Revista Vórtex | Vortex Music Journal | ISSN 2317–9937 | http://vortex.unespar.edu.br/
https://doi.org/10.33871/vortex.2024.12.9521

ORIGINAL ARTICLE – DOSSIER "NEW SOUND ECOLOGIES"

Ecosystem [512]: Acoustic Ecology Surveys as Music

Garrison Gerard

University of North Carolina at Pembroke, Department of Music | Pembroke, North Carolina, USA

Resumo: A ecologia acústica é uma disciplina em rápida expansão que pode revelar detalhes sobre ecossistemas que não são visualmente aparentes, e a gravação de campo tem sido utilizada de formas cada vez mais variadas em composições musicais, abrindo portas para projetos colaborativos que transitam entre a ciência e a arte. O projeto *Ecosystem [512]* está enraizado em levantamentos acústicos que realizei ao longo de nove meses nos Parques Nacionais da Islândia. Esses levantamentos forneceram mais de 10.000 horas de gravações que foram processadas, comprimidas, combinadas e analisadas antes de serem usadas como material fundamental para uma composição musical para eletrônica e clarinete. A peça utiliza múltiplas estratégias, como composição algorítmica, para lidar com a transição das gravações acústicas brutas dos levantamentos para o produto musical final. O processo de composição desta peça revela insights tanto sobre os processos de ecologia acústica quanto sobre considerações de composição musical.

Palavras-chave: Composição Musical, Ecologia Acústica, Música Eletroacústica, Bioacústica, Gravação de campo.

Abstract: Acoustic ecology is a rapidly expanding discipline which can reveal details about ecosystems that are not visually apparent, and field recording has been used in an increasing variety of ways in music compositions, opening the door to collaborative projects that walk the line between science and art. *Ecosystem [512]* is rooted in acoustic surveys that I carried out over nine months in Iceland's National Parks. These surveys provided more than 10,000 hours of recordings that are processed, compressed, combined, and analyzed before being used as the fundamental material for a music composition for electronics and clarinet. The piece uses multiple strategies such as algorithmic composition to handle the transition from rough acoustic survey recordings to final music product. The process of composing this piece reveal insights about both acoustic ecology processes and music composition considerations.

Keywords: Music Composition, Acoustic Ecology, Electro-acoustic, Bioacoustics, Field Recording.

Received on August 6, 2024. Approved on October 15, 2024. Available online: December 3, 2024. Guest Editors: Dr. William Teixeira; Dr. Jacek Smolicki.

M*y leg sunk waist-deep into the freshly fallen snow. Birds circled the cliffs above me on either side as I struggled to gain purchase on the slick rock face. Ribbons of long cooled magma traced an ancient path up the stiff cliff face on my right I couldn't follow. With some effort, I compacted the snow under my boot and made one ephemeral stair after another to climb into the valley. As the birds' mocking grew more audible, I finally hoisted myself into the ancient lakebed turned snow strewn meadow. Now visible were the countless waterfall festooned valleys that marked the path to the peak. I knelt beside the field recorder on the valley floor and cracked open the case.*[1]

1. Introduction

Acoustic ecology and music composition share a central focus on listening to our environment. This shared focus is evidenced in the many practitioners whose work has bridged both disciplines such as Murray Schafer (1977, 1994), Barry Truax (1999), and Hildegard Westerkamp (2002), whose work has leaned toward the musical side and Bernie Krause (2012), Bryan Pijanowski (2011), and Almo Farina (2013) who fall more on the side of scientific research. This overlap has led both to engaging musical compositions that use field recordings such as Truax's "Pacific" (1992) and scientific studies based in listening such as Krause's study of four areas in Sequoia National Park (2010). Following in the footsteps of these researchers and composers, I set out to compose a musical composition alongside an in-depth acoustic survey. The goal being to create a dialogue between elements of my work and create a final composition built on acoustic survey recordings and data alongside an acoustic survey whose analysis is informed by musical listening and observations.

[1] Description of the author's personal experience.

Soundscape ecology[2] as theorized by Brian Pijanowski lays out the ways that sound can provide insights into ecosystems, particularly highlighting how soundscape patterns elucidate the impact of human activity on biologic patterns (2011, p. 201). Acoustic ecology methods have seen rapid development alongside the proliferation of affordable recording technology (Hill *et al*., 2018, p. 1200); spurred on recently not in the least by the refinement of metrics such as acoustic indices (Bradfer-Lawrence *et al*., 2023). The acoustic surveys carried out in this project combine insights from acoustic indices with other metrics such as species identification enabled by the Arbimon platform (Aide *et al.,* 2013). In total, the acoustic surveys I carried out captured more than 10,000 hours of audio in the form of one-minute recordings from 15 field sites across Iceland's National Parks.

How do you compose an eight-minute piece of music using more than 10,000 hours of recorded audio? How do you work with one-minute recordings intended for acoustic analysis in a creative setting? Both of these and others are questions that accompanied my composition process of *Ecosystem [512].* The piece is written for clarinet and live electronics and is both an attempt to communicate about acoustic research and a personal exploration of Iceland's soundscapes through music and compression of my experiences in these ecosystems.

Ecosystem [512] compresses a large amount of data and recordings into a much shorter musical composition. The fixed media portion of the piece uses pre-defined algorithms to combine the field recordings and multiple compression strategies are used so that real recordings are the fundamental organizing principle in the piece. The choices of how to compress that data—which data is included, how is it layered, etc.—constitute the fundamental composition decisions in this piece.

This project builds on both past acoustic surveys and musical compositions that take a variety of perspectives on natural sounds. Acoustic ecologists such as Bernie Krause (1993, 2012)

[2] There is a voluminous and worthwhile debate on which term best represents the practice of using recorded sound to investigate the acoustic dimension of ecosystems: the terms acoustic ecology, soundscape ecology, landscape ecology, ecoacoustics have all been used by various practioners with varying shades of meaning. (not to mention the related disciplines of bioacoustics and sound studies). Without wading too deeply into these debates, in this paper I use the term acoustic ecology to broadly refer to the practice of using sound recording to study ecosystems, except where referencing a researcher who uses a different, specific term. I do not mean to use the specific definition of acoustic ecology as set forth by Murray Schafer and the World Soundscape Project (Schafer, 1977, p. 271).

and David Betchkal (US Department of the Interior, 2018) have used field recordings to examine the acoustic dimension of ecosystems. Their and others' approaches provide a blueprint for methods such as temporal audibility analysis and acoustic indices that I use to examine the impact of tourist noise on Iceland's National Parks. The use of natural sound has a long history in music composition stretching at least as far back as Ottorino Respighi's *Pines of Rome*. *Ecosystem [512]* piece departs from some earlier approaches by using recordings intended for research rather than field recordings made with high-fidelity microphones as its fundamental audio element, but it maintains an awareness toward musical aspects of natural sounds.

The first section of this paper deals with the passive acoustic monitoring surveys that provide the recordings principally used in the composition of the work. I carried out these recordings over the course of nine months; they provide a unique window into natural ecosystems of Iceland and the impact of human sound on them. While incredibly useful from a data perspective, using passive acoustic monitoring recordings in a creative setting introduced multiple issues including the recordings length, audio quality, and overwhelming quantity. The data derived from these recordings also influences the piece on multiple levels.

Next, I discuss the creative and philosophical considerations that preceded and accompanied the compositional process. The composition of this piece was necessarily iterative and revolved around monthly trips to the recording locations. The role of location and place is central in the piece—not only in that the recordings are from real locations, but in the influence of these places on the composition of the piece itself.

Finally, I examine the technical considerations in composing the electronics, clarinet part, and visuals for the piece. The electronics use multiple strategies that rely on the recordings themselves to provide the musical form. For instance, the fundamental sound layer relies on a pre-defined algorithm to autonomously combine the field recordings into a soundscape piece—this algorithm relies on real-time analysis and will produce a different soundscape piece every time. The other layers of the piece, such as the clarinet part, build on the preceding layer, creating a complex web of sound.

2. Acoustic Surveys

The research underlying this piece is a nine-month acoustic survey of Iceland's National Parks investigating the impact of anthropogenic noise on ecosystems within the parks. This survey was possible because of the support of a Fulbright-National Science Foundation Fellowship. In each of the three National Parks—Jökulsárlón, Snæfellsness, and Þingvellir—field sites were selected both close to and far removed from high human traffic areas (see figure 1). In total, there were 15 field sites: 4 in the west in Snæfellsness, 4 more centrally located in Þingvellir, and 7 in the east in Jökulsárlón (some site icons overlap in figure 1). Analysis of the recordings provided information such of noise level, species population, and acoustic indices at each field station, demonstrating part of the acoustic impact of tourism. A further goal of the project is to provide a baseline of noise levels and a set of recordings that document the current state of Iceland's natural ecosystems.

FIGURE 1 – Map of Acoustic Survey Locations

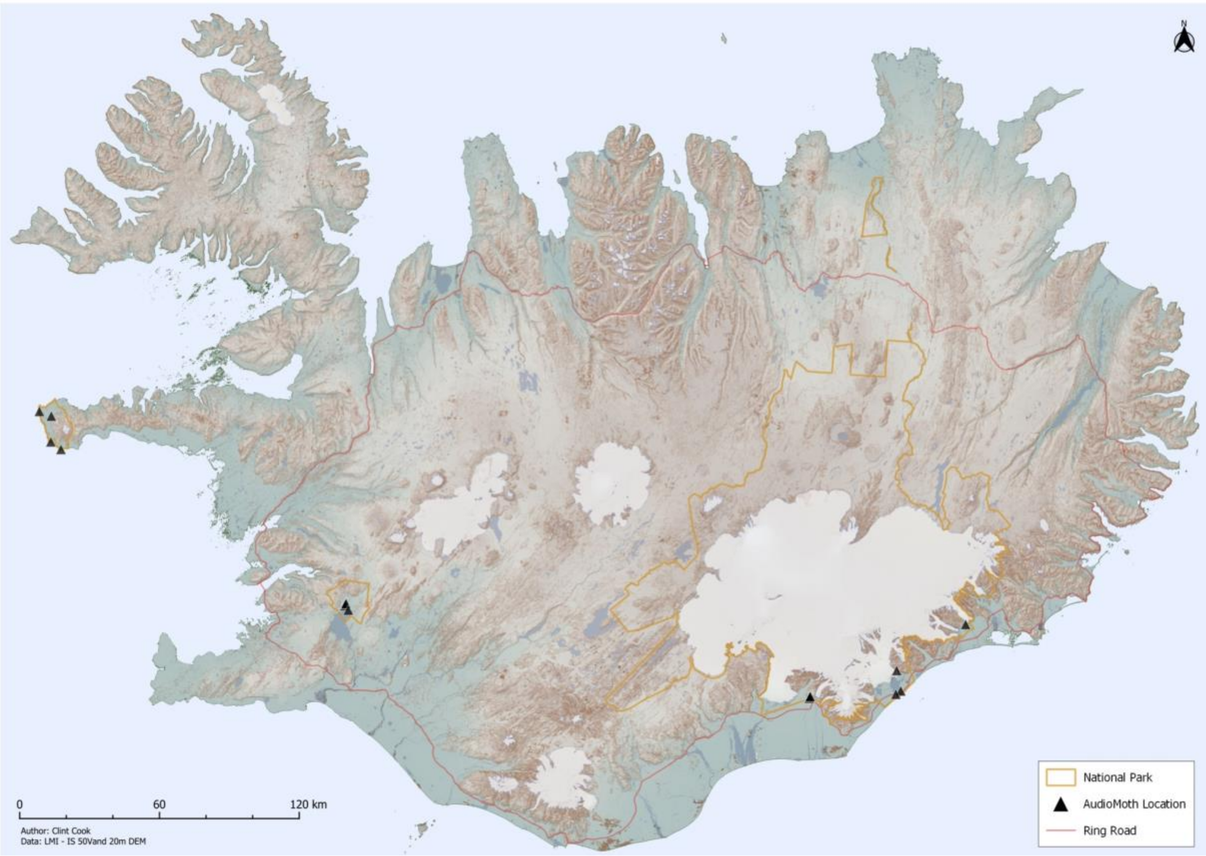

Source: Clint Cook (2024)

The surveys were conducted from September 2023 until May 2024. Not every site was continuously recorded over the entire period, but all sites were recording for at least three months, and multiple sites (especially those in Jökulsarlón National Park) were recording for the entire nine-month period. AudioMoth recorders were used to carry out the surveys. Each AudioMoth was configured to record for one minute and sleep for four minutes at a 48 kHz sample rate. This recording schedule allowed for the best balance of energy consumption and recording time, leading to a maintenance schedule of once per month.

Carrying out acoustic surveys in the unpredictable weather of Iceland introduced unique challenges, not the least impacted by the recording period covering the entire winter season when weather was considerably colder. Cold weather of -10 to 5 degrees Celsius necessitated the use of lithium batteries to prevent battery drain. The deforestation of Iceland also meant that mounting methods had to be adjusted since the default AudioMoth straps are designed to be affixed to trees—instead, posts were used alongside an adapted Velcro strap to affix the recording devices. All AudioMoths survived the weather, with the most abused being the recorders mounted in areas with loose rock such as at the Saxholl volcanic crater in Snæfellsness.

The recordings provide valuable data to analyze the ecosystems within the parks. The AudioMoths were all calibrated using a decibel meter,[3] allowing for noise level analysis using the PAMGuide analysis package (Merchant *et al.* 2015) (see figure 2). The Arbimon platform allows for species identification: by identifying a single sample of a target species, the platform can analyze all other recordings and automatically label the same species (Aide *et al.*, 2013). Beyond specific species identification, the Arbimon platform can also be used to identify sound events such as motor noise or rainfall. This is achieved by first detecting any audio event that departs from the standard deviation by a specified amount (60% was used in this study). These audio events are then

[3] AudioMoth recording devices were calibrated both before deployment and in the field by playing a normalized audio signal of 94 dB and comparing the value of the recorded sound file with the value of a calibrated decibel meter. This accounted for differences between devices and for the impact of the case and specific deployment orientation on sound level.

clustered by amplitude and frequency. Similar events will tend to cluster closer together, allowing for quick identification of sound groups.

FIGURE 2 – Median Loudness Across Principal Recording Sites in Jökulsárlón National Park, in Decibels

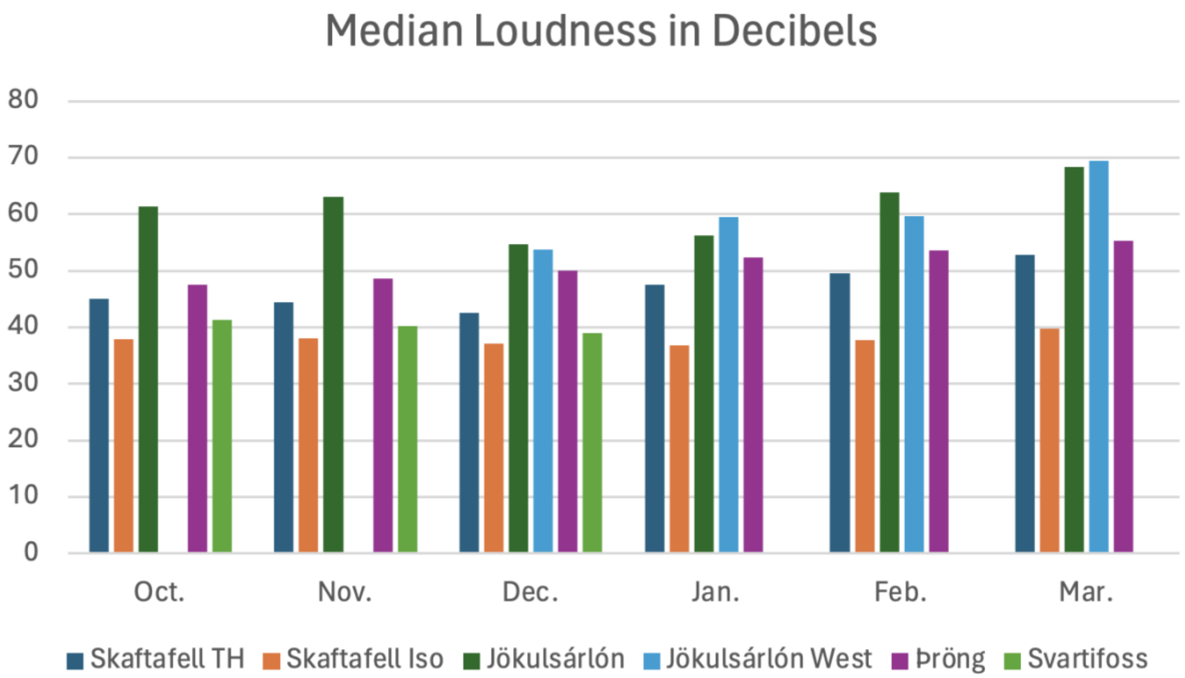

Source: Author (2024)

Acoustic indices provide a powerful tool to examine large swaths of data. While efficient and powerful, acoustic indices must be used with caution, particularly when using them as a proxy for biodiversity (Bradfer-Lawrence *et al.*, 2023, p. 2195-2196). Two crucial elements of using acoustic indices are to understand what the acoustic index is actually measuring and what impact adjusting the analysis parameters has on the output file. Two main acoustic indices were used in this study; as the goal was to examine the impact of anthropogenic noise, I used the Normalized Difference Soundscape Index (NDSI) and the Bioacoustic Index. The NDSI compares the amount of activity in the biophony band (2,000-11,000 hz by default) and the anthropogenic band (1,000-2,000 hz by default) (Kasten *et al.* 2012). The NDSI reports a value from 0. to 1.; 1. being sound recordings that are entirely comprised of biophony and 0. Being sounds comprised exclusively of anthrophony (see figures 3-5). The Bioacosutic index is similar, but it also takes into account how many frequency bins exceed the minimum level in the analysis band, thereby measuring how many species are responsible for the activity in the biophony band (Boelman *et al.* 2007).

The recordings from the national parks reveal fragile ecosystems with considerable anthropogenic noise. Across Jökulsárlón National Park, noise is directly proportional to the amount of tourist activity in the area, with the loudest sites consistently being around the glacial lagoon (Jökulsárlón) where there is constant vehicle and boat traffic (see figure 2). Acoustic index analysis reveals similar patterns: the normalized difference soundscape index compares the amount of acoustic activity in the anthropogenic band (1 kHz to 2 kHz) and the biophonic band (2 kHz to 11 kHz) (Kasten *et al.,* 2012). In areas with considerable tourist activity biologic activity peaks during the dawn and dusk choruses (see figure 3). Contrasting this, sites on the periphery of high human activity show much higher levels of biological activity during midday as animals move away from areas frequented by humans (see figure 4). Finally, remote areas with weather that is not conducive to biologic life show no distinguishable pattern of biophonic activity (see figure 5).

FIGURE 3 – Skaftafell Trailhead NDSI by Hour of the Day

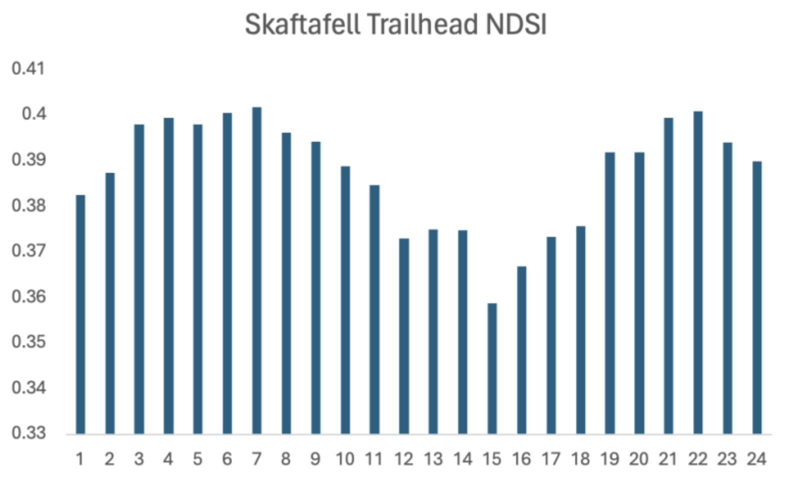

Source: Author (2024)

FIGURE 4 – Skaftafell Isolated Area NDSI by Hour of the Day

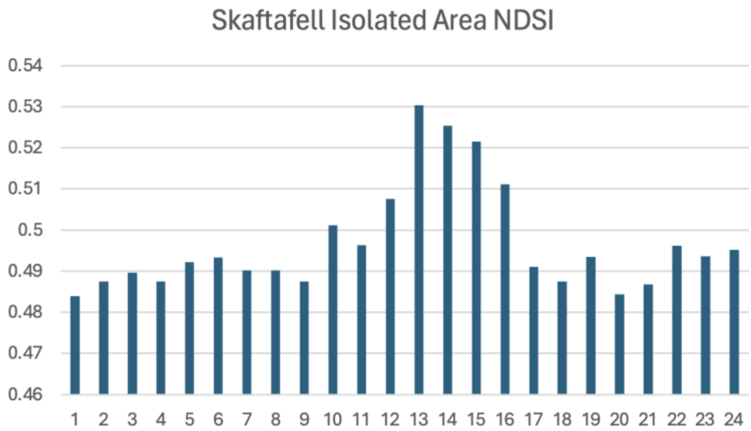

Source: Author (2024)

FIGURE 5, Þröng NDSI by Hour of the Day

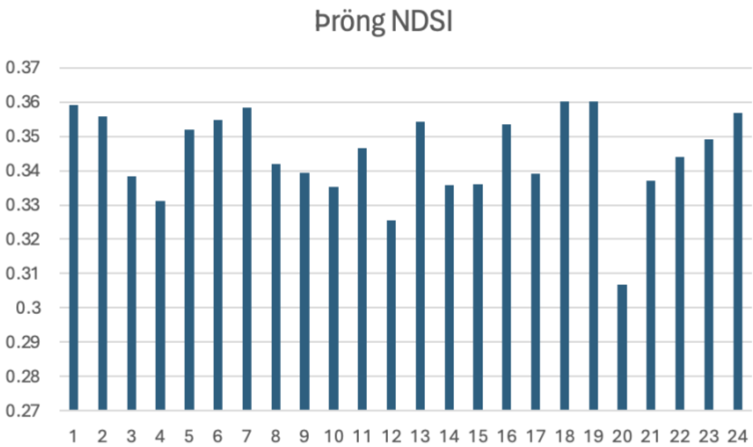

Source: Author (2024)

3. Connecting Acoustic Ecology to Composition

The AudioMoth recorders used in the acoustic surveys lasted for roughly 30 days before their batteries depleted, necessitating monthly trips to the recording sites. While logistically challenging to reach all the sites each month, this created a situation that allowed for repeated listening and reflection in each of the listening areas. These visits significantly enhanced the analysis and composition portion of the project: when listening back to the recordings, I was more easily able to visualize the surroundings of the field site and base my listening in fresh, real-world observations.

On these return trips, I was also able record the environment using higher quality field recording microphones.

The piece connects to the acoustic survey on multiple layers. First, it uses the actual recordings from the AudioMoth devices, but on a deeper level, the piece builds on the data retrieved from the analysis. Information such as noise levels and species present has a direct impact on the sound. This is true not only of the electronic sounds where data can be easily fed into digital parameters, but also of the acoustic clarinet element where connections can be more tenuous.

Going in the other direction, the composition of the piece influenced the progression and analysis of the acoustic surveys. Composing this piece necessitated listening to the recordings repeatedly in a different mindset and time scale than is required for normal analysis. On multiple occasions this led to discoveries that would not have occurred without the compositional process happening alongside the analysis. Multiple bird species that were infrequently present and not picked up by the analysis software were only noticed because of the repeated listenings involved in the composition process.

The composition process of the piece itself takes its form from my experience carrying out these acoustic surveys and from the ecosystems I recorded. The acoustic surveys show a complicated network of interactions that react in unpredictable ways to small changes: for instance, a small change in the number of tour boats on Jökulsárlón has an impact on the bird activity miles away. Mirroring this, the piece is constructed in layers where a small change in a base layer can ripple up through the other musical elements and have an outsized impact on the final sound.

Noise is present in a variety of forms in the acoustic survey and its recordings. First there is acoustic noise, both anthropogenic and geophonic noise is present to varying degrees at each of the recording sites. On one particular visit to the recording sites at Skaftafell, myself and a colleague who was accompanying me had to yell to be heard over the roar of the wind. There is also statistical noise; with this many recordings, it can be difficult to separate correlation from coincidences. Mirroring this monolithic attribute, the piece is designed to be overwhelming, particularly on the first listening. There are three electronic sound layers, each with their own internal and intra-layer impact on other layers. On top of the electronics, the clarinet score and its processing provide two more layers, and the visuals, which are themselves constructed of multiple layers, provide yet another layer of information. The overall effect then is one of a kaleidoscope of sound and information that can't be held simultaneously in full awareness. The goal is for the listener to be able to move within the kaleidoscope of sounds, ignoring some layers and focusing on others, just as members of these ecosystems have to filter out noise to communicate in the actual survey locations.

4. Foundational Sound Layer

The foundational layer of the piece consists of field recordings from across Iceland which are combined based on an algorithm implemented in Max. The first step in employing this algorithm is to define all of the sounds in a way which can be processed by the patch. Recordings are sorted into location nodes—12 different field stations and recording sites across Iceland serve as these location nodes (see table 1)—many of these sites combine multiple field recording sites from figure 1. Three (Highlands, Vik, and Westfjords) use only field recordings captured during separate trips and serve as a foil to the acoustic survey recordings. In each location node, sounds are classified according to their dominant ecosystem and levels of human and biological activity.

TABLE 1 – Location Nodes

Location Node	**Dominant Ecosystem**	**Level of Human Activity**	**Level of Biologic Activity**
Fladjökul Lagoon	Glacial Lagoon	Low	Low
Breidermerkujökul	Glacier	Medium	Low
Highlands	Tundra	Medium	Low
Ice Cave	Glacier	High	Low
Jökulsarlon	Glacial Lagoon	High	Medium
Vik	Urban	High	Medium
Skaftafell	New Growth Forest	High	High
Snæfellsjökul	Coastal	High	High
Þingvellier	New Growth Forest	High	Medium
Þröng	Tundra	Low	Medium
Westfjords	Tundra	Low	High

Source: Author (2024)

Once sounds have been classified and loaded into the Max patch, the system can begin to weave them together into a new soundscape. The system does this by combining multiple different

 | https://doi.org/10.33871/vortex.2024.12.9521

analysis tools and using dynamic goal values. The first variable is the "desired activity level." By default, this level slowly increases over time until a new location is triggered, then it resets to its default value. To determine if the current sound meets the desired activity level, multiple descriptors from the zsa.descriptors package, including spectral density, centroid frequency, and noisiness are measured (Malt; Jourdan, 2008). The system takes a constant running average of these values over both a five and ten second window. When a departure from the running average is detected, the system interprets this as activity and increases the "current activity level." In this way, the current activity level is a measure of how much change there is in the overall soundscape since the last location was triggered. If there is a varied soundscape with significant changes in the measured parameters, the current activity level will easily outrun the desired activity level, allowing location nodes to naturally end or be triggered off. Conversely, a relatively static soundscape will have a comparatively low current activity level, and the desired activity level will quickly rise above the current activity level, triggering a new location node.

The rate at which the desired activity level increases will change if the system detects certain soundscape characteristics. The system uses amplitude and spectral standard deviation to approximate three soundscape scenarios. If the current amplitude is low with a wide spectral standard deviation, the system concludes that the current sound is most likely either quiet noise or dominantly water sounds and increases the desired activity level three times as quickly. If the spectral standard deviation is low with a high amplitude, the system identifies that there is a single dominant sound such as an overpassing airplane; it then adjusts the desired activity level to decrease instead of increase, thereby triggering fewer location nodes and thinning out the density of the soundscape. Finally, if the amplitude is high with a wide standard deviation, the system labels the current state as having multiple dominant sounds and returns the rate of change of the desired activity level to the default level. Together, these analysis tools combined with controlling the rate of desired activity ensure that overall density will respond to current conditions and not remain stagnant.

The system begins by triggering one location node at random, from there it will rely on analysis levels and the pre-defined sound characteristics to create a soundscape. Once the current level of activity falls below the desired level of activity, a new location node is triggered. Every location node has a probability to move to any other location node: the highest probability connections are those sites that are geographically close in the real word. The second highest probability are sites connected either by their level of human activity or biological activity. Other connections are either programmed with a very high improbability or are simply not possible. The system keeps track of which locations are currently playing and will not re-trigger a location that is already playing. The cumulative effect of these probabilities is that the system tends to stay in one geographic area until it has exhausted all sounds in the areas and then moves to another area.

When a location node is triggered, one of three possible sounds is selected at random and is triggered at a random point from 0%-80% through the sound file. This ensures that similar paths through locations can sound drastically different. Other sounds within the location node can be triggered if there is a large shift in the current activity. The system can also remove the longest playing location nodes if the overall amplitude is too high crosses a certain threshold; this is primarily a failsafe if too many location nodes are triggered and the overall volume gets too loud. In practice, location nodes usually play out their full sound file length.

The principal recordings that are triggered by each of the location nodes are between 10 and 15 minutes. There are also shorter recordings within each location node that can be triggered if the desired activity level reaches a sufficiently high value. Whereas the principal recordings are mostly of the ambient sound at the location node, the shorter recordings are usually of sonically significant events such as an animal sound or motor noise. By dynamically triggering these more noticeable sounds, the resulting soundscape will be a dynamic representation of possible sound events rather than a simple layering of ambient recordings.

For *Ecosystem [512]* I ran the system twice creating two 30-minute soundscape pieces. I then selected eight minutes from both renditions and layered them on top of each other. On one of the renditions, I also added three side-chain convolution reverbs that use various recordings of ice falling as the sampled impulse response, creating a shifting harmonic reverb effect. These two renditions of the generative soundscape layer together form the foundational sound layer of the piece. To allow for further layering, the overall volume of the background layer is relatively low, but sonic events are still audible, and this layer provides a constant background for the other layers to build on top of.

5. Mid-ground Sound Layer

Where the foundational layer is controlled by a pre-defined algorithm, the mid-layer is comprised of blocks of sounds arranged to convey the narrative trajectory of the piece. These sound blocks are each made up of one month of acoustic survey recordings overlapped on top of each other. The AudioMoth recordings are all one minute long. I began by de-noising the files to prevent buildup of recording noise from the relatively noisy microphones in the AudioMoths, then I combined the 8000-9000 recordings from a given month into a single sound block. The resulting sound changes significantly from one recording site to another. In very windy locations the most prominent pitches of the wind emerge in the sound block. In quitter locations, individual animal vocalizations are audible. Density and sound quality of the different sound blocks then becomes a principal formal element throughout the piece.

Since the conceptual basis of the piece is to examine the mass of recordings from nine months of conducting acoustic surveys, I wanted to represent the corpus of recordings in at least a somewhat proportional manner. Layering one month of recordings on top of each other fulfills this condition. Locations with high levels of noise create dense blocks of sound where few details can come through. Conversely, recording sites with relative quiet allow any activity to move to the foreground.

The mid-ground layer builds on the foundational layer by using blocks of sounds from locations that are more audible in the foundational layer at a given point in time. The sound of the compressed blocks is substantially different from the sounds heard in the foundational layer, but the connection of location serves to emphasize the trajectory of the piece. There are moments in the piece without any compressed sound blocks, and other moments where up to three can be heard at once.

The effect of layering a whole month of recordings creates both surprising and predictable results. In the predictable category are sites where the predominant sound is wind—here the result is a mass of wind noise that sounds more cacophonous than a single wind recording by an order of magnitude. In order to not disturb the trajectory of the piece, these wind conglomerations are further processed using a variety of filters so that the complex changing wind conglomeration can still be heard without dominating the frequency spectrum. Some of the more surprising results came from the quieter ecosystems where almost any sound captured by the recorders would be audible in the sound block—this led to multiple new sounds being discovered in the recordings that weren't detected in the acoustic ecology analysis tools. Because of this, I will continue to adapt and use this process of combining files into a single sound block to more efficiently listen to acoustic survey results.

6. Surface Sound Layer

When I compose a piece of electronic or electroacoustic music, I often find myself obsessing over the surface layer sounds—gestures, pitch complexes, timbres—all the ways that I can exercise my own compositional judgment over the music and recorded sound. In *Ecosystem [512]* when I got to point of adding a surface layer on top of the largely generated sound world of the foundational sound layers and compressed sound blocks, I found myself at a loss. Intuitively, I thought the piece needed more, that everything done so far was too generative, and surely, I needed to add more sculpted sounds to lend the piece direction. And yet, even after stepping away from the piece for more than a month, when I returned to it, I found it did not require further compositional changes in my mind.

I did finally make two small changes and add a touch of other sounds to the piece. Early in the piece, I removed the foundational layer so that the filtered and compressed sound blocks of wind Jökulsárlón shines through for less than a minute. Marking the entrance of this moment, I used the sound of ice falling that was also used as the impulse response for the convolution reverb on one of the foundational sound layers. The second sound I added was of ice falling on the glacier during an early spring melt: the glaciers had made a deep impression on me at this point, and while many of the field sites were relatively close to glaciers, none were close enough to capture glacial sounds such as calving. I felt that including sounds that I recorded on the glacier were appropriate in this context.

These small changes then brought the piece to a point where it was time to add the living element: to compose the clarinet part that would accompany the fixed media portion.

7. Clarinet Score

From the outset, I knew that I wanted to write this piece for a solo performer alongside the fixed media part. As I worked on the fixed media portion, clarinet seemed like an apt choice given its timbral flexibility. Continuing to follow the layering process used in the fixed media portion of the piece, I completed the electronics before beginning to compose the clarinet part for this piece.

To begin composing the clarinet part, I first reframed the fixed media part. I listened to the fixed media repeatedly and drew a kind of map or score of what I heard. I was not overly analytical about what I heard, rather I drew shapes and wrote words of the first thing that occurred to me. For me, this was a good way of removing the bias that I had as the field recordist of the sounds where it is easy for me to fall into hearing where the sound came from rather than the sound itself. With this map of the piece in-hand, the next step was to compose the piece itself.

The process of composing the clarinet part was drastically different than that for the fixed media portion of the piece. Where the recorded sound was rooted in the acoustic survey data and recordings, the clarinet part began from personal reflections and experiences. I began the piece while staying in an isolated cabin in Hornafjörður on the southeastern portion of Iceland. I had been trying to begin the piece multiple times while in the capital, Reykjavik, but had never really gotten started. Perhaps it was the quiet of the cabin, perhaps the four hours I had spent shivering at the terminal face of a glacier early that day, but the piece began to flow.

One central idea emerged as the connective tissue in the clarinet part: rise and fall. The idea of rise and fall permeates the piece from the opening [0,2,5] motif to the rate of change throughout the piece. The piece is a loose theme and variations with multiple layers developing alongside each other. Some musical ideas will develop slowly—for instance, the balance of air sound to pitched sound which slowly varies throughout the piece. Where others will change instantaneously, such as the introduction of the spectral glissando in bar 14. Here I will examine some of these overlapping development strands and how they interact.

The opening motif, in addition to introducing the concept of rise and fall, serves two other purposes. It first removes the piece from 12-tone equal temperament with a quarter-tone glissando and trill serving to introduce the sound world of the piece. Second, it emphasizes the importance of timbre in the piece: subtle gradations of trill speed, especially on the timbral trills, serve to emphasize the importance of small changes. In m. 7, the second statement of the opening motif maintains the [0,2,5] set but changes the order of the notes appearance and uses octave displacement. The phrase ends with an exceedingly slow timbral trill.

The second phrase and first drastic change of the piece begins in m. 13. The tempo almost doubles and the spectral glissando in m. 14 heralds the arrival of a new sound world. But echoes of the first phrase hold the piece back; this echo in mm. 16-17 is all the more impactful because the live electronics are still echoing iterations of the first phrase at this point. This blurring between phrases is again accented as uneven rhythms are introduced in m. 18. Again, the delay lines used in the live electronics add a layer of complexity as the rhythms of the live clarinet mix with the echoed recording.

Following the second phrase, a new idea emerges in m. 25. Here a low repeated C# dominates for an entire bar, beginning to clear out some of the echoed clarinet sounds from the preceding bars. Air sounds also make their first appearance in this phrase. Throughout the piece, air sounds are used more often when wind noise is dominant in the fixed media track. Often, the air sounds are barely audible, however they contribute to the overall timbre of the piece; and because they are recorded into the delay network, air sounds will sometimes emerge from the texture, even some time after they were initially played. Mm. 27-40 are a slow transformation to total air sound. During this transformation echoes of the opening motif return in high timbral trill.

The fourth section of the piece returns to a focus on rise and fall. This section, from mm. 40-55, again returns to the opening motif, though now it has expanded its compass and verges on becoming a new motif all together. The rest of the section freely combines elements from all three preceding sections: lower repeating rhythms from the third section, spectral glissandi and off-beat rhythms from the second, and timbral trills from the first. This loose development marks a low point in the energy of the piece and mirrors one of the quieter parts in the fixed media track.

From this point until m. 87, the piece picks up energy as the tempo continually increases. Multiple metric modulations serve both to increase the speed of the piece and to further disconnect it form any clear metric grid. This entire section is compound melody with the rise and fall motifs interrupting the low repeated rhythmic motifs. This section also slowly removes all air sound, making it one of the most driving and forceful sections of the piece.

The piece abruptly slows at m. 88, but the two measure long spectral glissando serves to continue the forward momentum. This spectral glissando is especially impactful with the electronic processing causing it to become a cluster of moving harmonics rather than a single strident one. This extended climax of the piece inverts the rise and fall motif, becoming a fall and rise figure. Accented marcato notes in mm. 103-105 help to cut through one of the loudest parts of the fixed media track. Motifs freely mix such as the rise and fall motif combining with the rhythmic motif, until the energy of the piece abruptly dissipates.

As echoes of the preceding louder material continues to wind through the delay network, the clarinet plays some of the quietest material yet to conclude the piece. Air sounds and piano melodies wind over the most unprocessed portion of the fixed media track as the visuals change to unedited footage for the first time in the entire piece. The temporal interaction between the clarinet, fixed media, and electronics is left loosely defined: the clarinet is to begin at the same time as the fixed media, and since the clarinet score is metered, they theoretically should line up at the end. But to account for possible idiosyncrasies of performance an extended fade out ends the work and the piece is designed to function regardless of whether the final sound is clarinet or fixed media.

8. Electronic Processing

The combination of clarinet with the fixed media detailed above posed multiple compositional and technical problems. Not least of which was balancing the sounds for live performance: to address these issues, multiple live electronic processes were used alongside amplification of the dry clarinet signal. These focused on bringing the clarinet sound into the same sound world as the fixed media, but they also reflected the concepts of layering and feedback that underlay the piece as a whole. In addition, the use of a microphone for processing also allowed some dry clarinet sound to be added to the mix to further bring the sounds together.

The effects used on the clarinet sound include convolution reverb, delay, and feedback. The same convolution reverb that was used on the foundational sound layer is used on the clarinet, bringing it into a similar space as the fixed media track. The clarinet sound is also sent through a delay network with delay times ranging from 30 ms to more than 2 seconds. The delays feedback on themselves and into each other creating a complex network of echoes that makes it difficult at times to determine which sounds are original and which are coming from the delay. These delay times were adjusted specifically to facilitate the spectral glissandi used in the clarinet part, transforming the harsh overtones to more continuous spectra.

9. Visuals

There is a final performance element of video that accompanies the performance of the piece: a layered video using footage from the survey sites. In order to emphasize the layering form of the piece, the visuals are constructed by layering multiple videos from the field recording sites using a variety of color filtering and overlap methods.

The form of the video mirrors that of the piece with the density and flow of the piece varying over the course of the work. At any time, there are between two to five sets of footage that are layered on top of each other. One of the more common layering techniques used is to use one set of footage as a negative color filter over the other footage (see figure 6). This has the effect of emphasizing color difference between the two locations and highlighting motion in the footage. Hard color mixing is also used, which, when combined with footage with a single predominant color, creates dramatic monochromatic vistas (see figure 7).

FIGURE 6 – Still from *Ecosystem [512]*, Negative Color Filtering

Source: Author (2024)

FIGURE 7 – Still from *Ecosystem [512]*, Hard Color Mixing

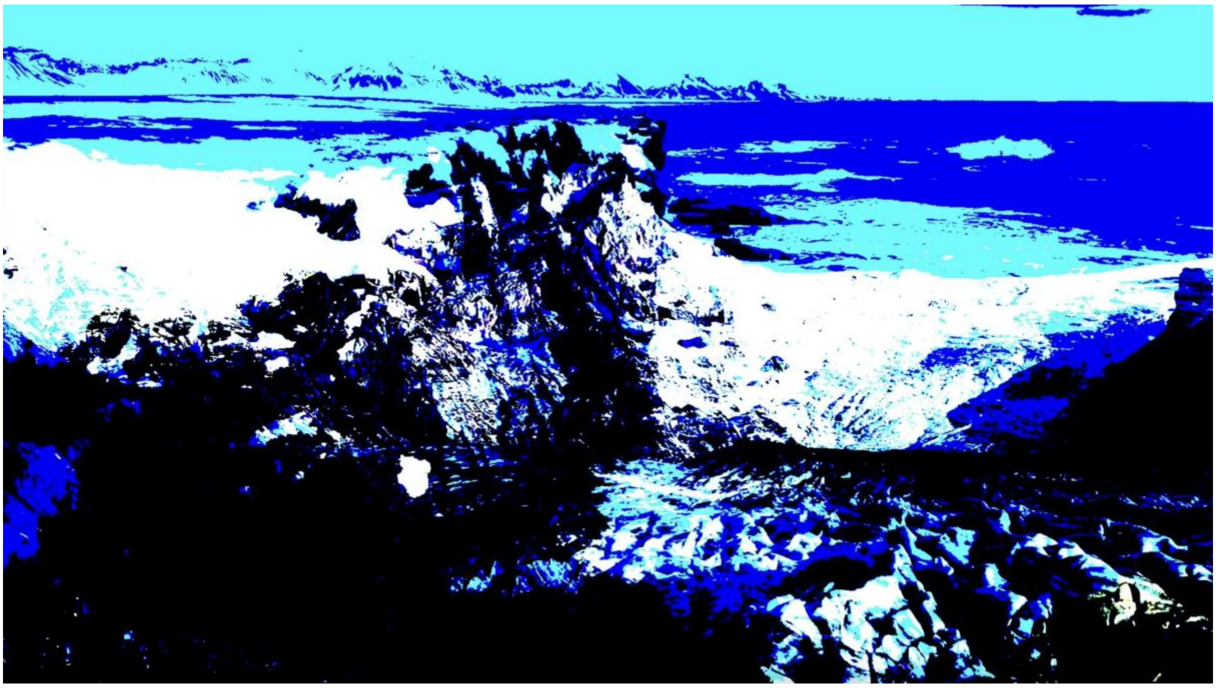

Source: Author (2024)

Matching the slow fade out of the fixed media and the quiet ending of the clarinet part, the visuals end with an elongated fade out. For the first time in the piece, the visuals use no processing or layering effects. Instead, a shot of Snæfellsjökull covered in fog ends the piece (see figure 8). The only indication of motion is the subtle wavering of the grass in the foreground and the slow motion of the clouds in the sky.

FIGURE 8 – Still from *Ecosystem [512],* Final Shot of Snæfellsjökull

Source: Author (2024)

10. Reflection

This project was an exercise in how acoustic ecology and music composition can co-exist and mutually benefit each other. In some instances, art may be a worthy vessel for conveying scientific ideas through pathways such as data sonification. I believe that art and science both have much to learn from each other. The composition of this piece revealed sound details that were not evident from an initial scientific analysis of recordings. And further, the process of composing this piece fundamentally changed my thoughts about the ecosystems I was studying and the goal of the research. While this project and piece were fundamentally personal, I believe that this piece points the way toward one future approach to combining acoustic ecology and music composition.

REFERENCES

AIDE, T. Mitchell; CORRADA-BRAVO, Carlos; CAMPOS-CERQUEIRA, Marconi; MILAN, Carlos; VEGA, Giovany; ALVAREZ, Rafael. Real-Time Bioacoustics Monitoring and Automated Species Identification. PeerJ, v. 1, p. e103, 16 jul. 2013. Available at: https://doi.org/10.7717/peerj.103. Accessed on: Dec. 1, 2024.

BOELMAN, N. T.; ASNER, G. P.; HART, P. J.; MARTIN, R. E. Multi-Trophic Invasion Resistance in Hawaii: Bioacoustics, Field Surveys, and Airborne Remote Sensing. **Ecological Applications**, v. 17, p. 2137–2144, 2007.

BRADFER-LAWRENCE, Tom; DESJONQUERES, Camille; ELDRIDGE, Alice; JOHNSTON, Alison; METCALF, Oliver. Using Acoustic Indices in Ecology: Guidance on Study Design, Analyses and Interpretation. **Methods in Ecology and Evolution**, v. 14, n. 9, p. 2192–2204, 2023. Available at: https://doi.org/10.1111/2041-210X.14194. Accessed on: Dec. 1, 2024.

FARINA, Almo. **Soundscape Ecology: Principles, Patterns, Methods and Applications.** Springer Science & Business Media, 2013.

HILL, Andrew P.; PRINCE, Peter; PIÑA COVARRUBIAS, Evelyn; DONCASTER, C. Patrick; SNADDON, Jake L.; ROGERS, Alex. AudioMoth: evaluation of a smart open acoustic device for monitoring biodiversity and the environment. **Methods in Ecology and Evolution**, v. 9, n. 5, p. 1199–1211, maio 2018. Available at: https://doi.org/10.1111/2041-210X.12955. Accessed on: Dec. 1, 2024.

KASTEN, Eric P.; GAGE, Stuart H.; FOX, Jordan; JOO, Wooyeong. The Remote Environmental Assessment Laboratory's Acoustic Library: an archive for studying soundscape ecology. **Ecological Informatics**, v. 12, p. 50–67, 2012. Available at: https://doi.org/10.1016/j.ecoinf.2012.08.001. Accessed on: Dec. 1, 2024.

KRAUSE, Bernie. **The Great Animal Orchestra: Finding the Origins of Music in the World's Wild Places.** Little, Brown, 2012.

KRAUSE, Bernie. The niche hypothesis. **The Soundscape Newsletter**, v. 6, p. 3-6, 1993.

MALT, Mikhail; JOURDAN, Emmanuel. **Zsa.Descriptors: a library for real-time descriptors analysis.** 5th Sound and Music Computing Conference, Berlin, Germany, jul. 2008. Berlin, Germany. p. 134-137. Available at: https://hal.archives-ouvertes.fr/hal-01580326. Accessed on: Dec. 1, 2024.

MERCHANT, Nathan D.; FRISTRUP, Kurt M.; JOHNSON, Mark P.; TYACK, Peter L.; WITT, Matthew J.; BLONDEL, Philippe; PARKS, Susan E. Measuring Acoustic Habitats.

Methods in Ecology and Evolution, v. 6, n. 3, p. 257-265, mar. 2015. Available at: https://doi.org/10.1111/2041-210X.12330. Accessed on: Dec. 1, 2024.

PIJANOWSKI, Bryan C.; VILLANUEVA-RIVERA, Luis J.; DUMYAHN, Sarah L.; FARINA, Almo; KRAUSE, Bernie L.; NAPOLETANO, Brian M.; GAGE, Stuart H.; PIERETTI, Nadia. Soundscape ecology: the science of sound in the landscape. **BioScience**, v. 61, n. 3, p. 203-216, 2011.

SCHAFER, R. Murray. **The Soundscape.** Vermont: Destiny Books, 1994.

SCHAFER, R. Murray. **The Tuning of the World.** New York: A.A. Knopf, 1977.

TRUAX, Barry. **Composing with Time-Shifted Environmental Sound.** Leonardo Music Journal, v. 2, n. 1, p. 37–40, 1992.

TRUAX, Barry. **Handbook for Acoustic Ecology.** 1999.

US DEPARTMENT OF THE INTERIOR. National Park Service. **Standard Operating Procedure (SOP) 4: Data Analysis Methods Version 1.0, Vehicle Management Plan, Denali National Park and Preserve.** By William C. Clark and Taylor A. Bracher. Dez. 2018.

WESTERKAMP, Hildegard. **Into India: A Composer's Journey.** 2002.

WESTERKAMP, Hildegard. Linking Soundscape Composition and Acoustic Ecology. **Organised Sound**, v. 7, n. 1, p. 51–56, 2002.

ABOUT THE AUTHORS

Garrison Gerard is a composer of electroacoustic and concert music. His work explores intersections between ecology and music. His music has been presented internationally with performances by groups such as [Mod]ular Ensemble, Fort Worth Symphony, and Nu Atmospheres Ensemble. An ardent collaborator, he has been commissioned by ensembles and soloists such as Andrew Cook, Spencer Byrd, the Avenue C Project, Atelier Piano Quartet, and Amorsima String Trio. In 2023 Gerard served as Fulbright Fellow at the University of Iceland. Gerard currently serves as Assistant Professor of Music at the University of North Carolina at Pembroke. ORCID: https://orcid.org/0000-0002-2790-5243. E-mail: Garrison.Gerard@gmail.com

APPENDIX

Link to Ecosystem [512]: https://youtu.be/woy69uuEZLM

Ecosystem [512]

for Brooke Miller

Rev Vórtex, Curitiba, v.12, p. 1-27, e9521, 2024. ISSN 2317–9937.
https://creativecommons.org/licenses/by/4.0/ | https://doi.org/10.33871/vortex.2024.12.9521


Ecosystem [512]
3
64
ff
p
f
mf
fp
f
67
D
tr
mf
f
ff
71
E
fp
f
p
mf
75
(♩ = c. 75)
f
78
F
p
mf
pp
f
81
(𝅗𝅥 = c. 50)
pp
mf
85
Suddenly Slower (𝅗𝅥 = 62)
p
f
p
f

Rev Vórtex, Curitiba, v.12, p. 1-27, e9521, 2024. ISSN 2317–9937.
https://creativecommons.org/licenses/by/4.0/ | https://doi.org/10.33871/vortex.2024.12.9521