of bird you were attempting to model. However, even with all that information, the best results can only be an approximation because there is no conclusive scientific evidence yet on how the syrinx physically functions to produce sound. Since I am looking to design a product with maximum ease of use, as well as the ability to be used for realistic, musical, and instrumental functions, the use of species-specific mathematical models and wavetables are most likely not going to be the best approach.

Bird vocalizations can be grouped into two main categories: songs and calls (Ritchison, n.d.). A birdsong is a longer, more complex, and more melodic vocalization that is usually employed by a bird for reproductive purposes like calling a mate or protecting territory. A birdcall is a shorter and more acoustically simple vocalization used for basic avian communication. For the purposes of this paper I am concerned with the synthesis of both birdcalls and birdsongs, as being able to produce both with a level of control is critical to synthesizing effective bird vocalizations. Birdsong could be composed for, much like a traditional instrument, because of its pitched nature; birdcall is much more useful as percussive sound, or for the creation of texture and environment. As I am seeking maximum control of synthesized bird vocalizations, my goal is to design a program that emulates both birdcall and birdsong.

The Structure of Birdcall in Regard to Synthesis

Much of human voice synthesis uses sine waves and noise, followed by series of filters in order to shape vowel and consonant sounds (Bonada et al., 2016). Markov chains, which are a computational statistical method of computing an answer based on a series of past inputs, are a good way to generate human speech because of the linearly dependent way that words and sentences are structured—they use all past inputs of a series to guess what is most likely to come next, much like the way you can guess what word a person might say next based on what they have already said. When a composer writes a melody, each consecutive note relates to all of the notes that precede it, and when a human speaks a sentence, each word is dependent on all the preceding words in order for any meaning to be constructed. In this way, Markov chains can be programmed to understand the basic general rules of melody-crafting or sentence-crafting, and can then use complex statistics to attempt to create their own logical sentences and melodies. Jordi Bonada, a computer music researcher in the Music Technology Group, Universitat Pompeu Fabra, Barcelona, proposes Markov chains would be similarly useful in the creation of bird song, with the knowledge that bird song is syntactically much simpler than human speech. This is an important consideration because most birdcalls are not simply one impulse, but a chain of gestures strung together in a way that, while it does not have the direct translatable significance of a human language, bears some sort of meaningful message for the bird. It would be a good idea for any musician, composer, or sound designer to take into account the types of gestures that are present in real-life birdcalls when attempting to either create convincing birdcall, or to subvert the expectations of what a birdcall can be.

One difficulty in the analysis of birdcall is the "rapid and strong frequency and amplitude modulations" (Bonada et al., 2016) common in many species. This can lead to an oversimplification of the calls because if the analysis software is not sampling the sound wave frequently enough, it will miss most of the nuance. Once the call is analyzed properly, it is then possible to look at the whole call, and break it down into smaller features. This method opens up the possibility of synthesizing each of those small features one by one and stringing them together after the fact, in order to create a more complex birdcall using one synthesis path. Bonada then uses this extracted information to train a Markov chain to produce statistically realistic, yet not predetermined, complex birdcall with synthesis. For the purposes of this paper, building a machine learning system to generate birdcall with a Markov chain is not within the usability scope that I am looking for because of the significant amount of programming model training, and computer memory, that are required to set up an artificially intelligent system. However, an analysis of Bonada's methodology leaves me with two possible methods to generate more complex birdcalls: 1) break the call up into bite-sized impulses, synthesize each separately, then string them together, or 2) develop a synthesis pathway with complex envelopes on many modulation parameters in order to hand-program each phase of the call. Whereas the second method would require a lot more time to set up each parameter and envelope in order to generate a complex bird vocalization all at once; the first method requires a more methodically planned approach. Constructing a synthesizer that could generate a complex vocalization all at once would be easier for a musician or composer to use as a versatile instrument, but alternatively, building that vocalization piece by piece would surely allow a level of detail that may not be achievable by the other method.

Previous Synthesis Models

We can look at the way that Andy Farnell uses the physiological information he gathered to inform his physical modeling synthesis of oscine, or songbird, calls (Farnell, 2010). He outlines his method as: "[treating] the bronchial passages as two separate pulse generators," and then "[combining them] with AM/FM and a resonant filter," in order to achieve desired effect. This process begins with three pulse oscillators, "one is used for each bronchus/labia port, and one for the syrinx simulation." The first two pulse oscillators are meant to represent the sound created by each of the two labia inside of the syrinx – remember that a bird can use each lung and bronchus independent of the other, which is why it is important to have two sound sources. The waves from these two oscillators are summed using ring modulation to represent the way that the sound produced by each bronchus and labia pair are combined within the syrinx. This combined waveform is then used as the basis of a frequency modulation operation. After the creation of this complex wave, Farnell sends the signal through a series of band-pass and high-pass filters intending to simulate the spectral effects of a bird's trachea, and beak. One advantage of this method of birdcall production is the extent to which avian physiology is taken into account. If one was looking to create as realistic a birdcall as possible, surely the best way is to attempt to recreate the actual sound-shaping processes that are occurring inside of the bird. One disadvantage of this method is that the construction of this synthesis pathway results in eleven parameters that need to be tweaked and adjusted in order to create a compelling timbre. Furthermore, if one is looking to create a truly realistic instance of oscine song, or if one is not hindered by inordinate levels of complexity, it would necessitate envelopes on each of those eleven parameters, and possibly even require the sound designer to string multiple short fragments together in order to comprise a full song. To create complex birdcalls would require even more parameters and envelopes. Farnell adds:

> Going back to our earlier observations on living things, they make sounds that have semantics and even simple living things will make astonishingly complex neural patterns to control movement and voicing. Let's suppose we could simplify birdcall to one envelope per parameter and make each sound a short phrase lasting less than one second. Then we could connect together these phrases into more elaborate calling. If we control each envelope that has an attack time, decay time, initial level, and final level we will have 44 parameters to play with, a lot for what will be unsophisticated behavior (Farnell, 2010).

In short, if a sound designer were looking to design a bird song and were willing to put the time into fine-tuning numerous parameter envelopes, then this would be a great method. Another consideration with this method of synthesis is that while it is well suited to the synthesis of pitched, oscine birdsong, it is not very adaptable outside of this realm. This may not be the method you would want to use in order to synthesize the birdcall of other families of birds like ducks, crows, or parrots.

Another implementation of this model of synthesis is by Pendharkar, an engineer from with Audio Technologies, Singapore. Although Pendharkar's version is somewhat simplified as compared to the original model laid out by Farnell. Pendharkar uses three oscillators: one to carry the frequency, one to modulate frequency, and one to modulate amplitude. A series of five envelopes are used to control modulation levels as well as oscillator frequencies. The implementation of this model (Figure 2) employs a user interface, which contains fifteen knobs, each controlling a different parameter of the birdcall synthesis.

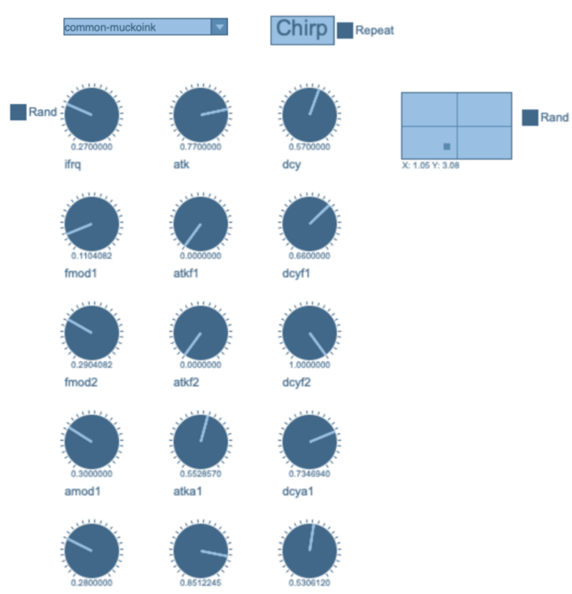

Figure 2 (Pendharkar, 2010)

While the design of the synthesizer is fairly straightforward, Pendharkar's user interface creates some limitations in function. The advantage of this style of birdcall physical modeling is that it takes the basic FM/AM structure of Farnell's method, and boils it down to its simplest parts. One major limitation is that there is no way to adjust the length of the birdcall. Additionally, while the use of dials (for frequency, attack, and decay of the carrier signal, two layers of frequency modulation, and two layers of amplitude modulation) create a very straightforward user interface, it leaves much to be desired in terms of customization.

Controller Methods

There are also a variety of ways in which each individual birdcall could be triggered. One way would be using a user interface to trigger the generation of each call. All of the parameters would be preset, and unless any of the parameters were changed between triggers, every individual triggered call would generate the same call. This is useful for the fine-tuning of the timbre and shape of the call, as well as for building a more complex call out of multiple smaller ones. These preset parameters could be controlled by envelopes, making it possible to create a great variety of complex calls. The downside of this method is that it requires more setup time to produce each individual call. Depending on which implementation method you utilize for your synthesis, there could be potentially dozens of parameters that need to be set before a single effective call can be produced. This setup time might make this method less useful as a live instrument, but more useful as a sound design tool, or a tool that is used to build a sample library that could be played live.

Another method would be triggering the calls with a MIDI keyboard. The keyboard could control the pitch of the birdcall as well as the duration. A method like this would enable the musician or composer to literally use the birds as an instrument that could create composite melodies and harmonies with the individual calls. While this would increase the creative control of the use of the birdsong, it could limit the ability to have calls with complex movements as well as the ability to string shorter calls together to build a more complex call. Any complexities or changes in timbre over the course of the triggered sound would either depend on pre-set envelopes, which could add overhead setup time to the use of the keyboard, or would have to be controlled by pitch and modulation wheels on compatible keyboards. Using the pitch and modulation wheels for expressive changes within the call, while potentially effective for the creation of a complex call, run the risk of being incredibly difficult to control without extensive practice, and even then it could be incredibly difficult to make convincing birdcall in a live setting using physical controllers, unless the parameters that they affect are extremely controlled.

Another more experimental controller involves the use of a mouth controller (Gamhewage et al., 2004). The authors propose that “the use of the mouth or vocal tract for controlling audio synthesis” is advantageous because it is natural for humans to use our mouths to control sound. In this method, a camera tracks the user’s mouth, and the size and shape of the mouth is mapped to parameters relating to synthesis using the physical model of birdcall. In their paper, it does not appear that the authors have fully implemented the connection between the face tracking and the synthesis but are instead proposing that this would be a good method. While this is a very interesting thought process, I am not sure that simply the shape of the human mouth would be able to be mapped to birdcall generation in a meaningful way, primarily because this method does not seem to account for a way to control pitch or loudness. Additionally, birds have a very different physical method of producing their calls than humans do, with individually controlled lungs and a syrinx, instead of a larynx. Therefore, while I think theirs is an interesting approach towards controlling the generation of birdcall, I am not sure if it would be the most versatile or best fit for musical and creative scenarios.

McCrosson Approach to Birdcall Physical Modeling

In this section I describe my approach to the physical modeling of birdcalls, which utilizes a synthesizer patch I developed in Max/MSP. My implementation reflects bits and pieces of examples previously discussed. As seen in Figure 3, the basis of my patch is one sine wave oscillator, frequency modulated by another sine wave. The pitch of the carrier sine wave, the modulation index, and the ratio of the carrier frequency to the modulating frequency, are all controlled by envelopes with continuous curves. After the initial FM synthesis, the signal is fed through an amplitude modulation process, seen in Figure 4. The AM frequency is a static value, while the depth of the AM synthesis is controlled by another envelope. Next, the signal is passed through a resonant low-pass filter.

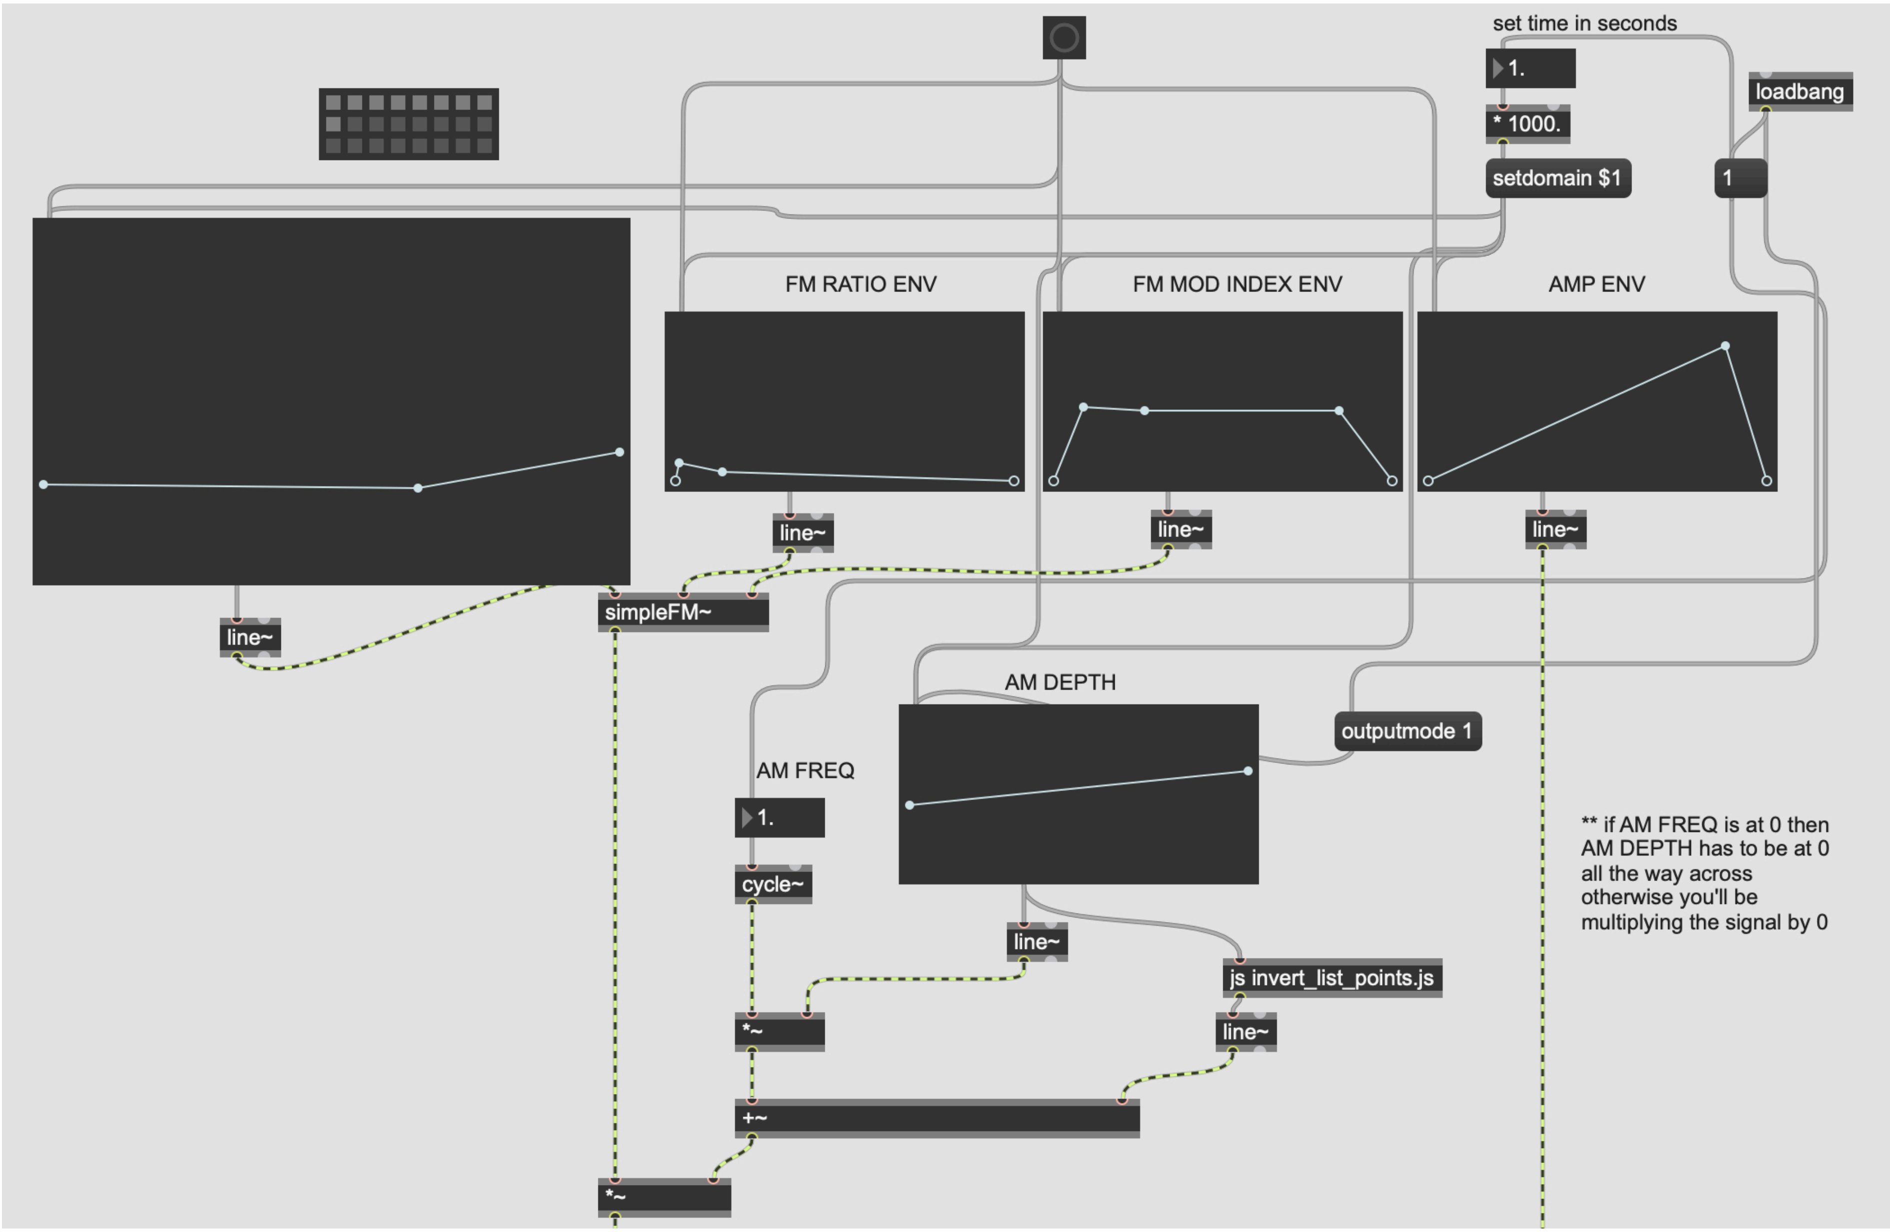

Figure 3. User Interface, McCrosson Birdcall Physical Modeling

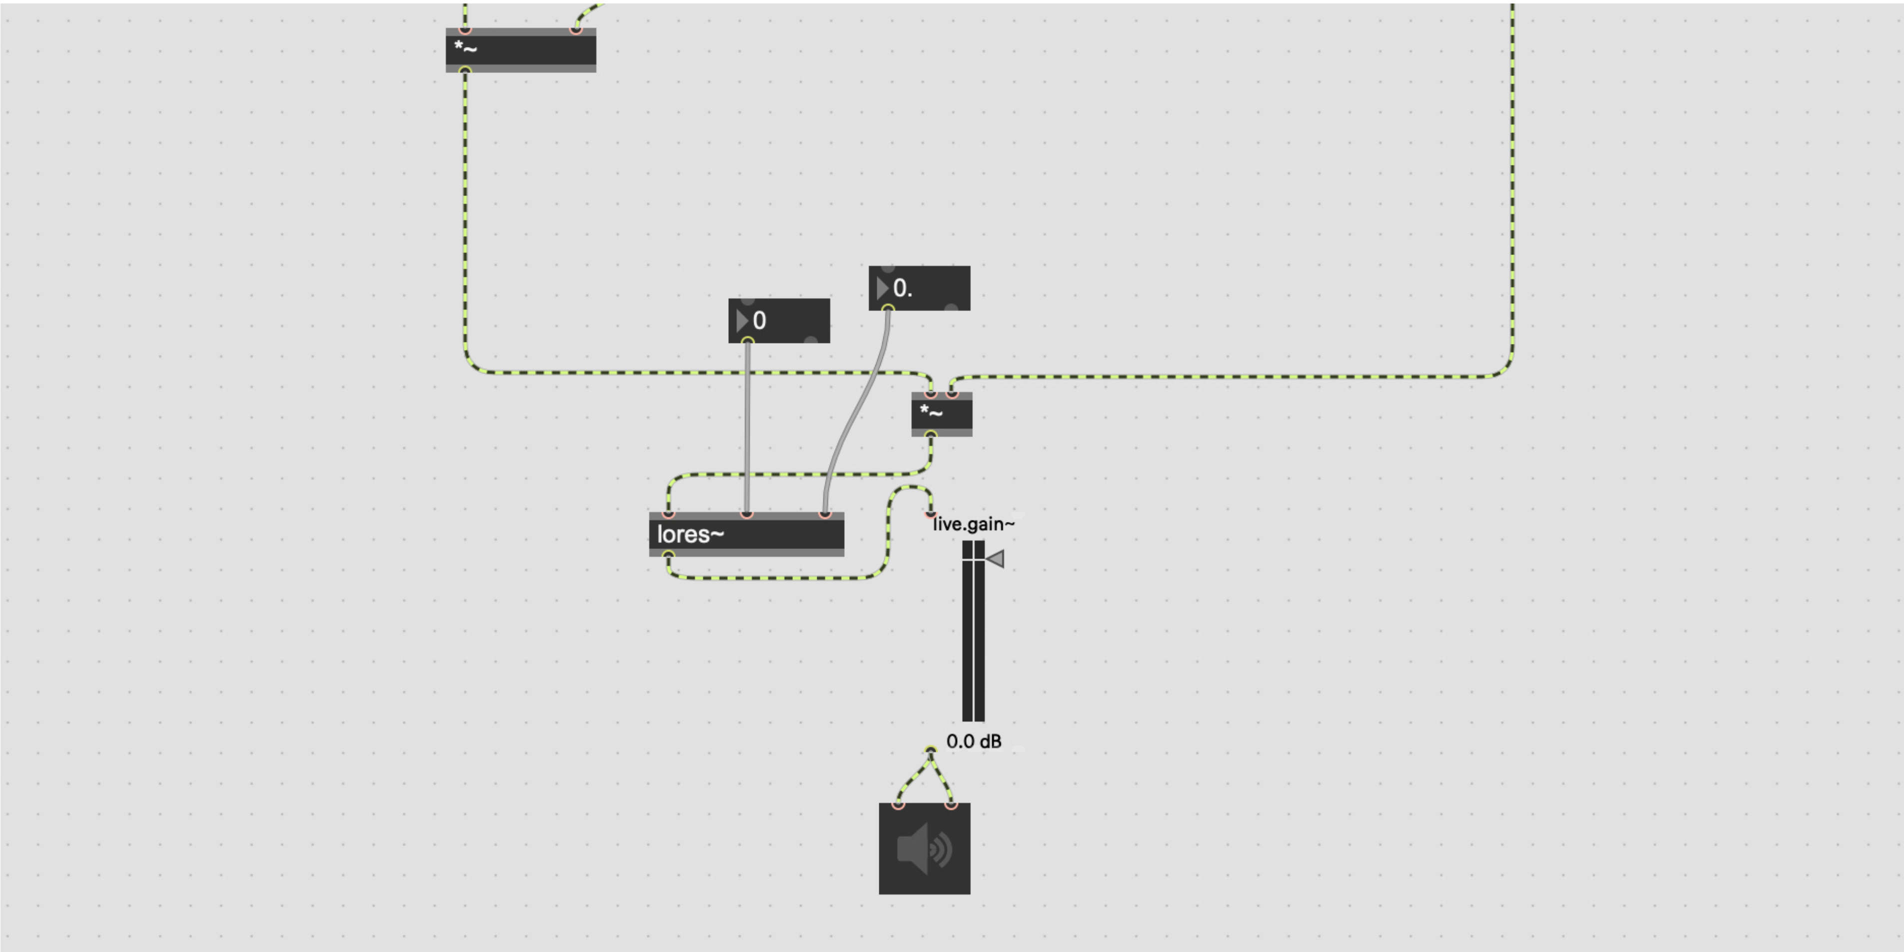

Figure 4. Amplitude Modultation, McCrosson Birdcall Physical Modeling

In terms of developing the timbre of the call, pitch, frequency modulation, and amplitude modulation simulate the activity of the lungs and syrinx in the physiology of a bird. The resonant filter simulates the effect of beak shape and size. The overall duration of the birdsong is set in seconds, which adjusts the lengths of all the corresponding envelopes so that all the envelope activity is synchronous. Additionally, there is an amplitude envelope that controls the overall volume of the birdcall over time.

At this point, readers may recognize how my approach to birdcall physical modeling is informed by: 1) a cursory knowledge of the avian vocal tract; 2) the implementation of other researchers before me; and 3) a desire to create a synthesis engine that is capable of generating both very realistic birdcall, and also more musical, artistic, and melody-based calls. Through a period of trial and error, I have developed an understanding of how different synthesis parameters effect the end sound, and how this knowledge can be used to my advantage. This solution combines the Farnell's physical knowledge of bird anatomy with the adaptability of a simpler software synthesizer, containing only five enveloped parameters as compared to Farnell's eleven. My implementation allows for the generation of a call of any duration and any pitch, critical for the utility of this synthesis tool in a compositional setting. Using continuous envelopes on these parameters means that there is a great deal more customizability than a traditional attack-decay-sustain-release (ADSR) envelope as in Pendharkar's implementation. Currently, none of these parameters are heavily constrained in order to give composers and musicians freedom to create both realistic avian vocalizations and synthetic ones. However, it might be helpful to document value ranges that create certain effects in order to cut down the amount of experimentation users of this synthesizer would need towards generate meaningful sounds. My implantation solves this problem in part by allowing the user to store a number of presets, so that they may save a configuration if it is found to create a desirable sound. Another advantage is that this system is versatile enough to create sounds that resemble both birdcalls and birdsongs; i.e., it can be used to produce sounds that function percussively or melodically.

One possible improvement or future development of this synthesizer would be to make it compatible with a MIDI keyboard. Controlling the pitch and activating the ADSR envelopes with a keyboard would make it much easier to interact with the synthesis model. This way, a musician would be able to interact with the synthesis in real-time, creating birdcalls of various lengths and pitches. The drawback from this method would be that it would be more difficult to implement complex changes in pitch during the call because the pitch would no longer be controlled by an envelope. The possibility to play chords and execute multiple bird songs at once might outweigh the resultant limits in pitch control. This interface could allow for a more musical interaction with the birdcalls; thinking of each call as a note is a more traditional way to craft a composition. This method, however, would not be well-suited for the creation of more complex calls through the stringing of multiple shorter segments. Instead, it would be limited to shorter, simpler birdcalls. Ultimately, when confronted with the choice between adapting the synthesizer to be played with a keyboard versus maintaining the flexibility of the complexity of the birdcalls, I would opt for the flexibility. The non-keyboard version would still allow musicians to build a library of birdcalls that could be used with a sampler to create a similar effect.

Methodology of Composition

Using my Max/MSP patch, I created a variety of birdcalls that I used in a composition called *digital_forest*. Some were modeled off of real birdcalls, like the cardinal, pigeon, and mourning dove, while others were created just to sound like a generic songbird or tweet. I synthesized short calls in one shot; for longer calls I broke them up into smaller bits, synthesized each section individually, and then strung them together later.

For the beginning of this process, I focused on building a small library of realistic and convincing birds that I could then use as compositional material. The beginning of *digital_forest* is composed with intention to sound like a soundscape, as if it were recorded in a real-world natural environment. As the piece moves forward, I introduce rhythms with the birdcall, insect, and environment sounds that I synthesized, to bring attention to the fact that this is a composed soundscape. From here, I begin to explore ways to alter and distort the birdcall that I have synthesized. Slowly, realistic birdcalls and birdsongs transition back to the basic waveforms from which they are generated, rendering what used to be beautiful and natural sounding into something synthetic and quite grating. My intention is to display the versatility behind the synthesis method that I have built; that is, to share with the listener the power of this program to morph from convincing birds to any variety of synthetic, grating sounds. There is a moment of realization in the piece when the listener recognizes that the birdcall they have been listening to is nothing more than computer generated signals.

Summary

Birdcall has long been a source of inspiration for composers; first as a muse for crafting melodies and phrases, and later as actual compositional material due to the advent of recording technology. Physical modeling synthesis is the newest frontier of birdcall as compositional material, allowing composers to have full customization and control of the call. Taking inspiration from the physical functionality of a bird's vocal tract, many researchers have determined that the most effective approach to physical modeling of birdcalls utilizes frequency modulation and amplitude modulation (Farnell, 2010). By adapting the approach taken by Farnell and Pendharkar, I built a synthesis mechanism in Max/MSP that uses frequency and amplitude modulation, a low-pass filter, and a variety of envelopes to create realistic bird sounds. This method is versatile, and has the capability to create both oscine and non-oscine birdcalls. It can create calls of variable length and with a lot of complexity and internal movement. I offer my approach to birdcall physical modeling to be an effective and realistic implementation of birdcall synthesis that also gives composers the creative and artistic freedom to develop birdcalls both musical and abstract.

References

Bonada, J., Lachlan, R., Blaauw, M. (2016) Bird song synthesis based on hidden Markov models. Interspeech 2016, 2582-2586. DOI: 10.21437/Interspeech.2016-1110

Chiat, L. (2005). Olivier Messiaen's Catalogue D'oiseaux: A Performer's Perspective, PQDT - UK & Ireland.

Cornell Lab of Ornithology. (2014). *Anatomy of bird song slides*. Bird Academy. Retrieved from https://academy.allaboutbirds.org/anatomy-of-bird-song-slides/

deSilva, G. C., Smyth T., & Lyons M. J. (2004). A novel face-tracking mouth controller and its application to interacting with bioacoustic models. *Conference on New Interfaces for Musical Expression*, *June 2004,* 169-172.

Farnell, A., & ProQuest. (2010). *Designing sound* (pp. 571-578). MIT Press.

Ferguson, Donald N., Ferguson, Donald Nivison, & ProQuest. (1954). *Masterworks of the orchestral repertoire: A guide for listeners* (pp. 457-460). University of Minnesota Press.

Head, M. (1997). Birdsong and the origins of music. Journal of The Royal Musical Association, 122, 1-23.

Judd, T. (August 26, 2019). Handel's musical depictions of birdsong. *The Listener's Club*. Retrieved from https://thelistenersclub.com/2019/08/26/handels-musical-depictions-of-birdsong/

Kahrs, M., & Avanzini F. (2001). Computer synthesis of bird songs and calls. *The Journal of the Acoustical Society of America, 111(5), 2392-2392.*

Larson, J. D. (Novmeber 27, 2017). *Björk: Utopia*. Pitchfork. Retrieved from https://pitchfork.com/reviews/albums/bjork-utopia/

Pendharkar, C. (December 10, 2014). *birds*. Github. Retrieved from https://github.com/notthetup/birds

Ritchison, G. (n.d.). "Bird vocalizations: Production, function, and investigation." Eastern Kentucky University. Retrieved on April 3, 2020 from http://people.eku.edu/ritchisong/boisestate.html

Symes, C. (November 29, 2004). *Setting the record straight: A material history of classical recording* (pp. 52). Wesleyan University Press.

Truax, B. (2008). Soundscape composition as global music: Electroacoustic music as soundscape *. Organised Sound, 13(2), 103-109.

Westerkamp, H. (1992) *Beneath the forest floor*. Hildegard Westerkamp: Inside the soundscape. Retrieved from https://www.hildegardwesterkamp.ca/

Wilkinson, W., & Reiss, J. D. (February 10, 2016). A synthesis model for mammalian vocalisation sound effects. *AES 61st International Conference*.

Yalcinkaya, G. (June 19, 2018). *"Björk explains designs of 'magically utopian' sets and costumes."* de zeen. Retrieved from https://www.dezeen.com/2018/06/19/bjork-explains-designs-of-magically-utopian-sets-and-costumes/

Zysman, D., Méndez, J., Pando, B., Aliaga, J., Goller, F., & Mindlin, G. (2005). Synthesizing bird song. Physical Review. E, Statistical, Nonlinear, and Soft Matter Physics, 72(5), 051926.