UNIVERSITY OF CALGARY

Through a Window: A Networked Music Composition for Four to Six Instruments and Electronics

by

Naithan Shane Bosse

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE

DEGREE OF DOCTOR OF PHILOSOPHY IN COMPOSITION

GRADUATE PROGRAM IN MUSIC

CALGARY, ALBERTA

June, 2018

Abstract

Through a Window is a three movement networked composition for four to six variable instruments and electronics. In this work, musicians are distributed between three performance sites and connected by sending multi-channel audio streams across a high-bandwidth network. The composition explores how the networked setting together with live sound-processing, soundfile playback, amplification, and spatialization creates a unique sonic performance environment. By applying sound-processing differently at each location, the composition creates dynamic configurations in which elements of the work such as harmony and orchestration are perceived differently at each location.

As the physical distance between performance sites increases, the time delay required to send audio between the sites also increases. This poses a significant challenge for music performance since tight synchronization becomes impossible. In *Through a Window* I employ several practical strategies to accommodate ensemble performance in the presence of network latency such as composing audible cues within the music, adopting proportional notation, and using networked stopwatches.

Formally, the composition presents a variety of musical processes based upon evolutionary algorithms, recursive algorithms, and swarm algorithms. These processes occupy multiple sections throughout the composition and contribute to aspects of the composition such as the harmony and melodic contours. As the composition unfolds, these algorithmic materials recur in new configurations and contribute to create an interlocking macrostructure in which the musical tension increases and recedes in a pattern inspired by ocean waves.
Keywords: networked music performance, computer-assisted composition, algorithmic music, interactive electronics, interval cycles, telemedia, distributed performers, chamber music

Acknowledgements

Over the course of this project, several individuals have helped me to develop my creative practice and musical knowledge. I would first like to thank my supervisor, Dr. David Eagle, for his support, enthusiasm, and creative advice throughout this project. His input was very helpful in developing the composition to its finished state. I would also like to thank Dr. Ken Fields for many hours of jamming and experimenting over the network as well as both Dr. Laurie Radford and Professor Allan Bell for their insights and guidance in private lessons and courses.

I am grateful for the time, dedication, and musicianship of August Murphy, Edmond Agopian, Chinley Hinacay, Ethan Mitchell, Tim Borton, Rachel Kreyner, Abdullah Soydan, and Melike Ceylan in premièring *Through a Window* and Brian Garbet and Lauro Pecktor de Oliveira for helping in rehearsal. I would like to thank Conor Stuart, Ethan Cayko, and Dr. Eric Bumstead for collaborations, conversations, and resources in the early stages of the program which helped to inform my choices as I began planning this project.

I acknowledge and appreciate the financial support provided by the Social Sciences and Humanities Research Council of Canada and the University of Calgary School of Creative and Performing Arts.

Finally, I want to thank my family for their encouragement and support in my pursuit of a PhD in music. In particular, I would like to thank my partner, Kamila Pelka, who listened to many drafts of *Through a Window* and whose feedback was always valuable.

Table of Contents

List of Tables

List of Figures and Illustrations

Chapter 1 Introduction

Through a Window explores compositional opportunities that arise when performing music over the internet. In this work, musicians located at three different performance sites connect to one another using a high-bandwidth network to send uncompressed multichannel audio between the sites. These high-quality audio connections enable the musicians to perform together as an ensemble despite being in different locations. The work is composed for four to six instruments and electronics and totals approximately 23 minutes in duration.

Networked performance creates new musical possibilities. Over the network, it becomes possible to perform together across multiple time-zones, facilitate cross-cultural collaboration, and encourage new ensembles or societies that span the globe. In particular, the sense of closeness and the clarity of sound that was created through close microphone placements and multi-channel sound systems was captivating both while composing and performing *Through a Window*. Performances across great distances, spanning cities, countries, or continents ironically encourage an intimate and immersive sound quality.

Networked performance across multiple locations also inherently involves multiple perspectives. Differences among the three sites such as the local technology, speakers, room-sizes, acoustics, number of local performers, and audience size all contribute to a listening experience that is unique to each location. The delay in sending audio between the sites creates unique temporal alignments at each site. In *Through a Window,* using sound processing in different ways at each location helps to emphasize the differences between the performance sites. The listener experiences one possible realization of the work among multiple co-existing possibilities. The sound processing and analysis tools create scenarios in which musical parameters such as harmony, texture, and orchestration are realized uniquely at each location.

This approach is only possible in environments where sound is not shared acoustically between the performers but can instead be manipulated by the computer before being heard at a remote location. A performer may act as a soloist in one location and when processed, act as a spectral ambience in another.

Through a Window is inspired by an excerpt from Jane Goodall in her book *Through a Window: My Thirty Years with the Chimpanzees of Gombe.*

> "There are many windows through
> which we can look out into the
> world, searching for meaning …
>
> … Most of us, when we ponder on the
> meaning of our existence,
> peer through but one of these
> windows onto the world.
> And even that one is often misted over
> by the breath of our finite humanity.
>
> We clear a tiny peephole and stare through.
>
> No wonder we are confused by the
> tiny fraction of a whole that we see.
>
> It is, after all, like trying to
> comprehend the panorama of the
> desert or the sea through
> a rolled-up newspaper." (Goodall qtd. in Parabola.org, 2015)

The three sites are akin to three windows looking out at the same object from three different perspectives. The sound processing tools act like the panes of glass in the windows. The glass may be misted, obfuscating elements of the remote musician's true performance. This mist may slowly dissipate over time and the connection between the sites could become clarified. Different windows may be tinted differently, allowing certain colours of the light through and not others. Perhaps water droplets run along the surface of the glass, bending the light. This might be comparable to spectral delays, which scatter frequency components in time, or pitch shifting which can bend or transpose notes in various directions.

Goodall refers to science as one of the windows through which we can “look out into the world, searching for meaning.” In many sections, I composed musical processes based on evolutionary algorithms, swarm algorithms, and recursion. Evolutionary algorithms and swarm algorithms model processes from the natural world. Recursive functions are also commonly used to generate fractals and self-similar structures that resemble shapes found in nature. Conceptually, these three techniques connect to the excerpt from Goodall. These processes recur in varied forms in several sections throughout the composition. In the case of the recursive processes, the recurrences are snapshots of a greater underlying process. This is a musical expression of trying to comprehend Jane Goodall’s “panorama of the desert or the sea through a rolled-up newspaper.” It is a fragment that suggests the shape of the whole.

1.1 Overview of the movements

Through a Window consists of three movements entitled “Stained Glass,” “In Strange Lines and Distances,” and “A Twisted Pair.” Movement 1, “Stained Glass,” focuses directly on the decorrelation between distance and intimacy that is often created in networked environments through the use of close microphone placement, dry acoustics, and multichannel speaker arrangements. Even though the performers may be kilometers apart from one another, the sound quality is unaltered by this distance. The performers will sound as if they are close to the listeners in all three locations. In performance, I use different microphone lines for each instrument and position the microphones close to the instruments to create an intimate sound quality between the sites. Delicate sounds such as breath or small objects falling in a bowl are at times greatly magnified. By juxtaposing these intimate acoustic sounds with similar electroacoustic soundfiles, I create ambiguity as to the nature of the sound sources between performance sites. If the listeners are unable to see the sound sources, extended techniques or unusual instruments, such as deflating party balloons, become more mysterious. As the movement progresses and develops, the audio is routed through pitch-shifting modules to create sonorities that are unique to each location.

In order to accommodate the variable instrumentation, I chose to notate extended techniques by describing desired sounds directly rather than describing instrumental techniques used to achieve the sounds in performance. Notational symbols for pitched, unpitched, and partially pitched sounds are used along with suggestions in the performance instructions as to how these sounds may be achieved on different instruments. The musicians can then interpret these written textures in a manner that is appropriate for their instruments. If their chosen techniques are not naturally well-balanced with the rest of the ensemble, the computer performers can plan to adjust the mix to help create a more balanced texture. I explore this further in section 2.2.

Movement 2 begins to challenge the ensemble's ability to coordinate while performing over the internet. It often requires the ensemble to maintain metrical alignment and presents more well-defined melodic materials than "Stained Glass." In this movement, I begin to explore a number of practical approaches, including notational strategies, to help the ensemble maintain alignment within the networked context. For example, I frequently include "anchor lines" throughout movements 2 and 3. An anchor line is a melodic or rhythmic line that is included as a cue in each part. The score instructs each performer to synchronize with the anchor lines as closely as possible in their own location, while the anchor performer is responsible for maintaining the tempo despite the possible rhythmic misalignment in their own location. These techniques are discussed in section 2.1.

The title of movement 2, “In Strange Lines and Distances,” is a reference to the often quoted excerpt from Francis Bacon’s *The New Atlantis* from 1627 in which Bacon imagines sounds conveyed along “trunks and pipes in strange lines and distances.”

> We represent small sounds as great and deep; likewise great sounds, extenuate and sharp; we make divers tremblings and warblings of sounds, which in their nature are entire … We also have divers strange and artificial echoes, reflecting the voice many times, and as it were tossing it; and some that give back the voice louder than it came, some shriller and some deeper; yea, some rendering the voice, differing in the letters or articulate sound from that they receive. We have also means to convey sounds in trunks and pipes, in strange lines and distances. (Bacon qtd. in Truax 111)

Several authors refer to this text in the context of electroacoustic music partly because Bacon’s imaginings turned out to be an accurate and poetic description of modern recording studios (Ernst, Truax, Hugill). This is also a fitting description of a networked music composition since sounds are conveyed across much stranger lines and longer distances than any other medium, running between continents along the ocean floor and encompassing the globe.

Movement 2 is much longer than movements 1 or 3, partly as a result of my compositional approach. I did not try to impose preconceived temporal proportions to the movements but instead let the processes, such as the evolutionary and recursive processes that I described above, influence the duration. The material and unfolding of the processes create the form. I discuss this in more depth in chapters 3 and 4.

Movement 3 is the most technically demanding of the movements both in terms of the complexity of the individual parts and the coordination among them. The title of movement 3, “A Twisted Pair,” symbolizes the networked setting as well as its compositional procedures:

first, twisted pair cables represent the network; second, the double helix shape of DNA is a twisted pair that symbolizes the evolutionary algorithm in the movement; and third, two large sections within the movement follow a recursive process that "twists" pitches within a registral span. I discuss these procedures in depth in chapter 3.

1.2 Computer performance

Throughout this document, I use "electronics" as an umbrella term to describe all sounds that enter into and emerge out of the sound system during the performance. This includes audio effects, soundfile playback, synthesis, and spatialization. In the score, I refer to the electronics as computers 1, 2, and 3 and in this document I use "computer performers" and "computer musicians" interchangeably to refer to the musicians running the electronics during the performance.

Through a Window contributes to a rich body of works for acoustic instruments and live or interactive electronics and also contributes to a smaller but growing body of networked music compositions involving electronics. Within this practice, the role and prominence of the electronics varies widely depending on the objectives of the composer. Projects range from simple setups involving an iPhone and two speakers, such as in Hans Tutschku's *Still Air I* for bass clarinet and electronics (2013), to large-scale productions for spatialized soloists, large ensemble, multichannel sound systems, and live sound processing such as in Boulez's *Répons* (1985).

The electronics can take on many roles within these settings. They may create a dialogue with the acoustic performer, acting as a ghostly second performer as in Boulez's *Dialogue de l'Ombre Double* (1984) or create an orchestra-like electronic accompaniment that follows and responds to an acoustic soloist as in Philippe Manoury's *Jupiter* (1989). The electronics can also create an imaginary space within which the performers exist. This is the case in Kaija Saariaho's *Lonh* (1996), where soundfiles and reverberation create a sense of distance that changes throughout the composition (Carnegie Hall, 2012). Electronics might also extend the timbral, registral, or durational possibilities of acoustic instruments through live sound processing as in Scott Wilson's *Flame* (2006) or Jonathan Harvey's *Speakings* (2008).

The electronics in *Through a Window* explore each of these roles. For example, in movement 3 from mm. 41-53, I use sound analysis tools and soundfile playback to create an interactive dialogue between the electronics and performers (see section 4.2). Throughout movement 1 and in movement 2 from mm. 274-331, I use soundfiles, spectral processing, delay effects, and spatialization to create a sense of an imaginary sonic environment. At various points throughout the composition, such as in movement 2 from mm. 227-223 and movement 3 from mm. 38-40, I use pitch-shifting, time-stretching, and delays to intensify the acoustic performance.

In a networked setting, the network itself can create delay effects by looping audio streams back to the source locations. This type of feedback loop can be useful to help performers and listeners gain an intuitive understanding of the distance between the sites, similar to how an echo indicates distance in an acoustic environment. A number of composers have used the network as a source for delay in performance. For example, in *Net:Disturbances* for synthesizers, alto saxophone, and one acoustic network of four channels (2008), Juan-Pablo Cáceres, Alain Renaud, and Justin Yang use network delay as a tool for both sound processing and rhythmic synchronization. Pedro Rebelo also uses the network to create feedback in *Net Rooms: The Long Feedback* (2008). In this work, multiple distributed participants each place a microphone and a speaker close to one another within their location so that their individual room resonances and ambient sounds contribute to a long feedback loop. Rebelo combines these multiple channels in performance to colour the sounds.

Sound processing can also help to create a sensation of a shared acoustic space over the network. For example, in Jonas Braasch's *Virtual Cistern Concert,* Braasch applies a reverberation effect to every instrument in the ensemble. Despite the potentially different acoustical qualities between the performance sites, the reverberation imparts a unified sound quality to the ensemble.

In *Not Being There*, Miller Puckette describes a contrasting approach. In the Global Visual Music group's networked concert "Lemma II", the computers perform an analysis on each musician's performance and send the analysis data between two sites. The data is mapped differently at each location, either controlling pianos via a MIDI interface or other synthesized sounds (410-411). Rather than creating a shared acoustic space, the performance explores multiple realizations (410). In *Through a Window* I also use electronics to explore multiple realizations, mapping soundfile cues to different sample banks and applying sound-processing differently at each location (see section 2.3).

Electronics performers often manage the performance from the center of the hall, where the audio balance and spatialization can be perceived well, and often control the electronics using a computer keyboard and trackpad. Since the offstage location physically separates the electronics performer from the rest of the ensemble, and since, from the audience's perspective, the performer's actions do not always establish a clear relationship between physical gesture and sonic reaction, the electronics performers often take on a background role more in line with that of a musically engaged technician. However, electronics performers may also occupy more focal roles in the performance, especially when interacting with gestural controllers such as iPad accelerometers as used in *Through a Window* (described below). For example, D. Andrew Stewart integrates two electronics instruments as discrete parts in a chamber orchestra in *Catching Air and the Superman* (2010). In this work, two musicians perform with T-Sticks, a digital instrument developed at the Centre for Interdisciplinary Research in Music Media and Technology at McGill University (Malloch et. al, 66). The musicians control the T-Sticks using a variety of controls such as accelerometers and pressure sensors (Malloch et. al, 67). At times, the T-Stick performers use visually striking gestures such as swinging the instruments over their heads. This broad action directs the audience's attention towards the T-Stick performers.

In Kiran Bhumber's *Raula* for Responsive User Body Suit (RUBS) and trumpet (2017), the performer wears a custom-made suit equipped with touch sensors. The digital controller becomes an improvisatory dance-based musical instrument. In both *Raula* and *Catching Air and the Superman,* the performers are not simply technicians but are rather a visible part of the performance.

In *Through a Window,* the role of the computer musicians varies between foreground performers, improvising soundfile playback in a visually noticeable manner, and background technicians, using buttons and sliders to advance presets or set volume levels. Before the performance, the computer musicians establish audio connections and set volume levels. In performance, each computer musician uses an iPad to control a Max patch via custom-built interfaces in either TouchOSC or the author's MaxComm (Cycling '74; Hexler). The Max patch controls custom-built digital audio effects, soundfile players, pitch-detection tools, stopwatch displays, and text instructions for the computer performers.

In performance, computers 2 and 3 mix the four to six acoustic musicians at their local site and ensure that the digital audio effects do not overpower the acoustic sounds. At certain points throughout the piece, computer 2 and 3 also use touch controls to cue samples from a bank of soundfiles in an improvisatory manner. Figure 1-1 shows the layout used in the première performance.

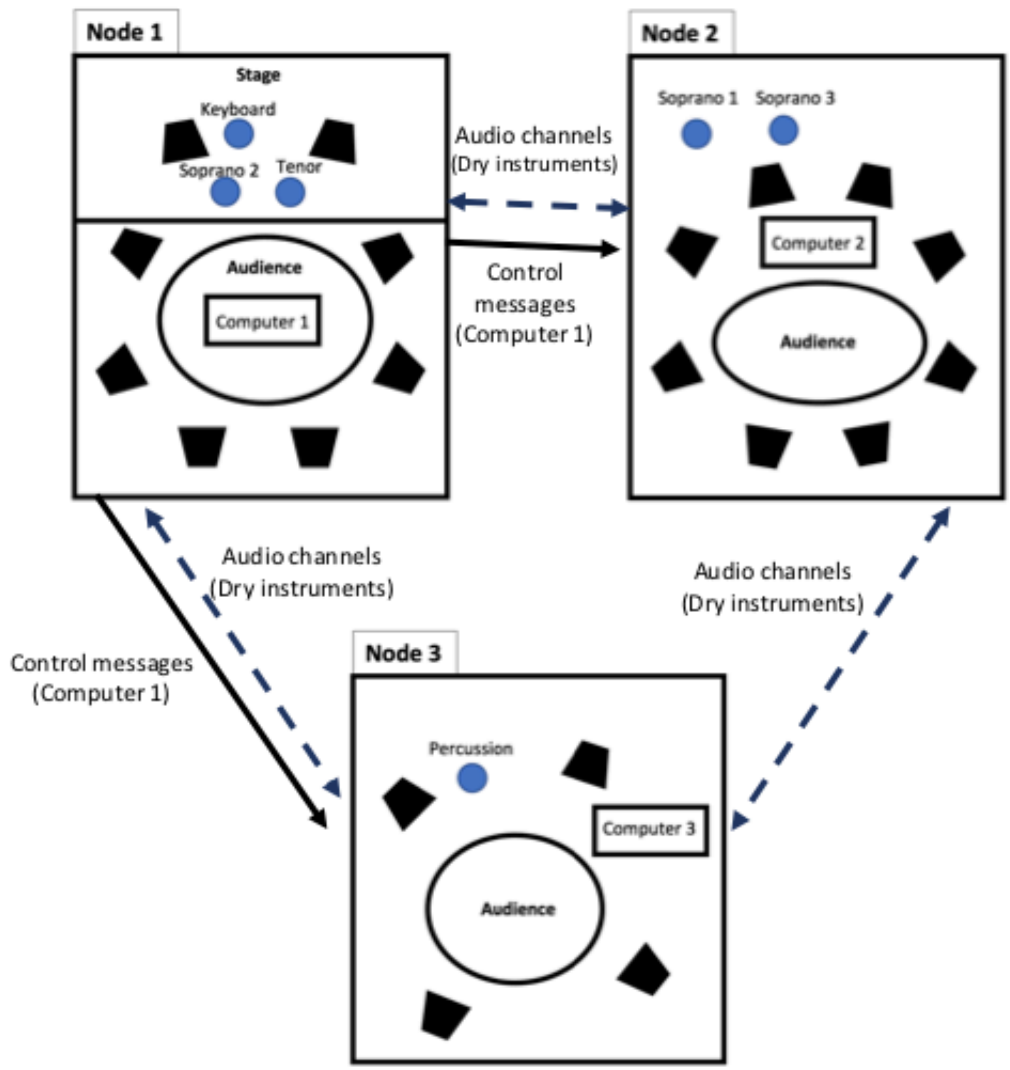

Figure 1-1 Layout for the première performance of *Through a Window*.

Computer 1 has a more demanding role. At the beginning of the work, the computer 1 performer interacts with the iPad's accelerometer and gyroscope controls to improvise with banks of soundfiles. The broad, visually apparent gestures, such as swinging the iPad in various directions to cue soundfiles, help establish a clear correlation between physical gesture and sonic reaction. This approach follows Simon Emmerson's suggestion that adopting a "consistent behaviour", whereby certain kinds of actions consistently produce the same kinds of sonic reactions, can help the audience to perceive the musician's agency in the performance (199-200). Additionally, computer 1 must advance effects presets throughout the composition, adjust spatialization settings, and manually control several digital audio effects. I networked the Max patches so that anytime the computer 1 part advances a preset or starts the stopwatches, the changes will be sent reliably to all locations.

Through a Window was premièred at the Forms of Sound Festival on February 8th, 2018 at the University of Calgary. The performance connected musicians across three buildings on the University of Calgary campus over the campus area network. The majority of the audience members chose to listen in the Eckhardt-Gramatté Hall in the Rozsa Centre; however, audience members could also choose to listen in the Doolittle Studio in Craigie Hall or the Telemedia Lab in the Art Building. Each site was equipped with multi-channel audio systems, individual microphones for each instrument, and stopwatch displays for the musicians. Additionally, two of the sites were connected through a video chatroom. The ensemble consisted of August Murphy, soprano 1 (flute); Edmond Agopian, soprano 2 (violin); Chinley Hinacay, soprano 3 (soprano saxophone); Ethan Mitchell, tenor (cello); Tim Borton, percussion; Rachel Kreyner, keyboard (piano); Naithan Bosse, computer 1; Melike Ceylan, computer 2; and Abdullah Soydan, computer 3.

1.3 Outline of the document

The following chapters detail the compositional approaches and formal organization of the composition. Chapter 2 introduces the larger context and practice of networked music performance and outlines several compositional strategies and considerations for enabling performance across multiple locations. I provide several examples of how I applied these strategies in *Through a Window* and also include representative examples of networked music projects to provide a wider context for the composition.

Chapter 3 examines the generative processes that I used to compose individual passages of music. As mentioned above, these processes were inspired by evolutionary, swarm, and recursive algorithms. Again, I provide a series of representative musical examples for each of these techniques in order to establish a wider context for the composition.

Chapter 4 expands the analytical scope to consider the composition from the perspective of its macrostructure. I show how sections connect with one another throughout the composition to create a cohesive musical whole. I also include statistical graphs to show how textures change over the course of the composition to create a wave-like formal structure.

The conclusion reflects upon the techniques outlined in the document, Appendix A includes the notated score for *Through a Window*, Appendix B diagrams the complete cycle of triple interval cycles described in section 3.1, and Appendix C contains software for running the electronics in performance and provides audio recordings of the première performance from the Eckhardt-Gramatté Hall and the Doolittle Studio.

Chapter 2 The network as a musical medium

Networked music performance (NMP) uses communications technology as an integral component in the construction or transmission of artistic material (Hajdu et al 2009). Networked projects range widely from purely acoustic performances, to live-coding, interactive art-installations, mobile-apps, and internet-based jam-spaces. Performance environments also vary widely from local performances involving a single room, to widely distributed performances connecting musicians across multiple continents, and mobile performances connecting moving musicians via wireless transmission.

Live acoustic performance over a wide-area network imposes limitations on traditional compositional methods, particularly those that rely on strict synchronization between performers. The inherent latency involved in transmitting data over a network can easily obstruct the ability of a musician to play in sync with other remote performers. This was demonstrated in *Effect of Temporal Separation on Synchronization in Rhythmic Performance* where latency of over 60ms caused a noticeable amount of rhythmic "asymmetry" between the performance sites and synchronization became unmanageable for the performers (Chafe et al 2010, 989). Network latency can easily exceed the 60 ms delay threshold when the locations are sufficiently remote.

Other challenges that arise in networked music performances include coordinating schedules across disparate time-zones, facilitating communication between remote performers during rehearsal, synchronizing in performance without the aid of visual cues, and increasing the prerequisite technical knowledge at each location. Despite these rather severe practical challenges, the network also provides several unique opportunities that are impractical or impossible outside of the networked setting such as real-time access to online databases, remote audience interaction, greater cross-cultural collaboration, anonymous performance, and intra-active digital instruments. The following sections introduce some of the obstacles and opportunities that emerge when adopting computer networks as a vehicle for collaborative musical performance and describe several compositional techniques that were employed in *Through a Window* both to take advantage of the network as a musical medium and to aid in maintaining ensemble accuracy between remote musicians.

2.1 Latency

Music is an extremely time-sensitive art. Even minor delays transmitting audio between distributed musicians can have a pronounced impact on the resulting performance. Chafe et. al show that delays as short as 25 ms can cause a gradual but continuous deceleration (989). The tendency to decelerate in response to latency is likely due to musicians unconsciously compensating for the remote musicians' seemingly late onsets. The musicians simultaneously compensate for one another at all locations, creating a continuous mutual rallentando effect. Once latency surpasses the 60 ms threshold, the temporal alignment is so strongly mismatched between the locations that the performance falls apart entirely (989).

Even if data were to travel at light speed, the distance between the University of Calgary and the Central Conservatory of Music in Beijing (CCOM) still imposes a minimum one-way latency of roughly 28 ms – just above the desired threshold for synchronous performance.[1] In practice, data transmission is significantly slower than light speed; light propagates through fiber optic cable at approximately 2/3rds the speed of light in a vacuum (Oda, 10). Network

---

[1] In *Displaced Soundscapes: A Survey of Network Systems for Music and Sonic Art Creation*, Barbosa suggests using light speed as the upper possible speed limit for data transmission (53). The distance from the University of Calgary to the Central Conservatory of Music in Beijing is 8378 km according to Google Maps. The transmission speed of light in a vacuum is 300 km per millisecond. The total latency is therefore 8378/300 = 27.93 ms.

technology and routing also adds latency – a typical one-directional latency between UCalgary and CCOM on CyberaNet is roughly 97 ms.[2]

Transmitting data from one computer to another is not the only significant factor contributing to the total latency in NMP. Audio buffer sizes, microphone placements and speaker distances also have a significant impact.[3] Table 2.1 estimates the latency contributed by each factor during the première performance at the Forms of Sound Festival at the University of Calgary on February 8th, 2018. Notice, the network transmission time was nearly negligible during the première. However, simply moving from a studio rehearsal space to a larger recital hall imposed enough acoustic delay to impact the performance, especially when the remote performers were projected from speakers at the back of the hall. Providing stage monitors or headphones for the musicians can help solve this issue.

Table 2-1 Significant causes of latency

Causes of Latency in performance	Estimated time (ms)
Sender	
Sound propagation from performers to mics	0-2
Audio buffering	7
Transmission between nodes	2
Receiver	
Audio buffering	7
Sound propagation from speaker to performers/audience	5-50
Total electronic latency:	16
Total acoustic latency:	5-52
Total:	21-68

[2] CyberaNet, Alberta's high-speed research network, is capable of sending data at 100 gigabits per second ("Network" 2016). The CyberaNet connects to the CANARIE Network which serves as the "backbone" for Canada's National Research and Education Network (NREN) ("CANARIE Network"). The CANARIE Network interfaces with many other research networks around the globe such as the Internet2 Network in the USA and CERNET in China ("Global Partners").

[3] See Rottondi et. al for a comprehensive breakdown of the factors contributing to latency in NMP (8830-8833).

Even in entirely acoustic settings, the acceptable levels of latency needed for strict rhythmic performance can be exceeded if musicians stand far apart from one another. Since sound travels approximately 0.343 meters per millisecond, a distance of 8.6 meters (28 ft) between performers in an acoustic environment can also cause a noticeable deceleration (Rottondi et. al, 8824). This distance is easily within the span of a concert hall. It is perhaps useful then to consider latency as a compositional problem as well as a technological problem and to adopt solutions that are already commonplace in purely acoustic settings.

One strategy for coping with latency in an acoustic setting is to establish an objective temporal reference for the musicians such as a conductor, synchronized stopwatch, or click track. Conducting is likely the most common method for providing an objective temporal reference to spatially displaced musicians. Since visual information travels at light speed, a conductor can establish a downbeat that will be perceived simultaneously for all participants regardless of their location. However, conducting is ineffective for strict synchronization over a network because video data is transmitted more slowly than audio data. An alternative approach is to use control data to generate synchronized animated visual references for the musicians. Visual references can take many forms including animated avatars (Oliveros, 434; Puckette, 411), scrolling scores (Hope and Vickery), and simple colour signals (Hamel 2014).[4] Visual references may be solely practical in nature or may be artistically generated and displayed to the audience in performance. They can function simply as visual metronomes or contain more complex performance information.

---

[4] In the première performance for Keith Hamel's *Full Circle* for trombone and electronics, trombonist Jeremy Berkman read the score off of an external monitor. The background colour of the score flashed yellow as a warning just prior to each page turn.

Through a Window uses synchronized stopwatches in several sections throughout the work to provide an objective visual reference for the ensemble (Figure 2-1). The stopwatch displays include large, brightly-coloured, animated sliders to indicate the position within the current system. The stopwatch interface also briefly illuminates at the end of each system to strongly signify the start of the upcoming system. The slider displays help create a visual interface that can be easily perceived by the musicians with only a peripheral glance (similar to a conductor). This is important because looking completely away from the notated score can be disorientating and can cause performers to lose their place on the page. 'Jitter cues' are another visual reference used in *Through a Window* and are described in section 2.5.

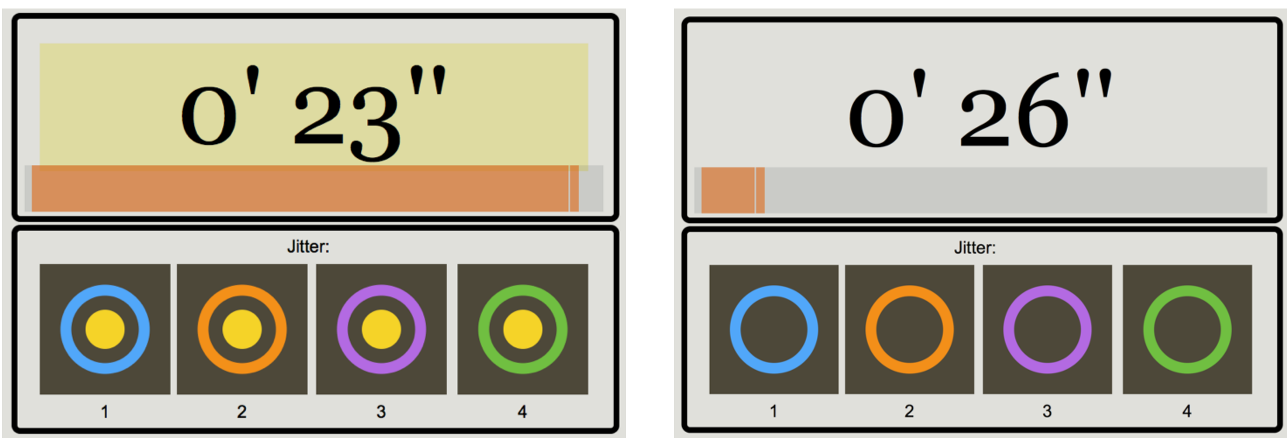

Figure 2-1 Stopwatch interface

Soundfile cues are a form of objective temporal reference that can translate well to a networked medium. Many compositions for acoustic instrument(s) and soundfiles, such as Jonathan Harvey's *Bhakti*, rely on rhythmic cues embedded in the soundfiles to help acoustic performers establish the correct tempo to synchronize with the fixed recorded components (Figure 2-2).

Figure 2-2 Excerpt from Jonathan Harvey's *Bhakti*

Soundfile cues used in this manner function as a musically engaging version of a click-track. As with a click-track, soundfiles can be problematic in that they force the performers to rigidly adhere to a predetermined tempo in order to remain synchronized with the mechanical component. Although this can feel very limiting in performance, the fixed element may be incredibly helpful in NMP to counter the tendency for performers to decelerate in response to latency. Initial drafts of *Through a Window* called for short soundfile excerpts that would establish tempos at the beginning of certain sections. However, I replaced many of these soundfile cues with purely instrumental cues in the finished version for purely aesthetic reasons.

Carôt introduces a third form of objective temporal reference that is only possible in computer-mediated environments dubbed the "Delayed Feedback Approach" (DFA) (Carôt, 5). The DFA adds an artificial delay to the performer's input signal where the delay time matches the network latency. The delayed signal is played back locally so that the local performer can hear how their performance will be aligned in the remote location. The performers must then anticipate their own output similarly to how organists must play ahead of the sounding performance due to the amount of time needed for sound to travel from the organ pipes to the keyboard (Rottondi et. al, 8824). This technique is most effective when using electronic instruments in which there is no acoustic component to the performance to clash with the delayed output (Carôt, 6).

A second common strategy for performing across large distances in an acoustic environment is simply to compose music that doesn't require tight synchronization between the distant musicians. R. Murray Schafer's *Music for Wilderness Lake* for 12 trombones (1979), for example, places trombonists at various points around a lake and requires only a very broad sense of coordination between the parts. Schafer notates durations for individual gestures using either approximate second values or simply specifying that the gestures should last for a full breath. He helps to coordinate the ensemble as a whole directly within the score by indicating that the performers must wait for cues from the other distant musicians before continuing with their own material. The musicians need to actively listen and respond to one another in performance in order for the composition to progress as written. When Schafer does wish the musicians to perform in stricter alignment, conductors cue the trombonists using coloured flags from a raft in the middle of the lake. However, due to the speed of sound and the distance between the musicians, even perfectly synchronized gestures will sound out of alignment from any point on the edge of the lake. Schafer accepts this misalignment as a unique characteristic of the performance setting and uses it intentionally within the composition:

> Sound travels at a little over 330 meters a second, which means that it would take up to 3 seconds for players on one side of the lake to hear anything from the other side. Right from the beginning this had to be a structural feature of the composition so that if I wanted things to sound together at a certain point they had to be written with the appropriate time delays. Simultaneous attacks […] would sound ragged in a deliberate way depending on which notes

> were assigned to the closest and most distant players. (Schafer, ii-iii)

Jordan Nobles has also composed several works for spatialized ensembles. In his open score collection, Nobles presents musical textures in which the ensemble is intentionally loosely coordinated or entirely uncoordinated. In *Lagrange Point* for chamber ensemble (2012), for example, the performers select freely from a pool of short musical cells with very little fixed organization for how the composition should unfold in time. The ensemble focuses entirely on listening and responding to one another within the texture:

> *Lagrange Point* is made of numerous melodic cells. Musicians perform any cell, in any order, in their own time, as expressively as possible, for the durations indicated. There is no synchronization required or desired in this piece. What cells to play, and when to play them, should be decided by each musician independently as long as they are responsive to the other players. […]
>
> Ideally the piece should start with a few quiet long tones played by the sustaining instruments and, after a short time, the various melodic material should be introduced. Likewise it could end with a few of the long tones trailing off and fading out after everything else has stopped. (Nobles, 132)

Scott Deal's *Goldstream Variations,* for variable instrumentation, is another example of a loosely coordinated work. In this case, the composition is specifically intended for performance over the network (2012). Instead of notating fixed pitches in the score, Deal notates pitch contours and allows the musicians to freely interpret their entries and the rhythmic alignment between their parts during the performance. Deal reduces the possibility for the musicians to clash with one another by using a single pitch collection throughout the entire composition and directing the musicians to freely map the collection onto the pitch contours.

In all three compositions, the performance environments and the physical distances involved mean the musicians cannot see one another well enough to rely on physical cues. The success of the compositions depends on the musicians' ability to actively listen and respond to one another solely through sound.

Several passages in *Through a Window* also adopt a loose approach to coordination between the musicians. Movement 1 from 1:12-3:15 and movement 3 from mm. 41-53, for example, both adopt an approach similar to Jordan Nobles's open score compositions, presenting pools of musical cells that the musicians select from freely during performance. In movement 2 from mm. 203-273, and movement 3 from mm. 2-19, and mm. 48-53, the textures are intentionally overlapping, widely spatialized, and tend to outline rising and/or falling contours to create an effect like overlapping waves. In movement 3, mm. 38-40 and mm. 168-181 present dense aleatoric streams of notes to create a texture inspired by swarms of insects and cascades of water droplets. In all of these sections, stopwatches help the ensemble maintain a broad sense of coordination from system to system.

Surprisingly, delays of over 100 ms tend not to affect the resulting performance tempo as severely as delays of 60-100 ms. Rottondi speculates that with such an extreme latency, the musicians begin to simply ignore their remote partners and perform their own part independently (8828). The musicians perform as separate independent entities rather than as a unified ensemble.

A significant difference between latency in a networked environment vs. an acoustic environment is that, in networked settings, large distances do not affect the resulting sound quality. Note onsets remain sharp and rhythmically defined. In purely acoustic environments, sounds are increasingly reverberated, filtered, and attenuated with distance. Since the sound quality remains unchanged when transmitted between network nodes, it is possible to adopt what Carôt refers to as the "Fake Time Approach" (FTA). In FTA, the performance tempo is derived from the network latency so that when a musician plays a note at one node, the note arrives at the remote node delayed by a beat within the performance tempo.[5] The musicians can now perform in sync but with the added requirement of offsetting the metrical position at one of the locations by a single beat which creates rhythmic asymmetry between the nodes (Figure 2-3). Other names for FTA include toporhythm (Cayko, 29), and distributed rhythmic patterns (Cáceres and Renaud, 244). Toporhythms can be created at tempos other than those derived from network latency by measuring the amount of latency and adding an extra artificial delay between the sites to match a beat duration at the desired tempo.[6] Ninjam is an example of an online jamming software in which all performers' inputs are delayed by one full measure at a user-defined tempo to allow interaction over a WAN between residential connections (Cockos).

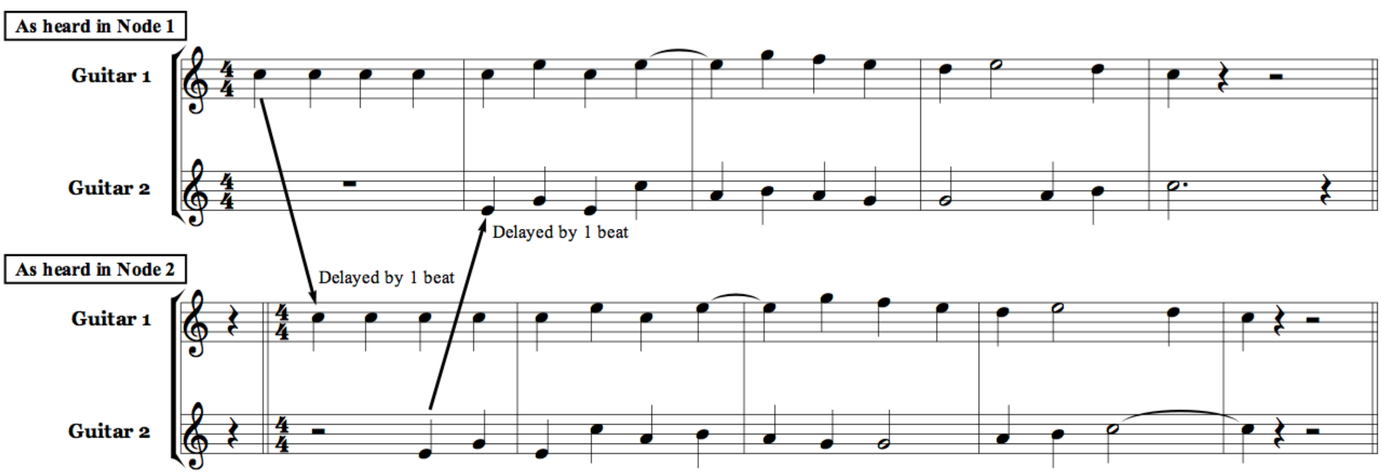

Figure 2-3 Two-node toporhythm. One-directional latency is equivalent to one quarter note. Note the barlines are misaligned to reflect the perceived beat structures in each node.

[5] 60000/latency = BPM
[6] To calculate the additional delay needed to perform at a given tempo: Artificial delay in ms = (60000/BPM)-latency.

In addition to stopwatch cues, *Through a Window* relies on “anchor” lines to help maintain alignment in performance. At specific points throughout the composition a musician will be assigned the role of anchor. The anchor musician is responsible for maintaining the tempo and the remote musicians are instructed to align to the anchor musician's performance in their respective locations. As a result, the anchor musician will perceive the remote musicians as lagging behind the written score. It is the responsibility of the anchor musician to maintain the performance tempo without regard for the lagging remote musicians. This technique is described by Carot as the “Master Slave Approach” (MSA).

Catch-up points are another technique that was used prevalently throughout *Through a Window.* A catch-up point is a fermata or other break placed at the end of a fast or rhythmically complex section to allow the ensemble to realign themselves and perhaps re-establish the tempo after a possible deceleration (Figure 2-4). Catch-up points also help maintain ensemble alignment in controlled aleatoric situations such as Witold Lutoslawski’s *String Quartet* in which phrases often end with rests between sections as well as aural cues.

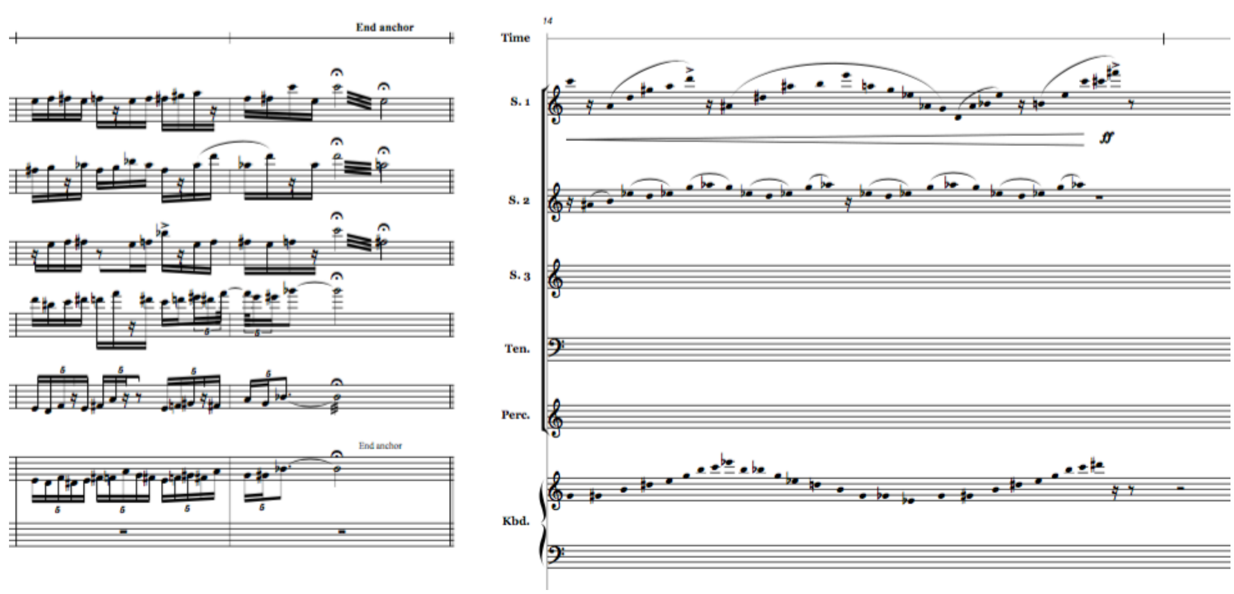

Figure 2-4 Two catch-up points. Movement 2, mm. 168-169 and movement 3, m. 14.

2.2 Presence

In *Tapping into the Internet as an Acoustical/Musical Medium,* Chris Chafe suggests that the network functions as a new sound propagation medium, similar to air or water (2009). However, unlike a traditional acoustic medium, the volume and timbral qualities of sound are not altered by increased physical distance. The sound studios frequently display dry acoustics and microphones are often placed close to the instruments to avoid feedback and to minimize latency. The sense of closeness created through microphone placement creates the opportunity to explore Sivouja-Gunaratnam's concept of an "aural magnifying glass," that she uses to describe Kaija Saariaho's use of electronics in her composition, *Lichtbogen* (2005). In *Lichtbogen*, sounds that are not audible in traditional concert performance such as the scratching of nails on a guitar string or breath through a flute, become magnified simply by amplifying the ensemble.

In *Telematic Composition*, Smith refers to this quality of acoustic closeness in relation to "hyper-presence" (24). The illusion of presence exhibited by the remote performers can be actively controlled to create musical drama. For example, video projections allow the remote performers to dwarf the local performers in terms of physical size and brightness while similarly, amplification allows the remote performers to dwarf the local performers in terms of volume. The remote performers will consequently draw more attention from the audience than the comparatively quieter and physically diminished local performers. Subtle gestures by the remote performers such as key clicks or page turns can potentially over-power a local fortissimo. The remote performers therefore become "hyper-present".

The manipulation of presence can be used to explore orchestrations that are not effective in acoustic performances. For example, a remote fortissimo snare drumroll could be filtered and attenuated live to fit within a texture that would otherwise be overpowered if the drumroll were performed locally. This approach has been used in acoustic performances through physical placement of the musicians within the performance venue. Hope Lee's composition *Imaginary Garden IV* for recorder and trumpet places the trumpet player offstage in part to help maintain a reasonable dynamic balance between the two very different instruments. However, unlike balance achieved through physical spatialization, the volume of remote instruments over a network can be manipulated independently of timbre.

In addition to the practical reasoning described above, placing performers offstage can add an element of mystery or drama to the performance. The audience cannot visually anticipate when the offstage performer will begin to play or whether the performer will change instruments or locations. The ambiguity can be emphasized further in NMP through sound processing, unusual instrumentations, extended techniques, or soundfile playback.

The role of amplification to facilitate unusual orchestrations is especially important in *Through a Window* due to the work's variable instrumentation. The anchor lines and cues, in particular, must be heard clearly by all performers. In the première performance, the soprano 1 part was performed by flute, which was problematic from mm. 123-135 in movement 2 during co-located acoustic rehearsals. This section is written in a soft register for the flute (mainly between C4 and F#4) and it was therefore difficult for the ensemble to follow this line due to it becoming buried under the other instruments. I was able to mix and spatialize the flute during networked rehearsals to mitigate the problem. If this same line were performed on violin instead of on flute, the problem would naturally disappear and the section would be mixed differently.

In addition to 'hyper-presence,' I will use the term "mediated presence" to refer to any networked environment in which the dynamic, purity, or intelligibility of the remote performers is intentionally manipulated to create an orchestration that is different or impossible in a parallel acoustic environment. Subcategories of mediated presence include hyper-presence, hypo-presence, and ambiguous presence.

Hypo-presence, the natural counterpart to hyper-presence, uses the network to attenuate or filter the input of the remote performers. The snare drum example above provides one possible manifestation of hypo-presence. A more complex example could be to strategically mute audio between nodes to create site-specific harmonizations (Figure 2-5). Similarly, muting audio between nodes could be used to allow musicians to perform at slightly varied tempi without influencing one another. A third node could receive audio from the two disconnected nodes to hear the remote performers drift out of phase with one another. Another example is the use of reverb or spectral delay on the remote inputs to create an illusion of distance or blur.

Figure 2-5a An example of using the network to accompany node 1 in both C Major and e minor. Node 2 is muted in nodes 1 and 3. Node 3 is muted in node 2.

Figure 2-5b The composite performance as heard in each node.

Initial drafts of *Through a Window* called for several instances of hypo-presence that were removed in the final draft. In each case, these instances required the audio to be fully muted between the nodes. For example, mm. 52-61 in movement 2 was composed with two possible interpretations in mind depending on the amount of latency. If the latency is sufficiently long, the connections between the nodes are fully muted. Harmonization sound-processing modules are added to both compensate for the thinner texture and to explore node-specific variations of the scored parallel harmonies. This approach was tested in rehearsal and although the resulting texture was sonically effective, muting the connections between the nodes was disorientating for the ensemble and seemed to create concern that a technical problem had occurred. The latency was short enough during the première performance that unmuting the remote nodes during this section created a pleasing slapback echo effect.

Once presence is established in performance, it becomes possible to actively distort the purity of this presence through the addition of sound processing. Ambiguous presence refers to the use of computer processing, recording, or instrumentation to intentionally obscure which sounds are performed live by the acoustic musicians and which sounds are pre-recorded or the result of sound-processing. If no visual contact is established between the local performer and the remote musicians, it can become unclear how many performers are located in the remote locations and which instruments are being played (particularly if extended techniques or unusual instruments are used such as in the early pages of movement 1). Pre-recorded samples of the remote musicians can also be used to create ambiguity as to which sounds are recorded and which sounds are performed live. Pre-recorded sounds combined with reduced visual connectivity can create the illusion of a greatly increased ensemble size perhaps made up of unfamiliar instruments or timbres that extend beyond typical performance constraints such as register or duration.

Projecting the remote nodes from fixed groups of adjacent speakers within the local node helps enforce the illusion of a larger interconnected composite space. As a listener, I begin to identify different physical locations within the local space as belonging to specific remote sites. Once the remote nodes are associated with a static location, simply moving the remote performers to new locations creates ambiguity and destabilizes the sensation of presence. I explored dynamically spatializing the performers during rehearsal and was surprised by how strongly disorientating the effect was within a large room.

Mediated presence was a foundational concept in movement 1, *Stained Glass*. This movement in particular creates a sensation of aural magnification through unconventional instrumentation, extended techniques, sound processing, spatialization, and soundfile manipulation. The movement contrasts the inherent disconnectedness of the networked performance environment with intimate sounds such as breath and speech. The intention is to amplify the breath-like sounds and quiet voices to create a soundscape that is both magnified and uncomfortably intimate – to use amplification to represent everyday objects as larger than life. *Stained Glass* is inspired, in part, by an excerpt examining the invention of David Edward

Hughes's microphone in comte Théodore Achille L. Du Moncel's 1879 book *The Telephone, The Microphone, and the Phonograph* (146) (Figure 2-6). I was attracted to the excerpt because of how closely Du Moncel's observations can describe both NMP as well as the intimate, magnified textures that I wanted to evoke. I extracted short fragments of the text to be read by the performers. Live and recorded deflating balloons create an immersive oceanic soundscape connecting to the "fly's scream" and percussive soundfiles are used to abstractly recreate the "steps of a fly". The magnified elements are an example of hyperpresence. Presence is also challenged from 0:24-1:12 as it becomes unclear which elements are recorded and which are performed live. Particularly in node 1, due to the lack of visual input, the sounds performed through the speakers are partially disconnected from the gestures and materials used to create them. Therefore, both hyper- and hypo-present elements are pursued in the first half of the movement. Ambiguous presence is created in the second half of the movement (2:00-3:10) through the addition of pitch-shifting modules to the remote performers (described below).

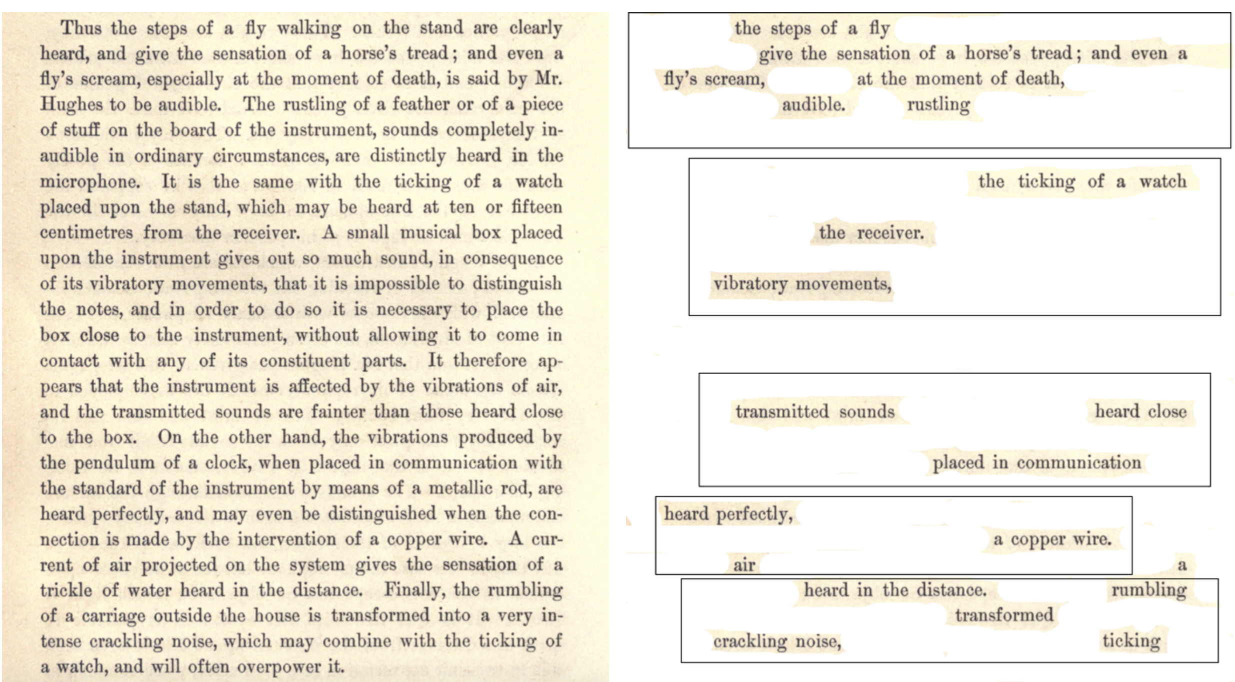

Thus the steps of a fly walking on the stand are clearly heard, and give the sensation of a horse's tread; and even a fly's scream, especially at the moment of death, is said by Mr. Hughes to be audible. The rustling of a feather or of a piece of stuff on the board of the instrument, sounds completely inaudible in ordinary circumstances, are distinctly heard in the microphone. It is the same with the ticking of a watch placed upon the stand, which may be heard at ten or fifteen centimetres from the receiver. A small musical box placed upon the instrument gives out so much sound, in consequence of its vibratory movements, that it is impossible to distinguish the notes, and in order to do so it is necessary to place the box close to the instrument, without allowing it to come in contact with any of its constituent parts. It therefore appears that the instrument is affected by the vibrations of air, and the transmitted sounds are fainter than those heard close to the box. On the other hand, the vibrations produced by the pendulum of a clock, when placed in communication with the standard of the instrument by means of a metallic rod, are heard perfectly, and may even be distinguished when the connection is made by the intervention of a copper wire. A current of air projected on the system gives the sensation of a trickle of water heard in the distance. Finally, the rumbling of a carriage outside the house is transformed into a very intense crackling noise, which may combine with the ticking of a watch, and will often overpower it.

Figure 2-6 Original text (left) and excerpted fragments (right) (Du Moncel, 146)

2.3 Anamorphosis

In visual arts, anamorphosis describes images that are “distorted” or “projected” through space so that viewers are only able to perceive the undistorted images from a specific position or with a viewing implement such as a curved mirror (Figures 2-6 and 2-7) (Clarke). Anamorphic art is naturally experienced from multiple perspectives. As viewers move within a space such as in Figure 2-6, they will see the work as various “broken fragmented shapes” as well as in proper alignment (Varini, 2014). The work was intentionally designed to be viewed in both distorted and undistorted forms.

In *Through a Window*, the network enabled me to explore anamorphosis within the domain of music. Like anamorphic art, NMP also naturally involves multiple realizations. Each location in a networked performance possesses unique characteristics that cannot be shared between the remote sites. For example, the performers, room sizes, acoustics, and audio balances will almost always be different at each location. Rhythmic alignments will also differ if the sites are sufficiently remote.

Toporhythms, described in section 2.1, are conceptually similar to anamorphosis in that both anamorphic artworks and toporhythms rely on the viewer’s or listener’s position in physical space to perceive material in a particular alignment. For instance, note onsets may alternate in one node and occur simultaneously in another. In this case, rhythmic alignment is directly connected to the physical distance through which the sounds are projected.

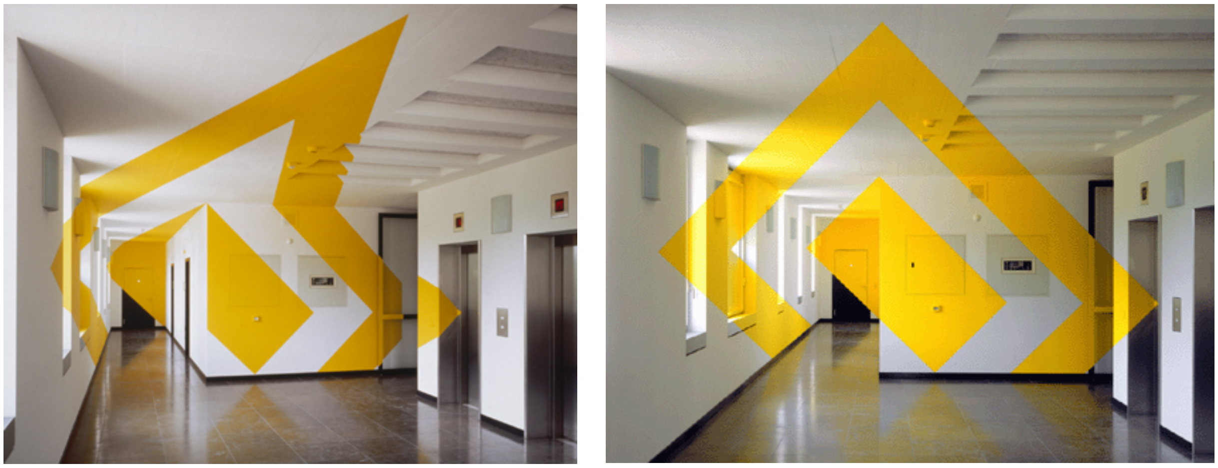

Figure 2-6 Anamorphic art by Felice Varini ("Rettangoli Gialli Concentrici Senza Angoli Al Suolo", 1997)

Digital audio effects can also distort sound to create performances that are realized uniquely at each node. Since the remote performers' entire sonic contribution is experienced solely through the sound-system, the remote performance can be altered by the computer with no parallel unaltered acoustic experience. This characteristic creates the opportunity to expand upon anamorphic rhythms to explore site-specific realizations of harmony, texture, tempo, and orchestration in a manner that would not be possible in a non-networked environment. *Through a Window* does not have a single correct realization, but instead has three different realizations which are interconnected, dependent upon on one another, and realized simultaneously. In this way, anamorphosis emphasizes an inherent characteristic of the network to create multiple realizations. Performances become "distorted projections" in the remote nodes. The use of sound processing modules to facilitate site-specific realizations is akin to the role of a curved mirror in anamorphic art.

Comparing anamorphic rhythms to anamorphic images might suggest that one rhythmic interpretation is prioritized above the others; for instance, one primary node is 'correctly' aligned while the remaining nodes are locally misaligned in order to be heard as correctly aligned in the primary node. However, in *Through a Window,* anamorphosis is used to create three valid variations of a musical idea. No location is considered preferable to another. The approach is closer to anamorphic works by István Orosz in which the distorted projection is entangled in an undistorted image to convey a “double meaning” (Orosz, 1). In Figure 2-7, notice the raven and table in the top-down image (left) becoming Edgar Allan Poe’s face and vest in the reflection of the curved mirror (right).

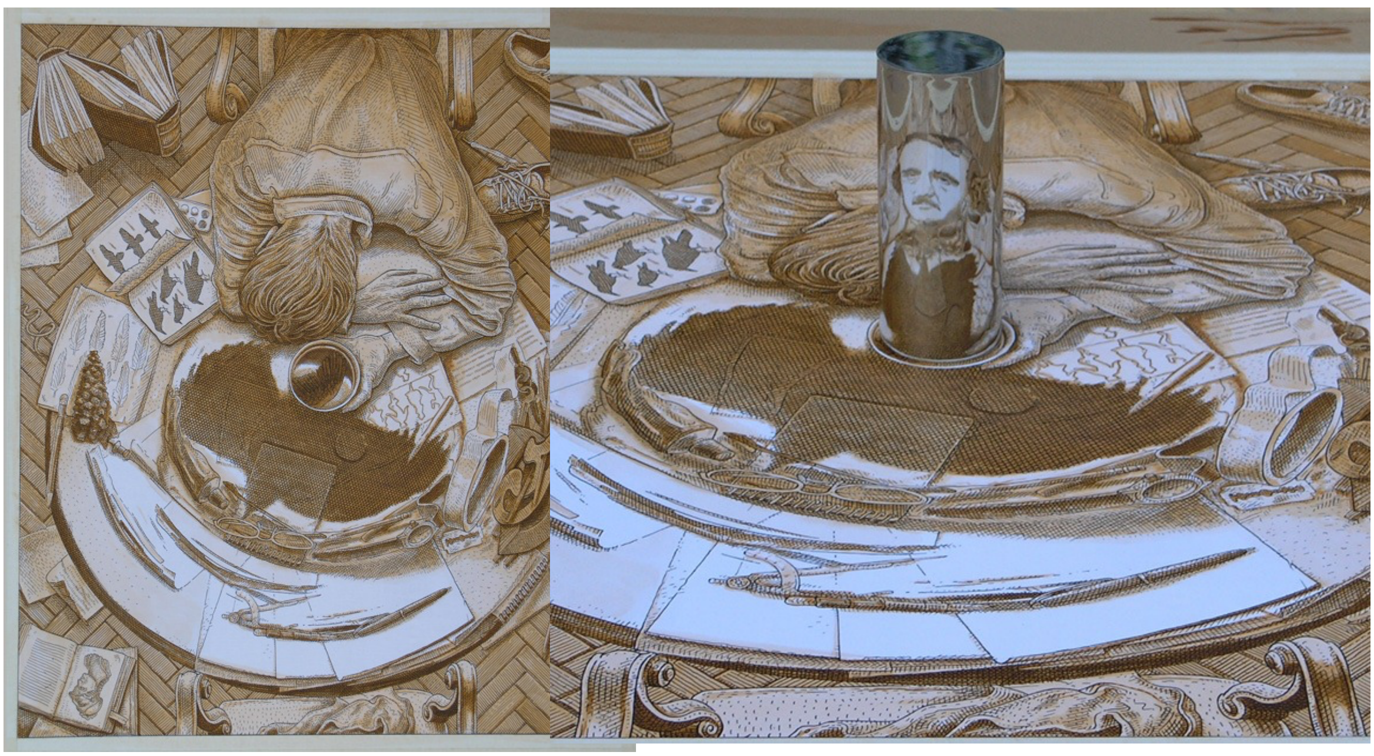

Figure 2-7 An example of mirror anamorphosis István Orosz’s *The Raven* (243-245)

Unlike the visual art examples from Varini and Orosz, in *Through a Window,* I am able to change the nature of the anamorphic distortions from moment to moment throughout the composition simply by changing the parameters of the sound processing modules. The effect is comparable to mirror anamorphosis if the mirror's size and shape was malleable rather than fixed. If the mirror changed shape the image would distort in different ways.

Several sections of *Through a Window* use pitch-shifting modules to create site-specific harmonizations. This is explored prominently in the section from 1:12-3:10 in movement 1. The incoming audio from the remote performers is pitch shifted by different intervals at each location. The nodes each present a slightly different harmonic quality. Figure 2-8 provides a short potential realization of 2:50-2:55 as harmonized at each of the three nodes.

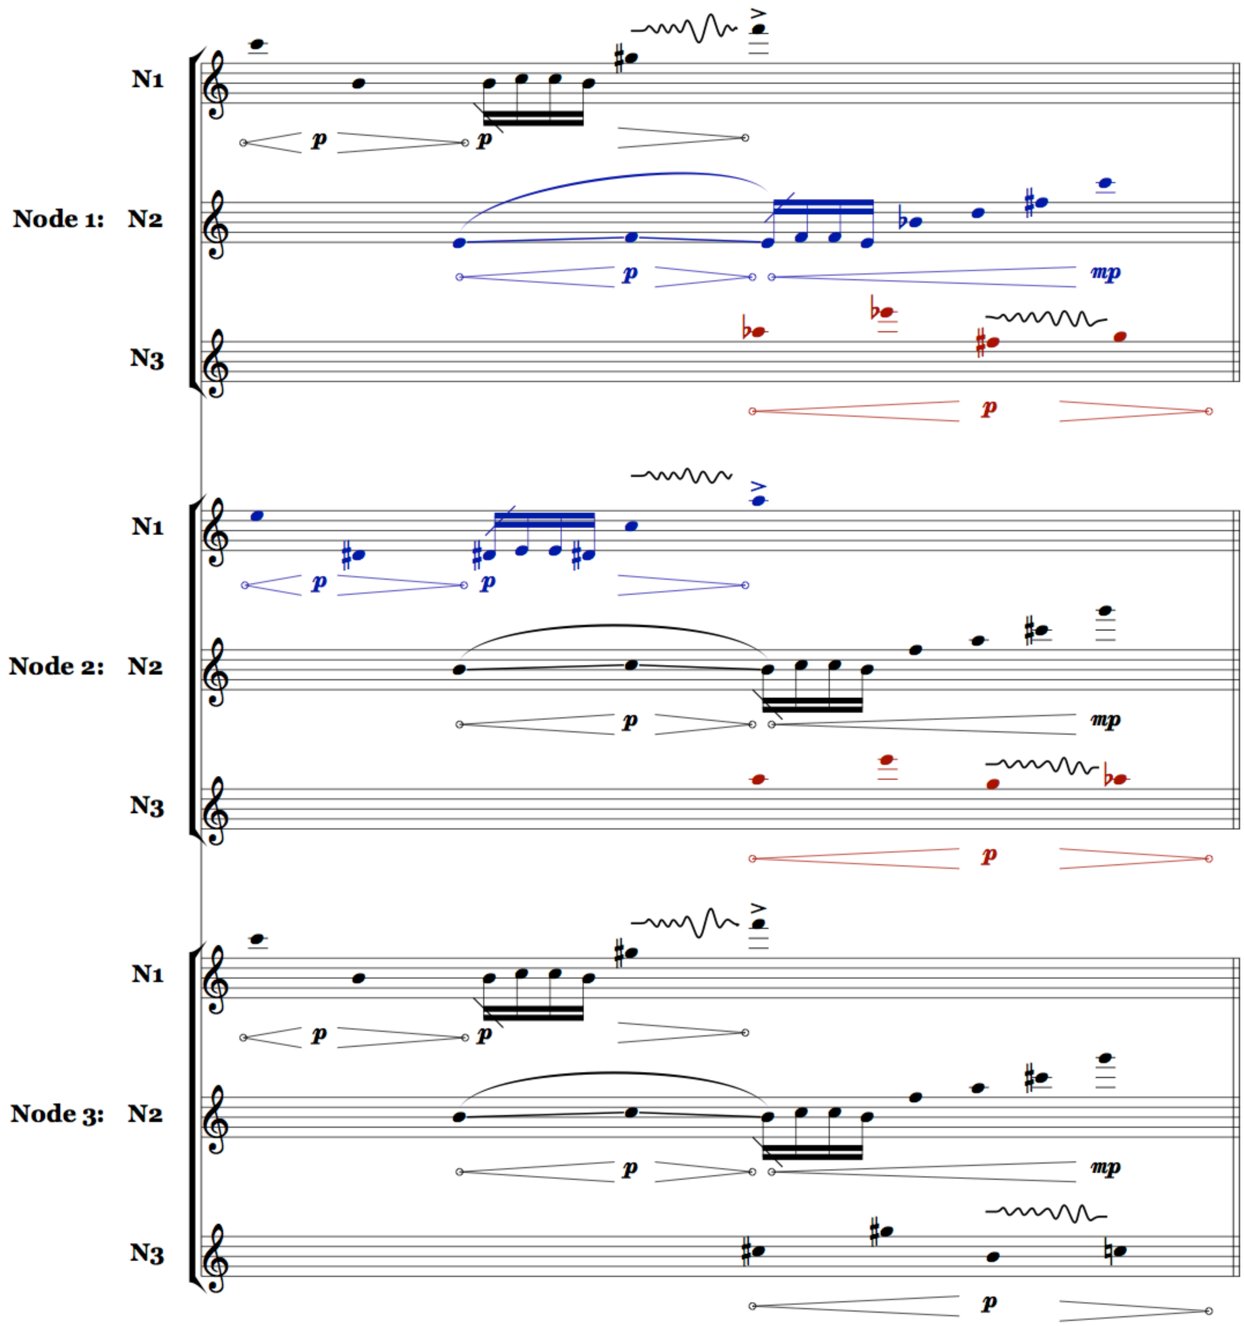

Figure 2-8 One possible realization of *Stained Glass* from 2:50-2:55 as harmonized in 3 nodes. Blue notes are artificially transposed down and red notes are transposed up.

Other examples of site-specific sound-processing can be subtler. For example, the granular synthesis effects in movement 3, *A Twisted Pair*, use slightly different pitch transpositions, grain sizes, and trigger rates at each node. Other sections in *A Twisted Pair* use site-specific delay times and time-stretch factors.

Through a Window also uses site-specific soundfile playback, drawing from several different banks of soundfiles throughout the performance. Each bank is generated using a different sound source. Soundfiles are commonly performed independently at each location and are not shared remotely. The computer performers cue soundfiles using their iPad's touch interface or motion sensors. Additionally, the computers automatically cue soundfiles using sonification data in movement 1 and using pitch analysis modules to respond to instrumental input in movement 3. The specific sample banks heard in each location are selected using the iPad at node 1. Three of the banks are mapped so that they will always be different at each node. If node 1 selects a bank of bicycle samples, then node 2 will use bottle samples and node 3 will use extended piano technique samples. The bank assignments rotate, so that if node 1 selects bottle samples, node 2 receives piano samples, and node 3 receives bike samples.

The performance layout for each location also contributes to a site-specific experience. The different room sizes, acoustics, channel numbers, musician placements, and video projections all contribute to a unique performance experience at each site (Table 2-2). The uniqueness of each location is intentionally pursued and emphasized using the electronics.

Table 2-2 Through a Window, Performance layout

Location	Node 1 (Eckhardt-Gramatté Hall)	Node 2 (Telemedia Lab)	Node 3 (Doolittle Studio)
Room size (sqft)	5555	620	1189
Local musicians	Violin, Cello, Piano	Flute, Soprano Sax	Percussion
No. of Channels	8	8	4
Video projection	No	Yes	Yes

2.4 Affordances of the network

In *Network Musics: Play, Engagement and the Democratization of Performance,* David Kim-Boyle argues that the inherent limitations and aptitudes of a network prompt a fundamental shift in "the role of the composer to that of [a] designer" who crafts a musical system affording performers or participants with some degree of agency within the musical framework (363). In such an approach, music is experienced not through the reception of a linear 'narrative' but rather through active experimentation and engagement with a responsive or interactive system. The experience could be as different from a traditional composition as a video game is from a movie. Both environments can be powerfully engaging and fulfilling to the participant/audience member but the nature of the experience is very different.

Weinberg proposes a taxonomy that highlights the difference between linear, fixed notated compositions and interactive non-linear environments (Figure 2-9). He categorizes projects as "process centered [sic] musical networks" and "structure centered [sic] musical networks." The goal of a process-centered project is not necessarily the "musical outcome" of the system itself but rather the form of interaction or process performed by the participants. This may be an exploratory interaction, such as a virtual reality environment in which actions within the environment result in sonic responses, or goal-oriented interaction, such as an online multi-player musical game.[7] In contrast, structure-centered projects aim to accurately realize a fixed, score-based composition.

---

[7] Note that Weinberg's use of the term "process" is unrelated to process as a musical formal device such as Steve Reich's use of phasing in *Piano Phase* (1967). *Piano Phase* would be classified as a structure centered project if performed over a network.

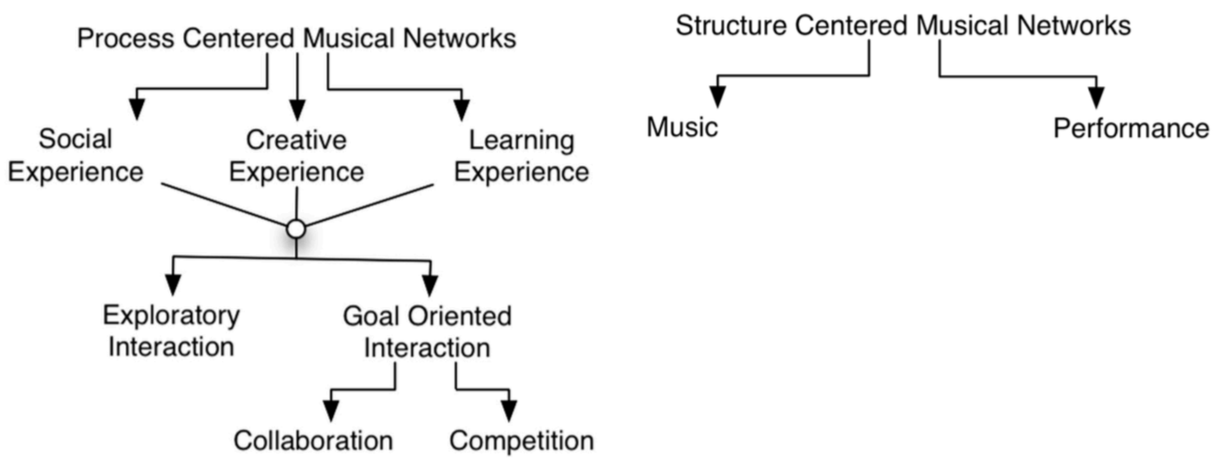

Figure 2-9 Process-centered and structure-centered musical networks. (Weinberg, 32)

Although Weinberg's two categories are not intended to be mutually exclusive, designing a project that fits within both paradigms is challenging. Process-centered projects require the composer to relinquish some amount of control over the musical outcome in order to provide agency to the performers or the participants. Consequently, these projects tend to undermine structure centered approaches in which the music itself is the ultimate goal. The author's *SoundScavenger*, a networked soundwalk app for iOS devices, is an example of a project that combines both process and structure. In this work, the user is placed on a map segmented into seven different zones (2016). The phone tracks the user's movements using its built-in location services. As the user wanders from one zone to another, the soundscape will change to reflect their current location. Structure is achieved by composing each zone to progress linearly as a parallel complement to the others. Although the foreground materials vary from zone to zone, there is a structural trajectory that is present throughout all zones and is created through the manipulation of broad musical characteristics such as the number of layers, attack density, overall dynamic, and sound source.

Similarly to Kim-Boyle, Kane advocates strongly for adopting the natural affordances of the network as structural elements in net-music. By using these affordances as the foundations for musical interaction or structure, the composer can develop approaches that are uniquely suited to the medium itself rather than approaching NMP as a substitute for acoustic performance. Kane identifies some examples of network affordances that can contribute to musical design such as anonymity, large-scale collaboration, and real-time access to online databases (4-5). Other possibilities include unlimited time-scales, intra-action, and multiple concurrent realizations (Föllmer 186, Kane 4-5).

Udo Noll's *Radio Aporee* uses both large-scale collaboration and database access to create a global soundscape map. Users from anywhere in the world can contribute to the project by uploading local soundscape recordings to an online database. These recordings appear as geo-tagged nodes on an interactive map of the globe using the Google Maps API (Figure 2-10). Visitors can interact with the map to listen to these soundscapes. The project engages users through active exploration of a virtual environment. Each user will have a unique experience of the project and will likely only perceive a small fraction of the total environment.

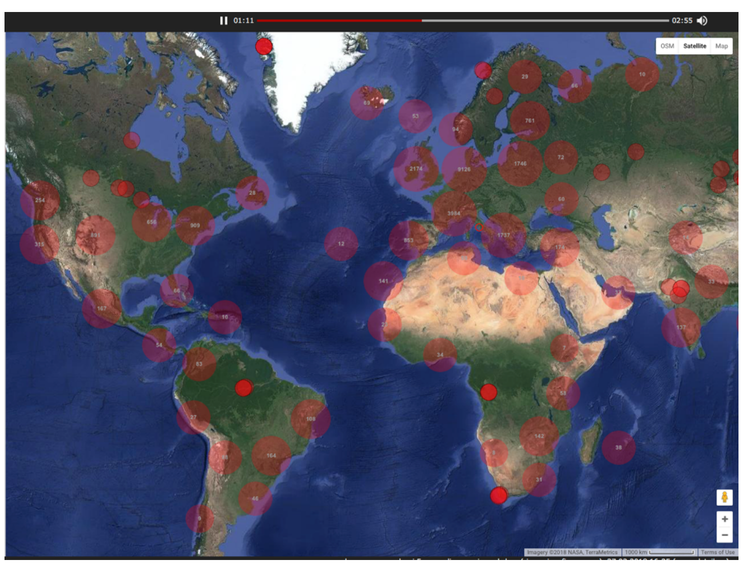

Figure 2-10 Radio aporee map (Noll). Red nodes depict clusters of soundscape recordings.

Networked Miniatures, one of my earlier studies in NMP, uses database-access in a performative context (2014). The work accesses news webpages to randomly select words from the current headlines. These words are displayed and spoken to the audience using video projection and text-to-speech synthesizers in Max (Figure 2-11). The intention is to abstract the words from their context and reveal any underlying emotional quality felt through the choice of language. The specific words vary from day to day as new stories are posted. The emotional tone of the words will also vary from day to day as well as from website to website.

Figure 2-11 A still from *Networked Miniatures* no. 4 (2014), Words are randomly selected from news headlines. Positions and sizes are selected using ping times as a source of semi-random numbers.

Intra-action is another concept that is particularly well-suited to NMP. Golo Föllmer defines intra-action as "processes in musical systems wherein single players' actions intercede into those of the other players, thereby establishing a strong presence" (186). An example of intra-activity includes networked instruments in which changing a parameter such as duration on one instrument also changes a parameter such as timbre on another instrument. The act of performing on intra-active instruments then becomes an experience of mutual feedback with your co-performers.

Wax Lips by Tim Perkis, a member of the pioneering network computer music group the Hub, explores a more complex form of intra-action. Each ensemble member designs a computer system capable of receiving, transforming, performing, and sending MIDI note messages. One performer seeds the system by sending out a MIDI note message to another performer. This next performer plays the incoming note, then transforms the pitch, channel, and velocity values for this note in some fixed, regular manner, and then forwards the transformed MIDI note to another performer who repeats the process using their own set of transformations. The transformations are not coordinated between the performers in advance; however, each transformation must remain fixed so that a given input value will always generate the same output value. According to Perkis, the goal is to “see if there is any emergent pattern to a static setup, where each station acts in a fixed, predictable way, but the interconnects are so complex that the overall behavior is still groovy” (Perkis). Members of the Hub argue that these complex intra-active systems create a “mind-like aspect” that would not be present if the systems were not interconnected (Bischoff et. al “Music for an interactive network of microcomputers” 28).

> The network action had an unexpected living and liquid behaviour: the number of possible interactions is astronomical in scale, and the evolution of the network is always different, sometimes terminating in complex (chaotic) states including near repetitions, sometimes ending in simple loops, repeated notes, or just dying out altogether. (qtd. in Gresham-Lancaster 43)

The “liquid behaviour” that Tim Perkis describes in relation to *Wax Lips* may be partly attributed to the limitations of the network technology. The network used by the Hub was fallible and MIDI messages did not always arrive at the destination as intended. Since each player simply responded to the output of the previous player, a single missed note message would stop an entire layer of sound.

> The main problem was one of plugging leaks: if one player missed some note requests and didn't send anything when he should, the notes would all trickle out. Different rule sets seem to have different degrees of 'leakiness', due to imperfect behaviour of the network. (qtd. in Gresham-Lancaster 43)

If this were not true, any message seeded into the system would continue infinitely. Instead, multiple note layers are created and die away one-by-one until more messages are seeded into the system.

I explored intra-action in early drafts of *Through A Window*. I initially notated movement 1 from 1:12-3:10 using pitch contours rather than specific pitches and the musicians selected from a changing set of pitches that were generated live and displayed on screen in performance. A Max patch continuously performed spectral analyses on the performers' inputs. For each instrument, the patch collected the most prominent partials over a given timeframe and translated these partials into musical pitches. The patch sent each spectrally-derived pitch collection to a performer at a different node. This performer selected a cell from the current generation and freely mapped pitches from the received spectral sonority onto the contour of the cell (Figure 2-13). Since each performer selected pitches using overtones from the previous performer's output, the process embedded a quickly rising pitch trajectory into the resulting cells. Each performer influenced the pitch content available to the next (Figure 2-14). The process was engaging in free improvisations but was technically challenging while using the pre-composed cells, particularly in later generations where the pitch contours become more complex. I removed this process from the final draft partly due to the difficulty, but more importantly, I felt that the harmonic content during this section would be more coherent in relation to the composition as a whole if it connected to the interval chains that make up the following two movements (see section 3.2).

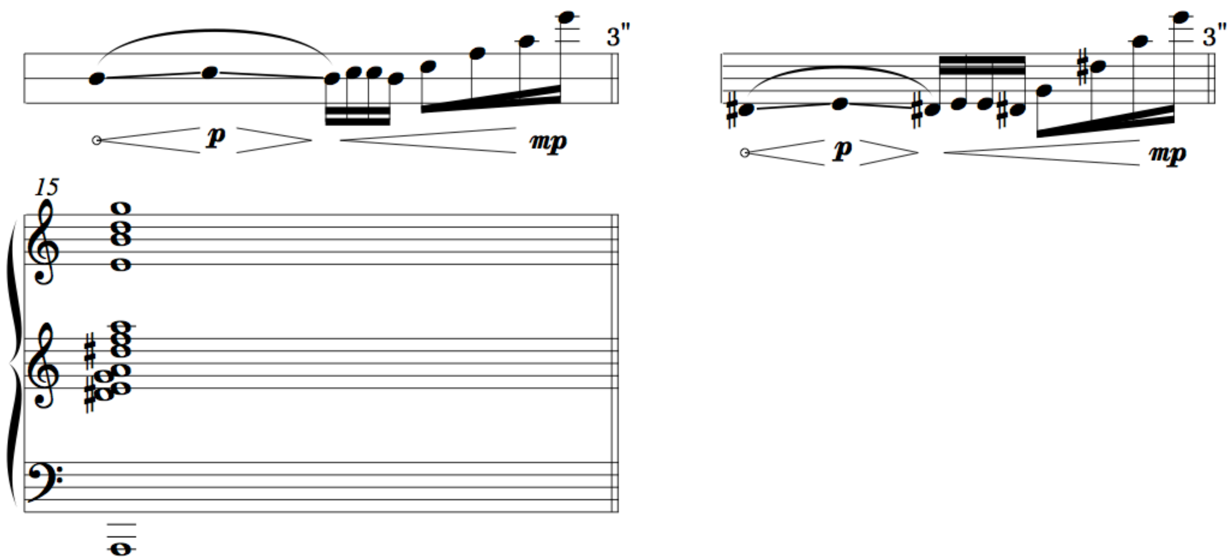

Figure 2-13 Unpitched cell (top-left) mapped onto a harmony generated through live spectral analysis of a remote performer (bottom-left) to create a pitched cell (right).

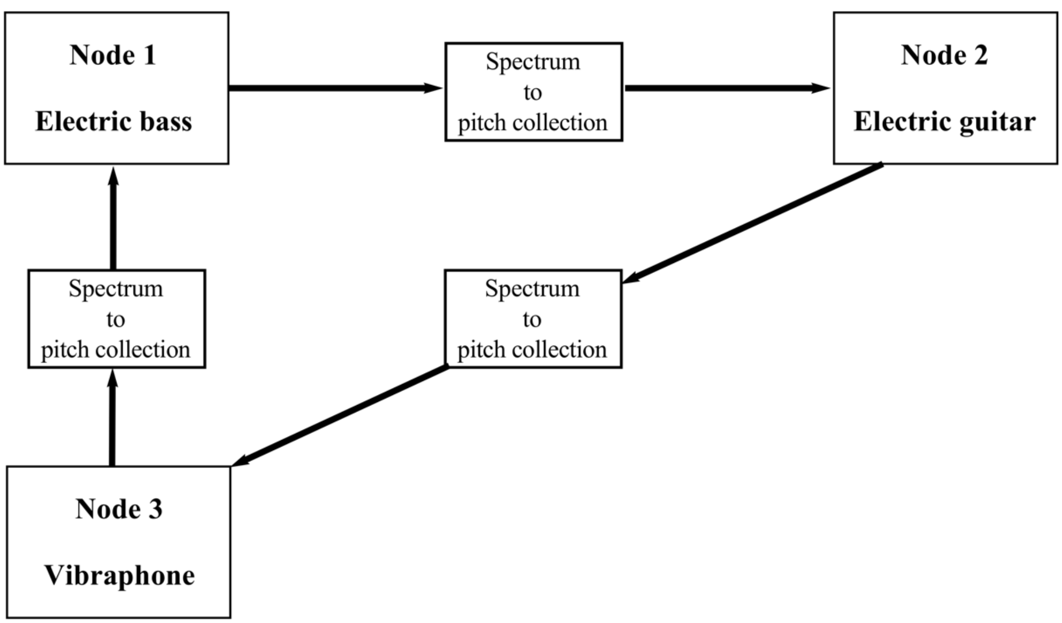

Figure 2-14 Intra-action overview in an early draft of *Through a Window*.

2.5 Sonification

Sonification is the translation of data into sound for informative or aesthetic purposes. It may serve simply as an auditory alert, such as a phone ring, or it may serve to expose patterns in data that are difficult to perceive visually. Chris Chafe has used sonification artistically in a variety of compositions and installations such as his *Tomato Quintet*, an installation in which CO_2 emissions, temperature, and light readings are measured over the course of ten days from sensors placed inside five vats of slowly ripening tomatoes. An algorithm sonifies the sensor data as live generative electronic music (Chafe 2007).

Föllmer identifies “interplay with network characteristics” as a commonly occurring aesthetic dimension in NMP projects and introduces the term “communication paradigm” to describe projects that translate the technology of network communication into the domain of sound (188-190). In the communication paradigm, musical structures are generated or informed by the performance medium itself. The Hub’s “leaky” network, described in section 2.4, is one example where network technology shapes the musical result.

In *Ping*, Chris Chafe employs the network as a musical instrument, using the transmission and reception of sound packets between nodes to “pluck” the network like a string (2009). The time-delay between the nodes corresponds to the frequency of the plucked note. Different configurations of nodes on a network are sonified as different harmonies.

Ping inspired me to use sonification as a technique to gain a more intuitive understanding of network characteristics and to challenge myself to create musical frames around extra-musical data. I measured latency and packet loss between my personal computer and various other nodes using ping, traceroute, Jacktrip, and Max (Cáceres and Chafe).[8] A 'ping' returns the amount of time in milliseconds for data to be sent to and returned from a user-specified address on a network. Ping times can range widely from less than 1 ms between machines on local networks up to and above 400 ms for slower and more distant connections. Since the range of audible frequencies in humans is approximately 20-20000 Hz, ping times up to roughly 50 ms can be represented as values in Hertz and notated as musical pitches to be performed by a synthesizer or musician.[9] 50 ms is equivalent to 20 Hz and therefore ping times above 50 ms will not directly map to musical pitches within the audible range. In the planning stages of *Through a Window,* I sent thousands of ping messages to various websites over WiFi, mapped the results to pitch in Max, and sent the pitch data from Max to NoteAbility Pro (Figure 2-15a-d) (Hamel, 2018). Close connections returned higher frequencies while more distant connections returned lower frequencies. My goal was to find addresses with noticeable differences in pitch range and uniformity under the 50ms latency threshold. Rather than perform these ping pitches directly within the composition, I instead looked for general characteristics within the data that I could use to inspire musical textures. The ping notes tended to gravitate towards an approximate 'main' pitch or pitch range with random chromatic deviations. The largest fluctuations away from this range tended to be ascending more often than descending. This is a result of the logarithmic nature of pitch perception. The time difference needed to cause a fluctuation of an ascending perfect fifth is roughly equivalent to the ping time difference needed to cause a fluctuation of a descending perfect fourth. Sonifying the pings directly as musical notes is more

---

[8] When sending information over a network, the data is segmented into 'packets'. Due to the unreliability of network technology, packets are sometimes lost and don't arrive at the remote location.

[9] Convert ping values to Hertz (cycles-per-second) using the following equation where T represents the ping time in milliseconds: $1000 / T = Hz$

sensitive to ascending deviations from the average ping time. The width of the deviations expands and contracts over time and the center pitch can also wander in response to network jitter.[10]

Figure 2-15 shows ping data from four different nodes translated into musical pitches. These examples loosely inspired the rapid chromatic textures from mm. 123-170 in movement 2 and mm. 122-166 in movement 3. In both cases, I established a 'home' pitch range and added ascending fluctuations derived from the network jitter. Unlike the ping material, the fluctuations away from the center during these sections are often strings of notes rather than single pitches.

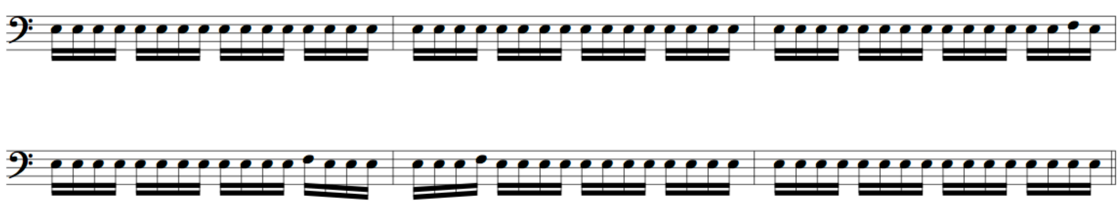

Figure 2-15a Ping times from node 1 translated to pitch

Figure 2-15b Node 2

---

[10] Jitter describes the fluctuation of latency on the network. It is used in combination with latency, bandwidth, and loss to describe a network's overall quality of service (QoS) (Tanenbaum, 405).

Figure 2-15c Node 3

Figure 2-15d Node 4

The traceroute tool is similar to the ping tool, except that it returns the roundtrip latency for each hop along the network towards a specified destination. Inspired by Chafe, I considered each hop along the path from source to destination as modes of vibration on an imaginary string. Again, I translated latency times to pitch to generate chords. Each chord tone is a single hop on the route from source to destination. Chords derived from traceroute data between my home network and the University of Calgary resulted in one high note followed by a cluster roughly three octaves lower with an occasional outlying low pitch.

In movement 1, from 0:00-1:12, ensemble members are directed to perform short musical fragments anytime they receive a 'jitter' cue. The jitter cues are a rhythmic representation of dropped audio packets on a network. In live performance, dropped audio packets are heard as short buzzing glitches in the remote audio and are a common occurrence when the network does not have sufficient bandwidth. I translated this unpleasant characteristic into a musical texture by preserving the delta timing between glitches and replacing the glitch sound with performed gestures. To measure the delta time between dropped packets, I created an audio connection between two laptops over WiFi and sent a signal with DC offset "1" from one laptop to the other and back. The looped-back signal contained an audio frame of 0's rather than 1's anytime a packet was dropped. I recorded the time between any dropped packets a minimum of 150 ms apart and applied a scalar to the resulting data. The process was repeated four times with different buffer sizes and redundancy parameters in JackTrip. The scaled time values are used to trigger visual cues for the musicians as well as sample playback in the electronics. Similarly to the process in *Wax Lips*, the texture is formed through the network's unreliability.

The final sonification experiment was a networked rendition of *I am Sitting in a Room* by Alvin Lucier. Rather than playing the recorded speech through a speaker and repeatedly re-recording the speech from the speaker output to pick up room resonance, I looped the audio between two computers on a network and recorded the loopback sound. The networked rendition creates a much more abrasive effect than the original. As network glitches slowly accumulate, the speech shifts from a clean recording to continuous distorted noise.

2.6 Through a Window

While planning *Through a Window*, I used the themes discussed above to assemble a palette of NMP techniques and selected freely from this palette during the compositional process. As a result, some techniques were used extensively while others were used in a more superficial manner or remained unused. Table 2-3 provides a list of the NMP techniques and includes a short description of each technique and an indication of whether it was used in the final composition.

The three categories of techniques in Table 2-3 reflect three roles that the network plays within the composition and presentation of the work: the network generates musical structure through sonification and intra-action, the network helps create drama and mediates presence through sound processing and amplification, and most importantly the network enables musicians to perform complex and nuanced music together across great distances through careful consideration of latency and aural cueing.

Table 2-3 NMP composition techniques

Title		Description	Used
Network characteristics as inspiration for musical material			
Ping		Measure transmission times between nodes on a network and translate the results from ms to Hz.	Yes
Traceroute		Translate transmission times for all hops between two nodes into chords or timbres.	No
Jitter		Cue events on dropped packets.	Yes
Intra-activity		One performer's input directly impacts another through the manipulation of performance parameters.	No
Network as a performance medium			
Mediated Presence (Audio and video)	Hyper-presence	Remote performers rendered 'larger-than-life'.	Yes
	Reduced-presence	Remote performers attenuated or distorted.	Yes
	Ambiguous-presence	Ambiguity between live and recorded sound due to lack of visual cues.	Yes
	Spatialization	Manipulation of spatial location of performers (to create ambiguous presence or otherwise).	Yes
Toporhythms		Ping time and artificial delays used to create rhythms which are offset by beat values between the locations.	No
Site-specific	Soundfiles	Assigning different soundfile banks to each location.	Yes
	Processing	Assigning different processing tools or parameters to each location to create unique harmonizations, textures, durations, etc.	Yes
	Sound analysis	Using different instruments to cue responses in the electronics at each site.	Yes
Practical approaches for aiding ensemble performance			
Synchronized stopwatches		Using stopwatches to provide a broad form of objective temporal reference.	Yes
Anchors		Assigning one performer to provide the temporal reference point for the other musicians.	Yes
Aural cues		Short, sharp cues to signify a change of section or establish a new tempo.	Yes
Soundfile cues		Electroacoustic click-track.	No
Catch-up points		Fermatas or other pauses to allow musicians to regain footing after a fast or rhythmically complex section	Yes
Mediated orchestration		Actively mixing the audio of remote musicians to ensure that anchor musicians are clearly audible.	Yes

In the composer's notes for *Music for Wilderness Lake*, R. Murray Schafer makes the following assertion:

> The history of music shows that although musical styles change constantly, the contexts in which music is presented vary much less frequently. […]
>
> The big revolutions in music history are those with the power to change performance contexts. It is these which govern performance rituals and legislate musical forms and instrumentation. (Schafer, i)

Composing for any particular instrumentation or performance venue should inform a composer's musical choices. Different instruments have different sonic possibilities and different practical limitations. Different rooms can create different amounts of reverberation as well as different amounts of intimacy between the audience and the musicians.

Network music certainly challenges many traditional aspects of music performance. Lack of direct visual communication, latency, presence, and even performance time zones are all challenges that rarely occur in the concert hall. The networked music performance environment informed my choices of timbre, instrumentation, and acoustics, as well as the general compositional themes and practical approaches.

However, within these limitations, there is still an incredibly large range of musical possibilities. Much of the pitch and rhythmic organization of the work is independent of the network itself. The next chapter will proceed deeper into the structural organization of the composition and reflect upon the role of computer algorithms as an aid for composing musical processes.

Chapter 3 Using algorithms to develop musical processes

3.1 Algorithms as musical process

In *Through a Window*, a collection of computer algorithms, often modelling processes in nature, informed the ideas, form, and textures in the music itself. In this chapter, I will show how recursion as well as evolutionary and swarm algorithms are expressed in the music in the form of processes that unfold over the course of individual sections.

In *Optimizing Future Perfect: A Model for Composition with Genetic Algorithms*, Geoff Holbrook uses the terms "descriptive" and "generative" to describe two common approaches to algorithmic music composition.

> By the descriptive approach, we control *what properties* the output must have. By the generative approach, we control *by what process* the output of the algorithm is produced. (Holbrook, 3)

Through a Window uses the generative approach to algorithmic composition. My goal was not to design software to write entire compositions or compose in a specific style, but rather to search for and uncover an underlying musicality within specific types of algorithms. I chose algorithms that I found to be conceptually beautiful and explored ways in which I could express these algorithms as musical processes.

In the algorithmic sections of the composition, I shifted my compositional approach from forming a collection of characteristic harmonies or motives and then arranging these elements in time to instead forming a collection of characteristic musical processes from which the musical materials emerge. The processes recur and are realized in different configurations throughout composition. The relationships between notes or gestures that emerge from these processes are like the movements of leaves caught in a gust of wind and witnessed through a window. The motions of the leaves reveal the presence of an unseen and unfelt force.

3.2 Recursion and self-similarity

Recursion describes any function that calls upon itself within its own execution. It is a basic principle in computer science and is used to perform a wide variety of tasks such as fractal image generation, sorting, and tree traversal. Figure 3-1 shows the result of a simple recursive function named “drawShinkingCircles”. The function draws a circle at a given diameter, multiplies the diameter by a value between 0.1 and 0.99, and then calls itself again with the newly modified diameter. The process repeats until the diameter has shrunken to a user-defined minimum size. Figure 3-2 shows two variations of the same function with spatial offsets applied to each circle to create self-similar patterns.

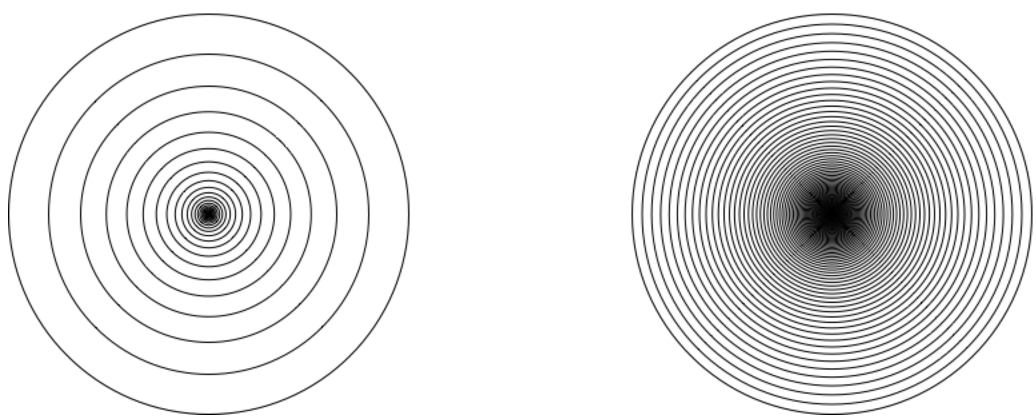

Figure 3-1 Output of drawShrinkingCircles. Diameter scaled by 0.8 (left) and 0.95 (right)

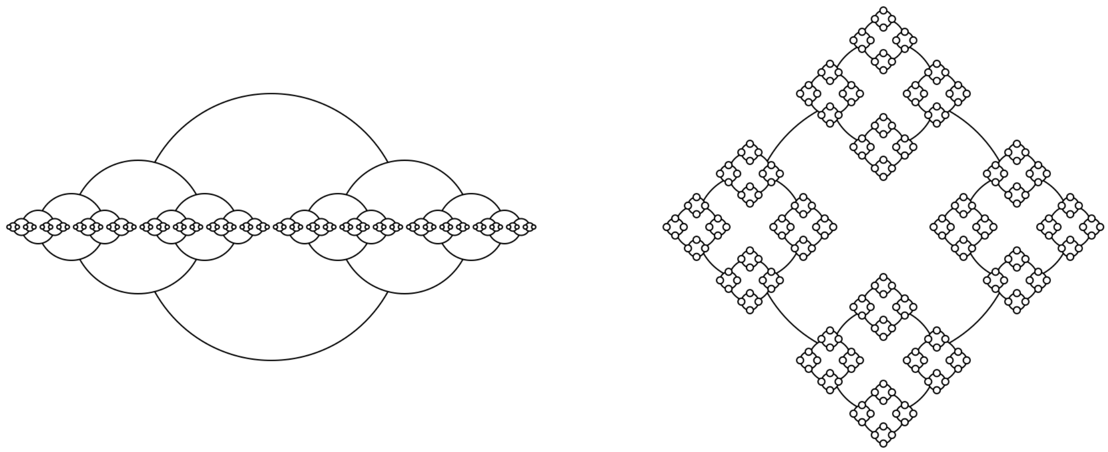

Figure 3-2 Two outputs of drawShrinkingCircles with spatial offsets.

Recursion and self-similarity have been widely explored in both music analysis and composition. In analysis, Lerdahl and Jackendoff's Generative Theory of Tonal Music (GMMT) use recursive tree structures to group music in hierarchical relationships (1996). Robert Morris's contour reduction algorithm also uses recursion to reveal underlying contour relationships in melodic material (212). In composition, Bruno Degazio's *On Growth and Form* (1988) and Charles Wuorinen's *Cello Variations* (1970) both use recursive structures to create self-similar forms. Degazio does so using Fibonacci relationships to develop material (i). Wuorinen builds nested formal structures based on the proportions of the opening pitches (McConville, 173). Alvin Lucier's *I am sitting in a room* (1969) is also recursive: each iteration of Lucier's text applies the same process to the output of the previous iteration (Simoni, 288).

The intervallic material in movement 2 from mm. 203-273, movement 3 from mm. 1-37, and movement 3 from mm. 54-121 is built from progressively expanding compound interval cycles generated using a simple recursive function implemented in the C programming language. In the following sections, I show how I combined this recursive function with two other cyclic functions governing registral span and rhythm to create a kind of clockwork musical form.

3.2.1 Interval cycles

George Perle defines an "interval cycle" as the pitch classes generated from "a single recurrent interval in a series that closes with a return to the initial pitch class" (1990, 21). Edward Gollin refers to Perle's definition as a "simple interval cycle" and uses the term "compound interval cycle" to denote cycles generated from a pattern of two or more intervals (Gollin, 146). The notation "*x*-cycle" denotes simple interval cycles, where *x* is the "generating interval", and "(x, ..., n)-cycle" denotes compound interval cycles containing two or more intervals. Using this notation, the (01245689t) pc-set used prominently in movement II from mm. 1-203 can be notated as a (1,1,2)-cycle (Figure 3-3).

Figure 3-3 (1,1,2)-cycle

Gollin describes the cycles as "compound" because they can be represented as a set of multiple "interlaced" simple *x*-cycles where *x* is the sum of the intervals (Figure 3-4). In my own listening, I also perceive compound interval cycles as pc-sets that are built on the notes of a simple interval cycle (Figure 3-15). I can hear the (1,1,2)-cycle as both three interlaced augmented triads and three (012) sets. I used the (1,1,2)-cycle prominently within the composition to represent the three nodes of the network. Any chord tone can be perceived as the root in an augmented triad; the three (012) sets, built on the tones of an augmented triad, symbolically represent equality between the nodes. Additionally, since the cycle is symmetrical, I can pitch-shift instruments by major thirds or minor sixths and still express the same collection.

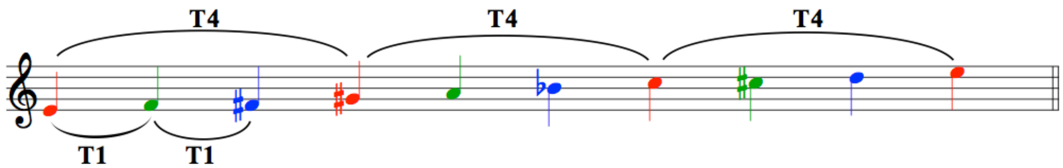

Figure 3-4 (1,1,2)-cycle represented as three interlaced 4-cycles

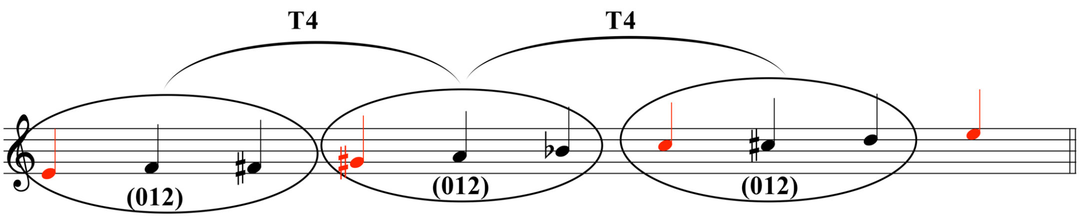

Figure 3-5 A (1,1,2)-cycle represented as three (012) sets built on the pitches of a 4-cycle.

The recursive function I used in *Through a Window*, titled *nb.cycle*, generates every possible triple interval cycle in which each of the three intervals are limited between a minor second and a major seventh. The process is similar to incrementing the hands on a clock: The third interval in the cycle (the second hand) is repeatedly incremented by one until it reaches the maximum value. At this point the third interval is reset to the minimum value and the second interval (the minute hand) is incremented by one. Similarly, when the second interval completes its cycle, the first interval (the hour hand) is incremented by one. The expansion process is therefore a cycle of cycles. Figure 3-6 illustrates a portion of this process; the three intervals are shown as the three hands on a mod-11 clock and the corresponding triple interval cycles are notated below the clock images.

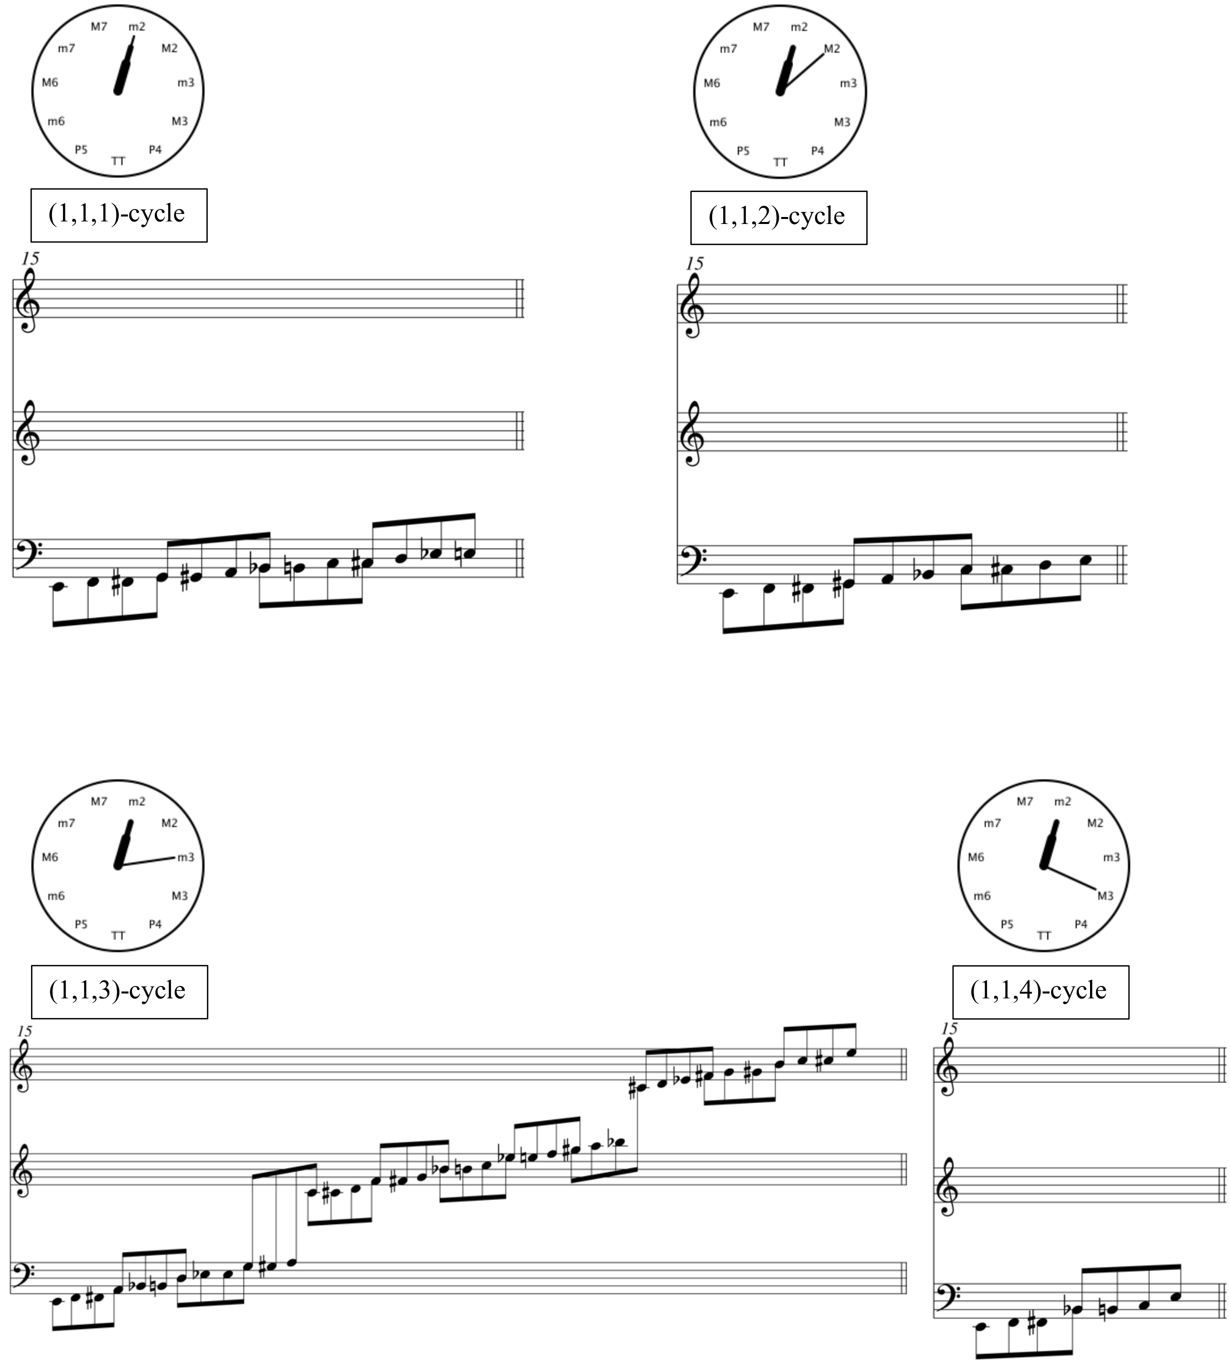

Figure 3-6 nb.chain interval cycle expansion process. (1,1,1)-cycle to (1,2,1)-cycle.

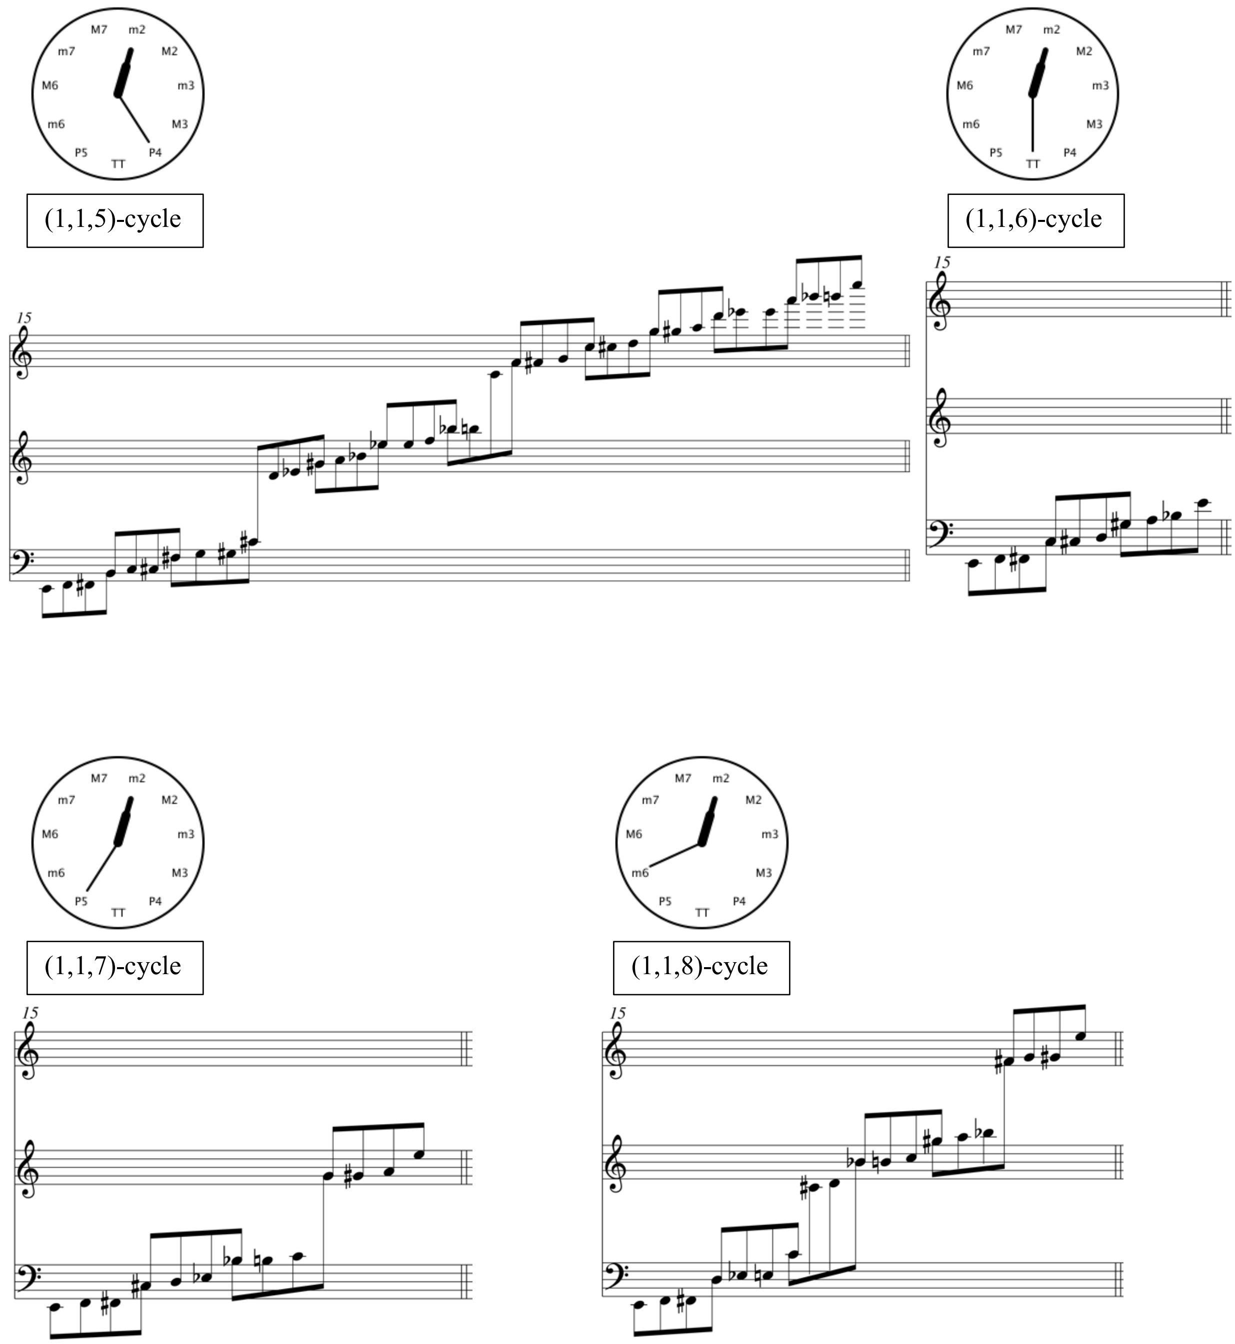

Figure 3-6 (continued)

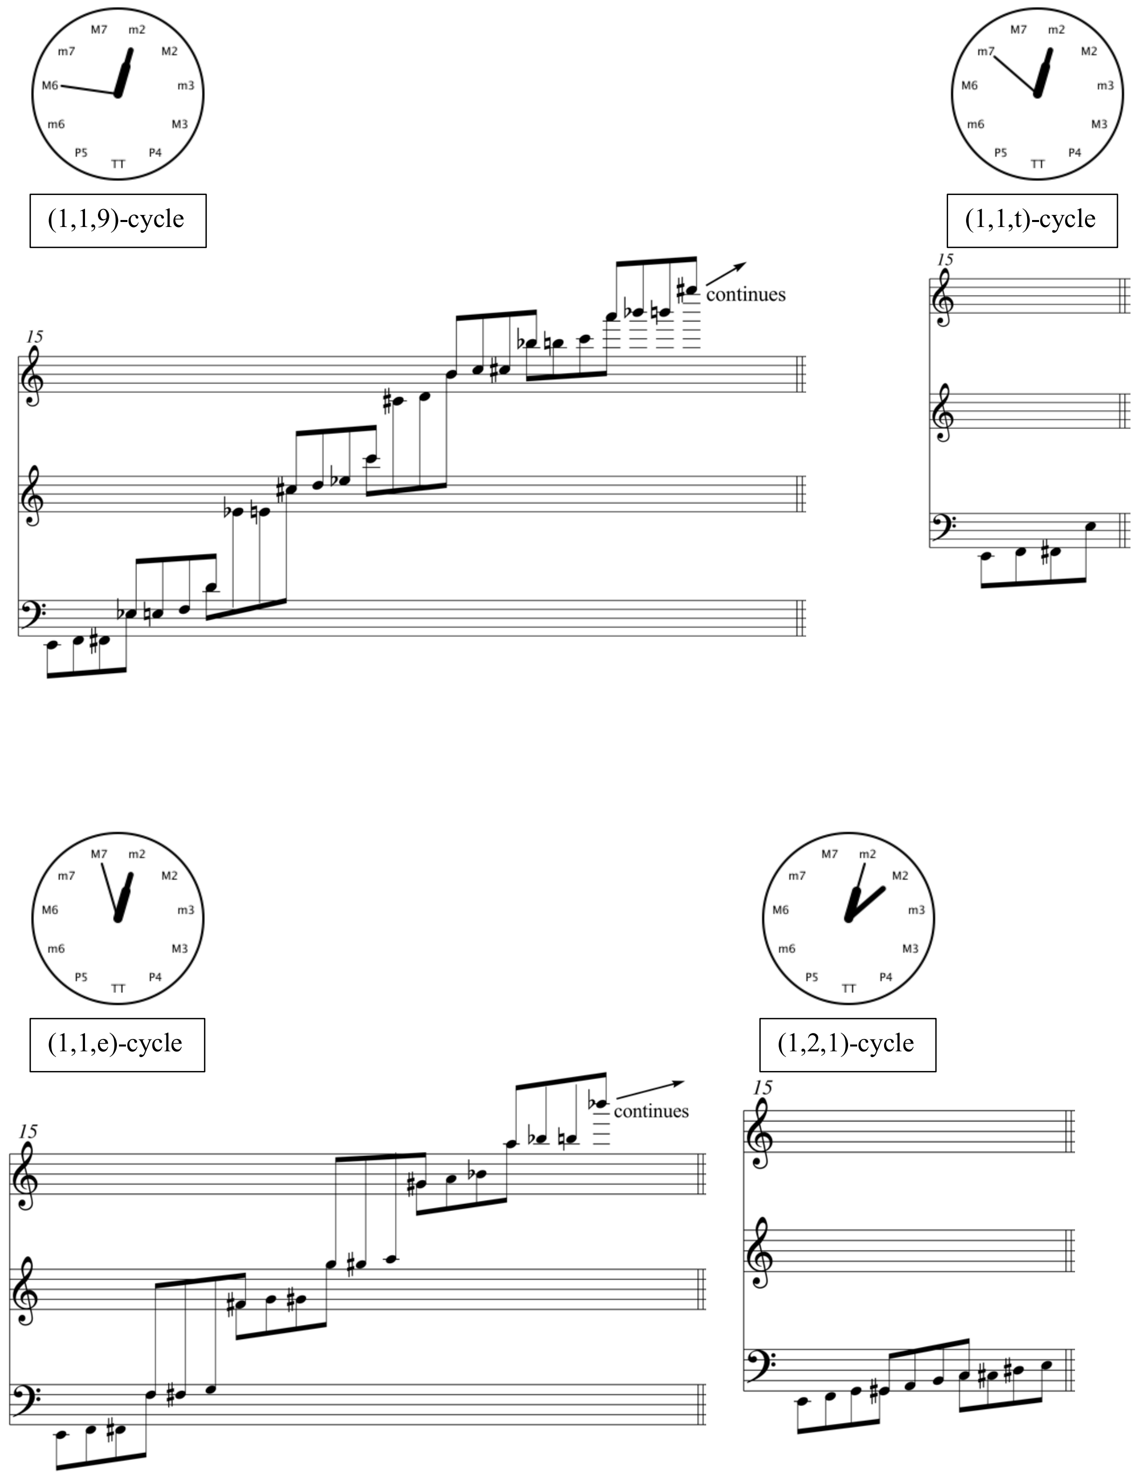

Figure 3-6 (continued)

The number of possible cycles for a given number of intervals is i^n where i is the range of intervals used and n is the number of intervals in the cycle.[11] In *Through a Window*, this equates to $11^3 = 1331$ compound cycles. The number of compound interval cycles can be reduced significantly if rotationally-related interval cycles are treated as equivalent. For example, if the (1,1,2)-cycle, (1,2,1)-cycle, and (2,1,1)-cycle are considered equivalent then the list of triple interval cycles drops from 1331 to 285.

Figure 3-7 graphs the sum of the three intervals in the cycle throughout the entire process. The sum rises as interval 3 increments but drops when interval 3 completes its cycle and interval 2 increments. The same process occurs for interval 2 and interval 1 on a longer timescale. Each interval slowly expands to its maximum value, creating wide gaps between adjacent pitches, before suddenly collapsing back to its minimum value, resulting in chromatic steps between adjacent pitches. The sawtooth-like pattern is self-similar across three timescales. The expansion and sudden collapse occurs for each of the three intervals. Figure 3-8 outlines the same process using interval classes rather than intervals graphed on a mod-6 space. This representation shows that the expansion process has a palindromic quality.

[11] This includes cycles that can be represented as a single interval cycle such as the (1,1,1)-cycle.

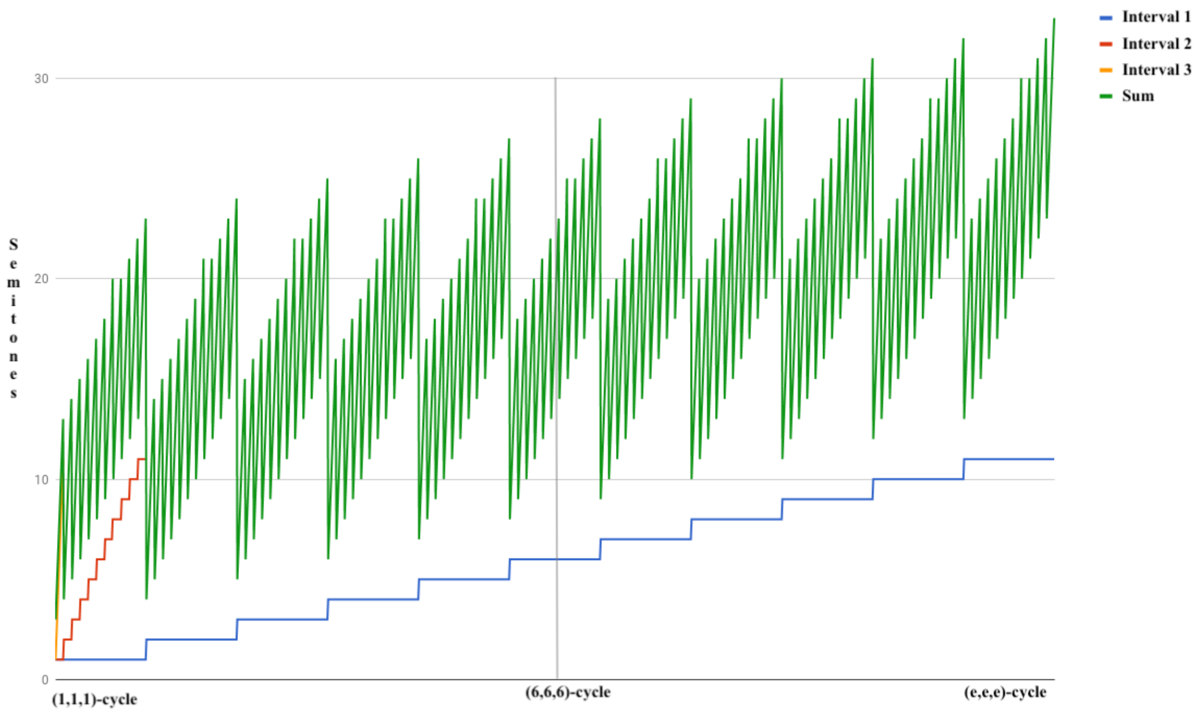

Figure 3-7 Sum of the three generating intervals in each triple interval cycle.

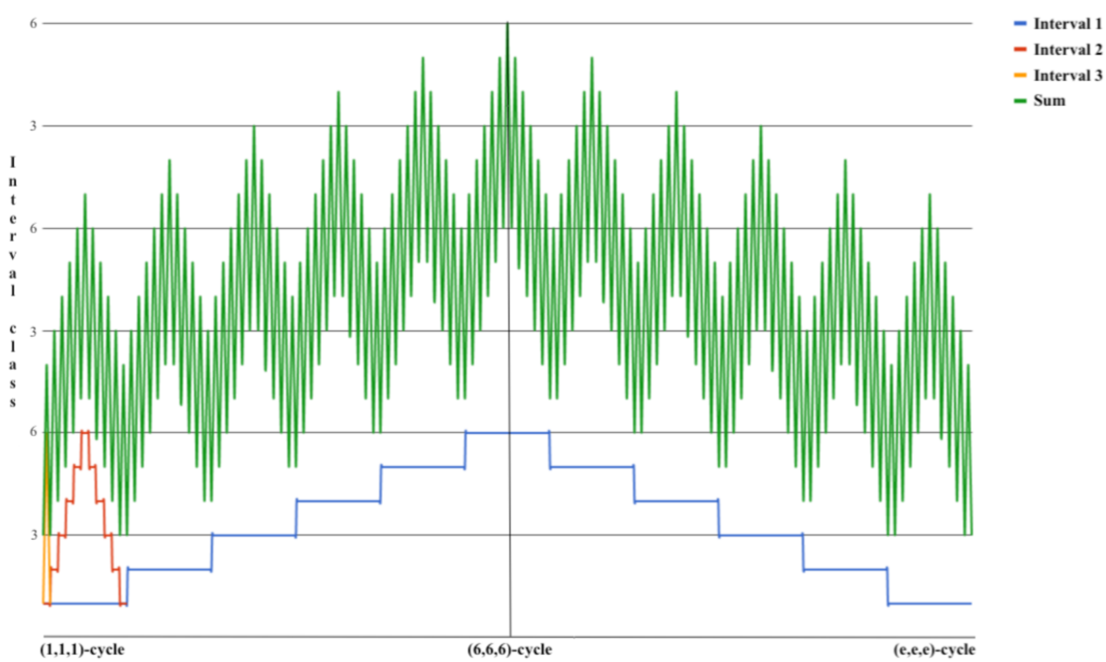

Figure 3-8 Sum of interval classes for all triple interval cycles. Y-axis loops in mod-6 to clarify the pattern.

3.2.2 Harmonic rhythm based on misalignment from the octave

Adjacent interval cycles in the process may evoke different tonal sensations. For example, the (1,1,e)-cycle in Figure 3-6 contains many more pitches and a much wider registral span then the adjacent (1,1,t)-cycle and (1,2,1)-cycle. However, the step-wise expansion process creates a continually rising chromatic quality which helps create a sense of consistency as the process unfolds. The sensation is as if the shadow of a rising chromatic scale is projected through the cycles.

Figure 3-9 presents an excerpt of the recursive process from (6,1,1)-cycle to (6,e,e)-cycle. The cycles are presented as a series of pc-clock diagrams with a single row representing a full cycle of interval 3. Row 1 progresses from (6,1,1)-cycle to (6,1,e)-cycle, row 2 progresses from (6,2,1)-cycle to (6,2,e)-cycle, and so on. At a glance there is a clear sense of symmetry around the center cycle (highlighted in red). The symmetry is related to both the interval class and the number of pitches needed to complete the cycle. In the following section I discuss this palindromic relationship and show how this embeds a symmetrical harmonic rhythm into the process as a whole.

The number of pitches needed to complete an interval cycle varies depending on how evenly the sum of the intervals fits within the span of an octave. Intervals that divide an octave evenly require fewer iterations to complete the cycle. The exact number of pitches can be determined using the following equation where *n* is the number of intervals in the cycle and SUM is the sum of the intervals: $\frac{12}{GCF(12,SUM)} * n$ (Gollin, 146).

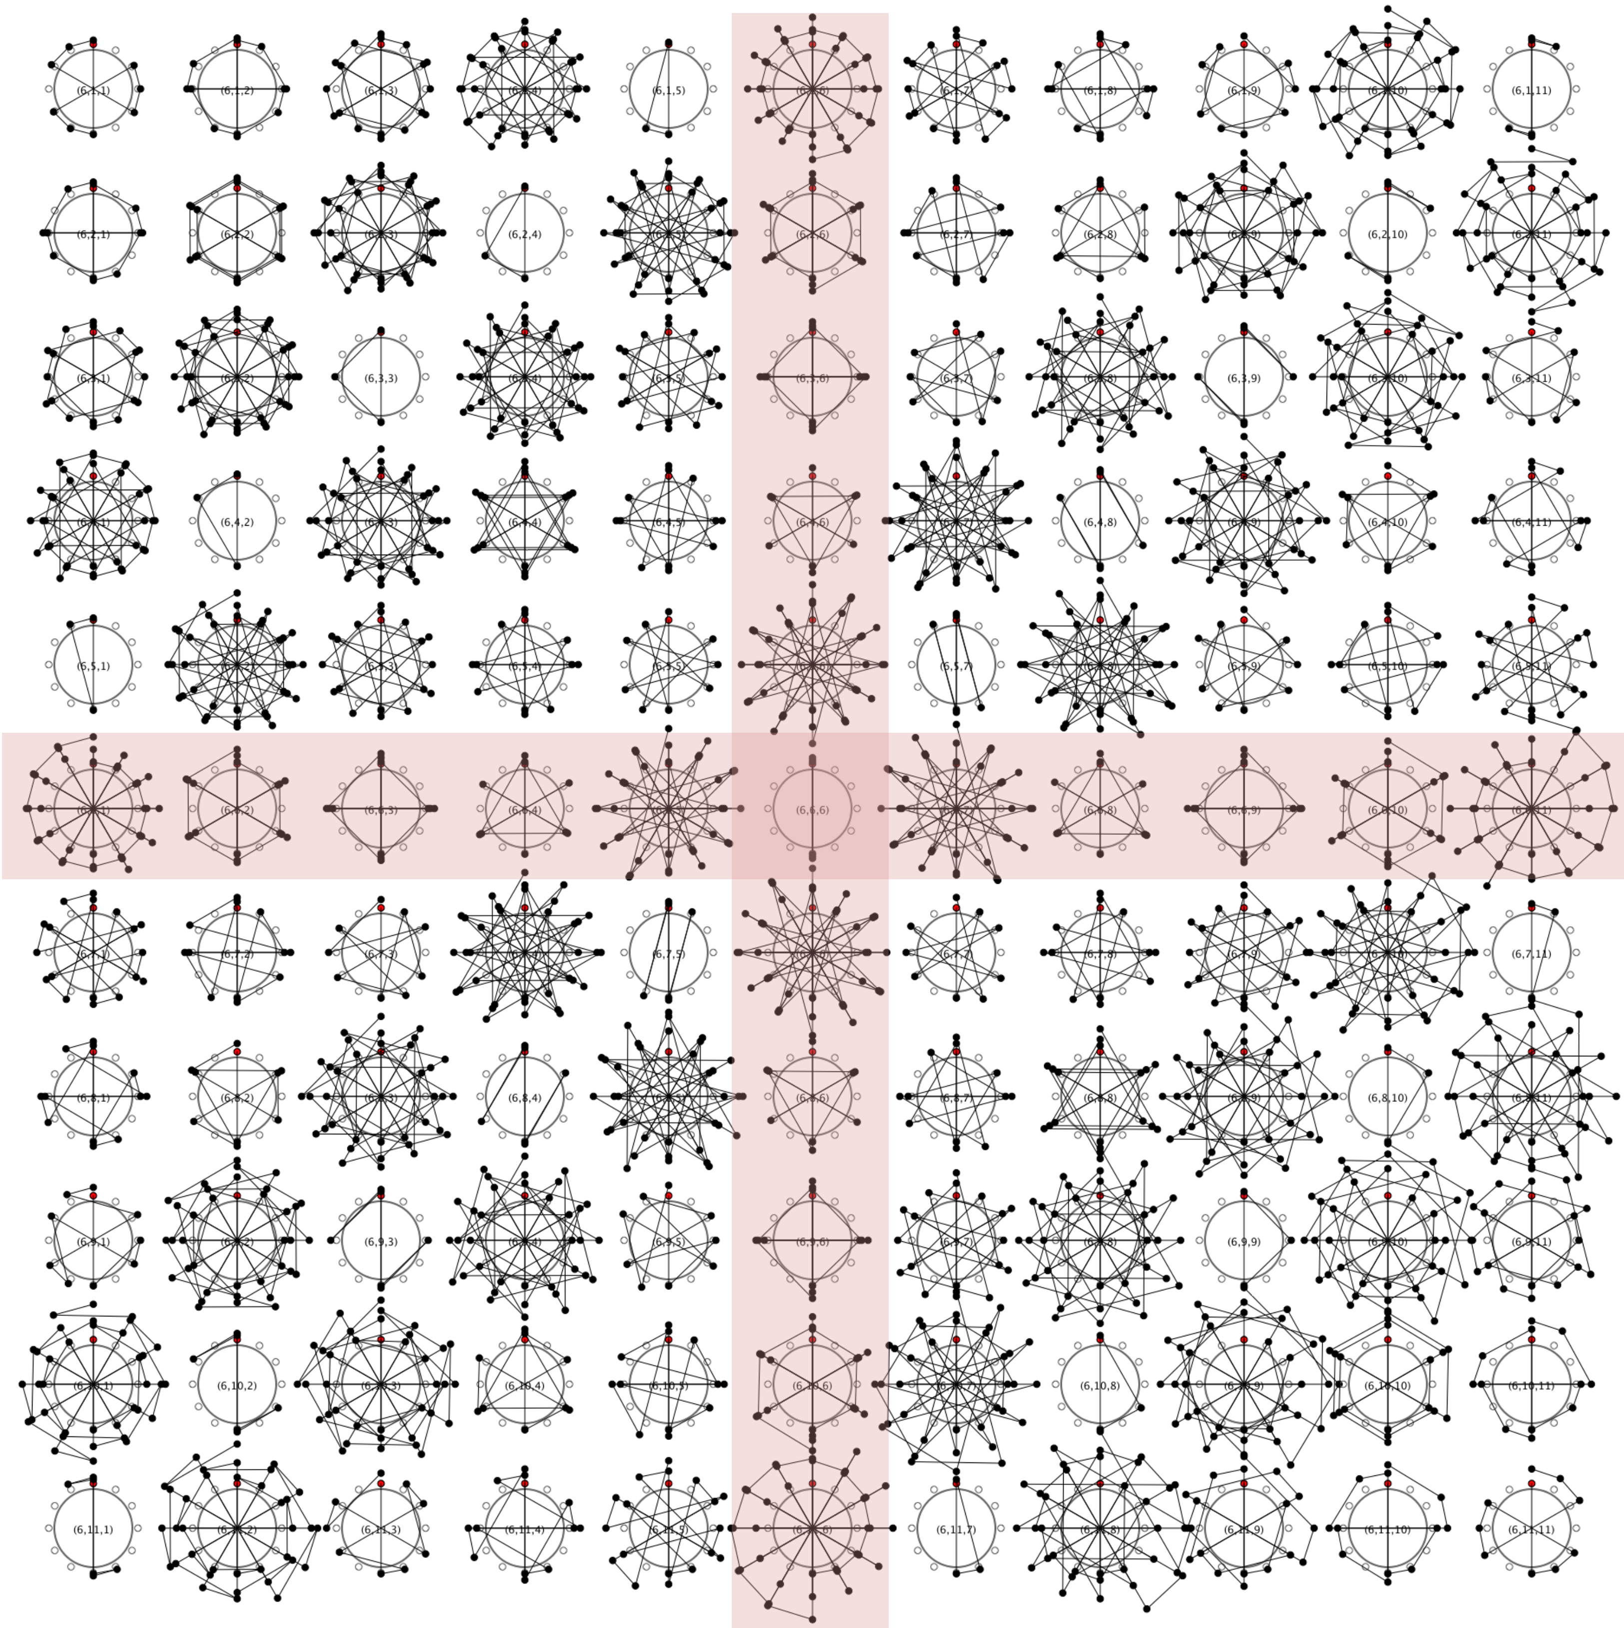

Figure 3-9 *nb.cycle* from (6,1,1)-cycle to (6,e,e)-cycle. Symmetrical pattern in red.

Since an octave has only six factors, there are only six different possible lengths of interval cycle for a given number of intervals. The possible lengths of a triple interval cycle are the factors of 12 multiplied by 3: 3, 6, 9, 12, 18, and 36.

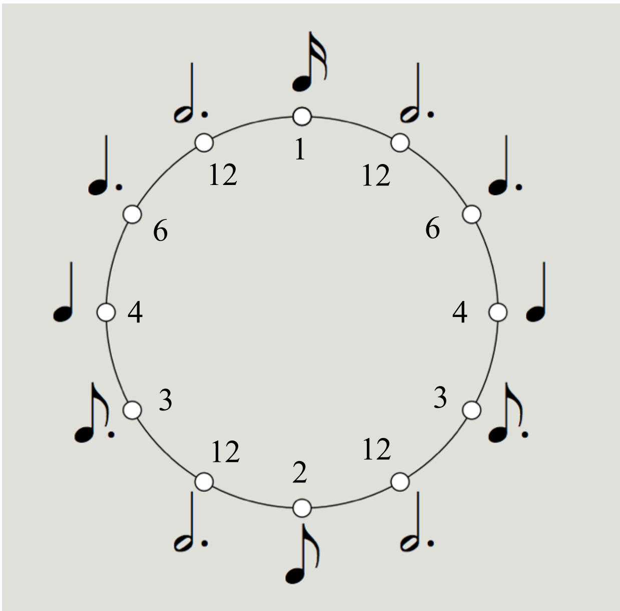

Figure 3-10 Proportionate lengths of progressively expanding interval cycles represented both numerically and with rhythmic proportions. Multiply the values by three for triple interval cycles.

If a series of interval cycles is presented where the sum of the intervals between adjacent cycles is steadily incremented by one, then the cycles will iterate through these six lengths in a palindromic pattern. Figure 3-10 notates this pattern on a clockface using rhythmic values to show the relative proportions of the cycles. If the expanding cycles are played back-to-back in a scalar pattern, then the circle in Figure 3-10 generates the harmonic rhythm notated in Figure 3-11.

Figure 3-11 (1,10,1)-cycle to (1,10,11)-cycle. The proportional harmonic rhythm for each cycle is notated above the staff.

The palindromic pattern creates an interesting balance between predictability and complexity. The 1-4 o'clock region creates a sense of acceleration and the corresponding 8-11 o'clock region creates a matching deceleration. The 12 o'clock and 6 o'clock values create a sudden rhythmic contrast in comparison to the surrounding dotted half note values.

The *nb.cycle* process generates the palindrome pattern while the third interval increases from minor second to major seventh. The starting point in the pattern rotates forward every time the second interval increments as shown in Figure 3-12. Each note represents an interval cycle and each row represents interval 2 incrementing. The rhythms in the columns match the rhythms in the row.

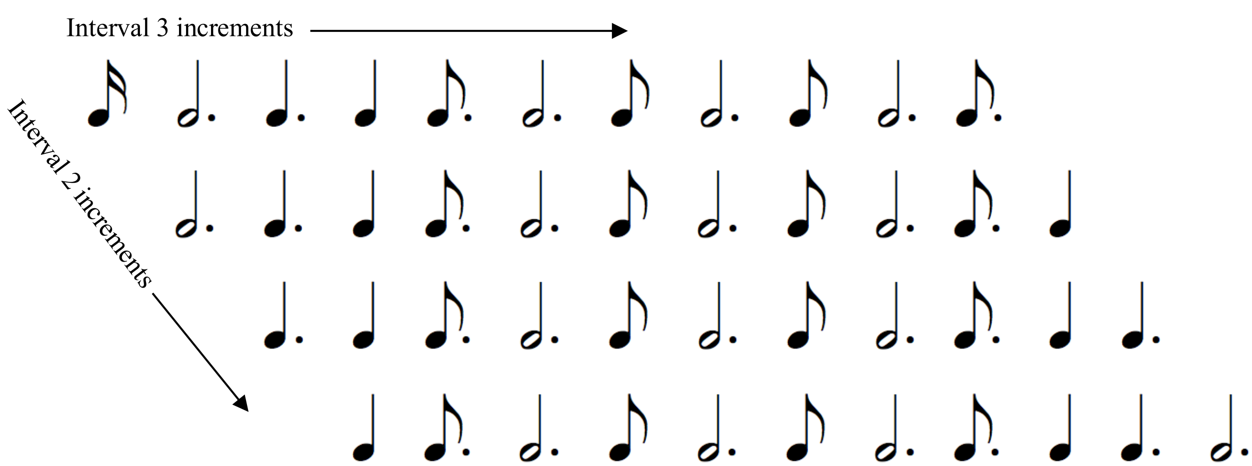

Figure 3-12 An excerpt of the *nb.cycle* harmonic rhythm cycle. Each note represents the proportionate duration of a cycle.

3.2.3 Boundary conditions

Since the triple interval cycles used in *Through a Window* can contain as many as 36 pitches per cycle and intervals as large as a major seventh, the triple interval cycles can easily exceed typical instrumental pitch registers. For example, a (10,10,11)-cycle requires 12 iterations of its three intervals for completion, extending upwards 31 octaves. 31 octaves is far beyond the range of human hearing. Therefore, the cycle can't be meaningfully represented as a single ascending audible pattern. Instead, the cycle needs to be constrained to a smaller register in order for all the pitches to be performed and heard.

To systematically constrain the cycles within a given pitch range, I included upper and lower pitch boundaries in the cycle-function and implemented several possible responses whenever a pitch lands outside of one of these boundaries. The user can select one of these responses when running the function. The chosen method strongly influences the musical outcome. Figure 3-13 illustrates two cycles that need to be reduced in registral span. Seven possible solutions are described below.

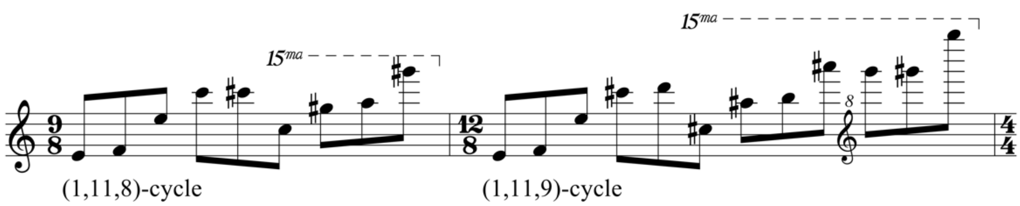

Figure 3-13, Two triple interval cycles beginning on E4 that are not constrained by register. Both cycles exceed most instrumental pitch ranges and the (1,11,9)-cycle reaches the upper limit of human pitch perception.

Truncate

Truncation simply terminates the cycle whenever a note exceeds the upper boundary, resulting in a series of ascending Figures (Figure 3-14). The progression from one interval cycle to another is clearly perceived as a wide descending leap from the final note of one cycle to the initial note of the following cycle.

Figure 3-14 includes rests in place of pitches that exceeded the upper boundaries. However, the rests could also be removed and the truncated cycles placed directly adjacent to one another. In this case, the wider intervals that occur in cycles at the end of the large-scale expansion process will exceed the pitch boundaries more quickly than the cycles at the beginning of the process, therefore leading to shorter, incomplete cycles and embedding a 'harmonic' acceleration over the course of the intervallic expansion process. In this case, pitch and form are intertwined. The harmonic rhythm is directly connected to the intervallic content.

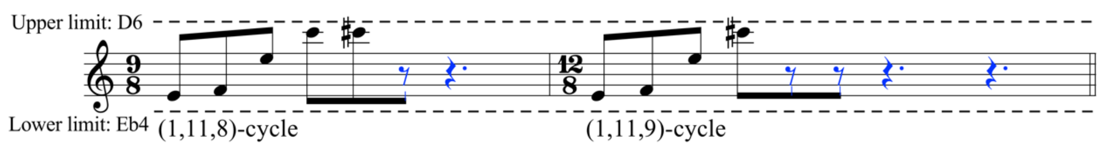

Figure 3-14 Truncation. Pitches are limited between Eb4 and D6.

Octave shift and wrap

Figures 3-15 and 3-16 both reset pitches back near the bottom of the allowable range whenever a pitch crosses the upper boundary. 3-15 preserves the intended pitch-class and continues the pattern using the modified octave. This creates a more complex pitch contour than truncation. The progression from one cycle to another is less clearly articulated than truncation.

In Figure 3-16, wrapping occurs whenever pitches would otherwise exceed the upper boundary. The function subtracts the upper boundary pitch from the intended pitch value and adds the resulting interval to the lower boundary pitch. Pitch class discontinuities occur unless the boundaries are related by an octave. For example, the sixth note in 3-16 should be a C7, however, the outer limits are related by a Major seventh, so the pitch wraps to C#5 rather than C5. The interval cycle is distorted by the space in which it is represented.

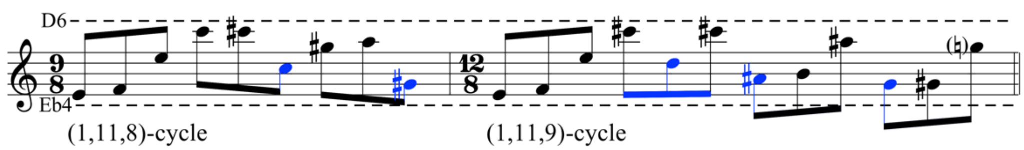

Figure 3-15 Octave-shift by maximum amount. Blue notes have been octave-shifted to the lowest possible matching pitch class within the Eb4-D6 span.

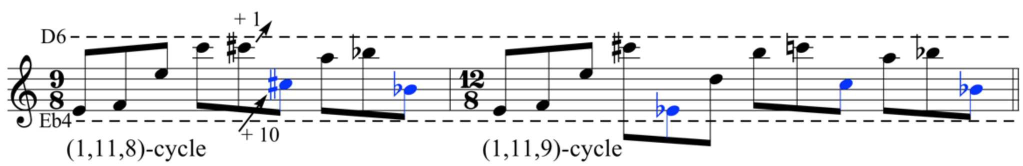

Figure 3-16 Wrapping. Blue notes have been wrapped between the boundaries, similar to a modulo operation.

Invert and fold

Figures 3-17 and 3-18 both reflect off the boundaries. If the intended pitch would otherwise cross a boundary, Figure 3-17 simply inverts the interval and continues in this new direction. Figure 3-18 measures the distance by which the intended pitch exceeds the boundary and then subtracts this value from the boundary to constrain the current pitch. As with wrapping, the operation introduces a pc discontinuity to the remaining portion of the cycle.

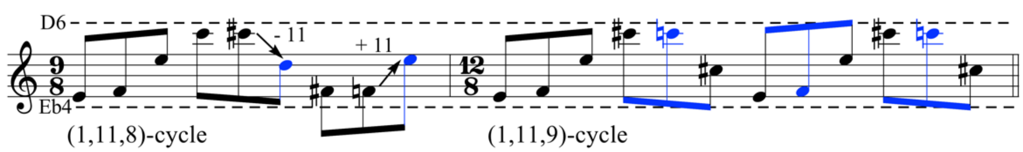

Figure 3-17 Invert and continue in new direction.

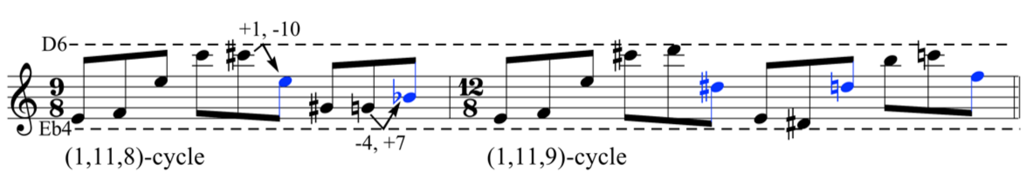

Figure 3-18 Fold and continue in new direction.

Registral sticking points

The boundary conditions shown in Figures 3-19 and 3-20 both result in ascending figures that become stuck in the upper register once a pitch crosses the upper boundary. Using these methods, the rising portion of the gesture acts as a perceivable cue at the onset of each new cycle.

Although the recursive section in movement 2 from mm. 203-273 contains the octave shifting method illustrated in 3-15, the onset of new cycles are articulated similarly to Figures 3-19 and 3-20. In mm. 203-273, *nb.cycle* places the first notes in each cycle below the lower boundary. The pitches then shoot immediately upwards into the allowable range and remain there for the rest of the cycle. The entries for each cycle are therefore articulated as a rising figure that becomes stuck.

The interval cycles are considered complete once the initial pitch material is revisited and the function would otherwise begin outputting redundant pitch content. However, since the constraint methods can alter the intended pitches in a cycle, it is possible that the cycles will not terminate after the expected number of pitches have occurred. If pc discontinuities occur, the function can either terminate the cycle after the expected number of pitches that would have occurred if the boundary method had not been applied, or it can continue until the initial pitch classes are revisited. If the discontinuities make it impossible for the original pitch classes to be revisited, then the method will terminate the cycle once a maximum number of notes have been generated.

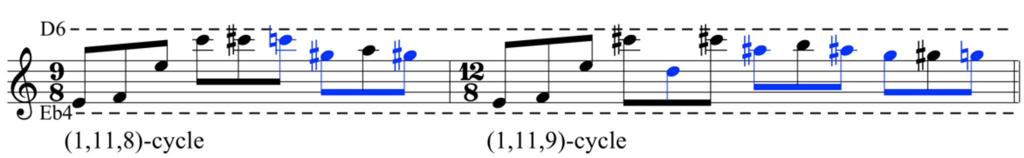

Figure 3-19 Octave-shift by minimum amount.

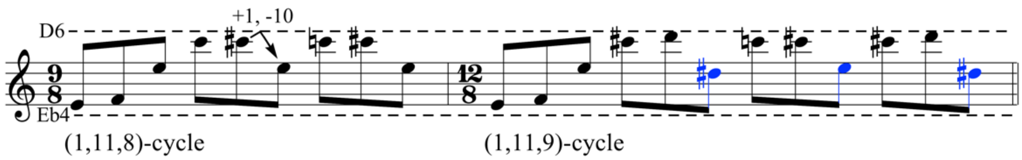

Figure 3-20 Fold without reversing the direction.

Cycling boundaries

Registral boundaries may be static or dynamic from cycle to cycle. Static boundaries remain fixed at their initial position throughout the entire process while dynamic boundaries change at the start of each new interval cycle. The three recursive sections in *Through a Window* all employ the same process of transposing the outer pitch boundaries in cyclic patterns. Figure 3-21 presents the outer boundaries for movement 3 from mm. 54-116. In this section, *nb.cycle* initially sets the upper limit to G#4 and incrementally transposes the limit upwards to B6 at a rate of a whole tone per interval cycle. Once reaching B6 (or above), *nb.cycle* sets the upper limit to C5 and the process continues. *Nb.cycle* initially sets the lower boundary to C4 and transposes the boundary upwards by a semitone per cycle until the boundary reaches G#5. At this point *nb.cycle* resets the lower boundary to C4. Anytime the lower boundary reaches a higher value than the upper boundary, *nb.cycle* swaps the two values and the process continues normally.

In some cases, the upper and lower boundaries are less than an octave apart. In these cases, the octave-shifting constraint method used in this section is not able to adjust outlying pitches back within the correct range. The pitches remain outside the boundaries. For example, see the first measure of Figure 3-22.

The sudden change in registral span that occurs when a boundary jumps from its maximum value to its minimum value is easily perceived while listening and helps to articulate the beginning of the cycle. For example, in mm. 86 in Figure 3-22 the registral span suddenly increases by 22 semitones. The upper and lower boundaries cycle at different speeds, functioning like differently sized gears. Different textures are created when the two boundaries come in and out of alignment with one another. They create tight registral spans when the two are in phase and wide registral spans when they are out of phase.

Figure 3-21 Outer boundaries. Each dyad represents the upper and lower pitch boundaries for a single interval cycle. The upper boundary ascends by whole tone towards B6 and is reset to C5 whenever it would exceed B6. The lower boundary moves by semitone towards G#5 and is reset to C4 whenever it would exceed G#5. The boundary values are swapped if the lower boundary surpasses the upper boundary.

Figure 3-22 'Raw' recursive patterns in *A Twisted Pair* from mm. 54-116. The onset of each new cycle is labelled and the outer pitch boundaries are notated on the upper staff. The excerpt applies the octave-shifting boundary method. Note that pitches will sometimes exceed the boundary limits when the distance between upper and lower boundaries is less than an octave.

Figure 3-22 (continued)

Figure 3-22 (continued)

Isomorphic roots

The (1,1,2)-cycle, used prominently in mm. 1-202 of *In Strange Lines and Distances,* is embedded as a root progression for the interval cycles in the three recursive sections. This creates another layer of self-similarity. The top staff in Figure 3-23 shows the complete root progression that I used to generate the 'raw' interval cycles. The progression is a sequence of four (1,1,2)-cycles with a slight deviation from the cycle at the end of the first three legs of the sequence (circled in blue). Each pitch in the progression represents the starting pitch of an interval cycle. The bottom staff shows the first five interval cycles transposed according to the pitches in the root progression. These five cycles appear in movement 2 at m. 203 and root progression continues as the section unfolds. Whenever the *nb.cycle* function reaches the end of the progression it simply jumps back to start and repeats the pattern in a continuous loop.

I added the deviations (circled in the figure) at the ends of the (1,1,2)-cycles in the progression in order to create a smooth transition between the adjacent (1,1,2)-cycles. Rather than suddenly leaping downwards by a major seventh, the deviation alters the endings from three large leaps to three smaller hops.

Some deviations from the root progression occur in the score. For example, the keyboard entrance on A in at m. 230 in movement 2 precedes the percussion entrance on D in m. 239. These two pitches are not adjacent in the root progression. The discontinuity reveals a seam where I cut out a portion of the 'raw' recursive process in order to prevent the process from becoming overly static or predictable. At this point, in order to allow the climax in m. 237 to resolve, I changed from loud notes, short durations, and high attack density to quiet notes, longer durations, and lower attack density.

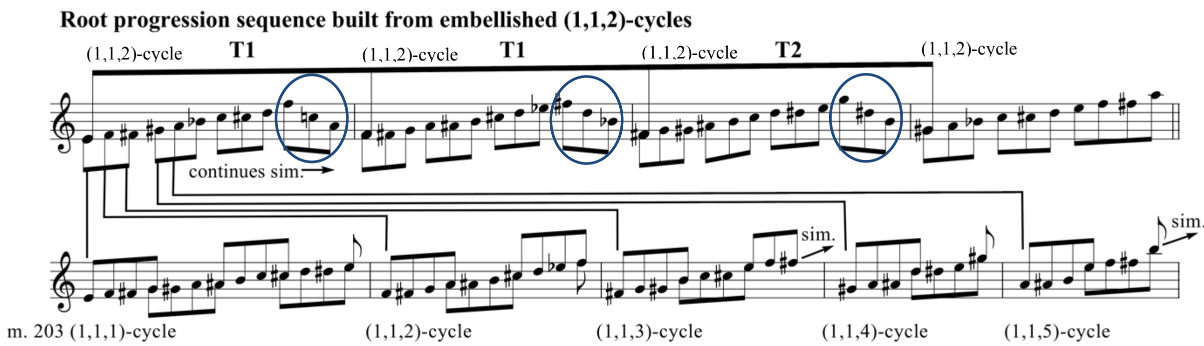

Figure 3-23, The root progression for the recursive materials (upper staff) is a modified sequence of four (1,1,2)-cycles. These four cycles are themselves related by T1, T1, and T2, indicated by beaming the first note of each cycle in the upper staff. The bottom staff shows the first five interval cycles transposed according to the root progression (see m. 203 in movement 2 in the score).

3.2.4 Cycling wave rhythms

In movement 2 from mm. 203-273, *nb.cycle* uses a sine function to set the onset times for the pitches in each interval cycle to create a wave-like acceleration and deceleration. The period of the sine wave matches the length of the interval cycle so that the attack density smoothly increases and decreases over the course of a cycle. The phase slightly increases from cycle to cycle so that the point of highest attack density rotates forward slowly over the course of the section. For example, compare the rhythm of the tenor in Figure 3-24 to the keyboard in Figure 3-25.

Figure 3-24, *In Strange Lines and Distances*, mm. 203. Sinusoidal rhythmic patterns are represented using feathered beaming.

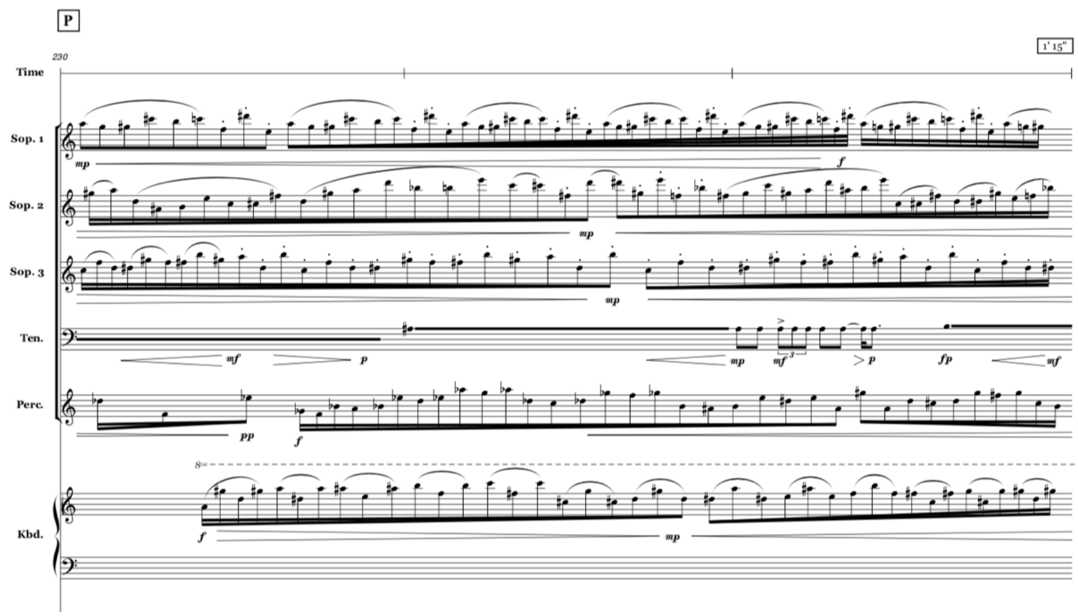

Figure 3-25 The phase values for the sinusoidal rhythms have rotated forward throughout the section. The point of lowest attack density is now at the middle of each cycle.

Nb.cycle also sets the entrance time for the cycles in movement 2. After the initial cycle, each cycle enters when the preceding cycle is 40% complete. Since the cycles can vary widely in length (see Figure 3-10 above), the 40% rule means that several short cycles can occur before a long cycle has completed. As a result, the definition between the cycles smears somewhat and creates a wave-like texture. At times I extended the length of some of the cycles to increase the amount of overlap and increase the tension within the material to create a sense of climax. For example, in mm. 230-236 I repeated notes within some cycles to extend their length and increase the number of simultaneous layers in the texture. This further emphasized the high intensity that was already present in the raw material.

The process as a whole is made up of many different cycles of varying length and varying parameters. The cycles controlling the interval content, root pitches, registral span, rhythm, and harmonic rhythm all shift in and out of phase with one another. While composing, I ran the function with various different settings which generated a variety of textures and patterns that I could then evaluate, compare, and revise.

The layers of cycles create a sense of a large-scale circularity within the raw data. Different kinds of textures recur but each recurrence is somewhat different from the previous iteration. I identified recurring patterns within the raw data and extracted excerpts to further shape and include in the finished composition. The excerpts represent a microcosm of the total recursive process. In chapter 4, I compare and contrast the three recursive sections to show how the sections reflect one another within the large-scale form.

3.3 Swarms

Swarm algorithms simulate the group dynamics of swarms of insects and flocks of birds. A swarm algorithm implements a collection of ‘agents’, where each agent represents a creature in the swarm and obeys a set of simple rules regarding how it will move in relation to the rest of the swarm and in relation to its environment. For example, Craig Reynold’s pioneering “boids” (or birdoid objects) follow rules controlling how strongly to avoid colliding with other boids in the flock, how strongly to gravitate towards the center of mass of the flock, and how closely to align to the average velocity of the flock (1987). Through these rules, the agents create dynamic patterns that emerge from their interactions as a group. Video examples of murmurations occurring in nature can be found in many online news sources such as the BBCs footage of starling murmurations in Israel ("Israel starling murmuration captured on video" 2015). Appendix C also includes short video examples of simulated swarms and Figure 3-26 shows two examples of swarm algorithms rendered as visual art. In both figures, the lines trace the paths of individual agents over time.

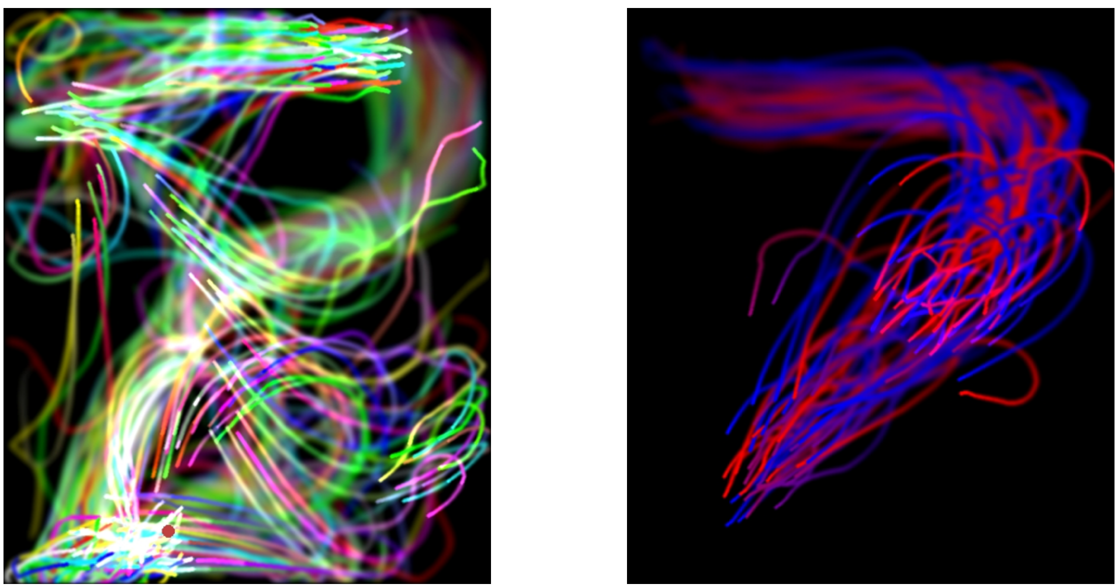

a. Carnival of Swarms **b. Swarm Swing**

Figure 3-26 Swarm art by Christian Jacob and Gerald Hushlack. (Jacob et. al)

Musical examples of swarm algorithms include spatialization, granular synthesis, and MIDI improvisation. David Kim-Boyle has mapped the positions of individual boids to spatialization parameters both for individual voices in a granular synthesizer and individual bins in an fft analysis/resynthesis (2006). Tim Blackwell has created a variety of musical swarm applications for interactive improvisation, and granular synthesis.

In Blackwell's *Swarm Music* and *Swarm Granulation* systems he creates an interactive improvisatory environment in which the dimensions within virtual spaces map to various musical parameters (2007, 207). The agent positions in these virtual spaces determine elements such as note duration, pitch, and dynamic in the resulting improvisation. The systems also listen to human improvisers during performance. Aspects of the incoming audio, such as loudness or pitch, set 'attractors' that the agents will gravitate towards within swarm spaces. The output then begins to reflect the performer's input.

In my own projects, I've used swarm algorithms to control spatialization, granular synthesis, dynamic stochastic synthesis, and soundfile playback. In each of these situations, I used the high-level parameters to manipulate general characteristics of the swarm which, in turn, controlled parameters for the various sound generation or spatialization tools. By changing high-level characteristics, I can easily interpolate between contrasting swarm states. My goal in *Through a Window* was to translate these high-level controls to notated music.

Table 3-1 Swarm rules and musical mapping

Swarm Parameter	Definition	Notated representation
Gravity point (or attractor)	A point in virtual space that the agents gravitate towards.	Average pitch height and dynamic.
Gravity strength	Strength of the agents' attraction to the gravity point.	Registral span and dynamic range.
Separation	Degree to which agents avoid one another.	Attack density and articulation.
Alignment	Degree to which an agent's velocity is influenced by the overall velocity of the swarm.	Similarity of pitch trajectories between parts.
Coherence	The tendency for the swarm to head towards the center of mass	Registral span.

Three sections in *Through a Window* are inspired by swarm algorithms. In these sections, each performer represents an agent in the swarm. The notated music depicts how agents might react when I manipulate high-level controls such as coherence, alignment, and separation. Table 3-1 provides a list of different swarm rules and describes how the rules were mapped to notation in the score. I did not attempt to transcribe the output from a swarm algorithm exactly but instead looked at general characteristics and intuitively translated these characteristics to musical texture. I will briefly describe how I realized the swarm parameters in movement 3 from mm. 38-40 (Figure 3-27):

First, the swarm becomes attracted to a gravity point associated with a low pitch register and soft dynamic. The agents slowly descend in pitch and dynamic as they head towards the gravity point. While this occurs, the coherence parameter is reduced and the registral span widens slightly; the agents are no longer as strongly attracted to the average pitch range. The separation parameter is also reduced and the agents become increasingly staccato and sparse to avoid "colliding". Once the swarm reaches the gravity point, the gravity strength is inverted and the agents flee to the outer pitch limits. The alignment value is also reduced and as a result the agents flee in both directions away from the gravity point.

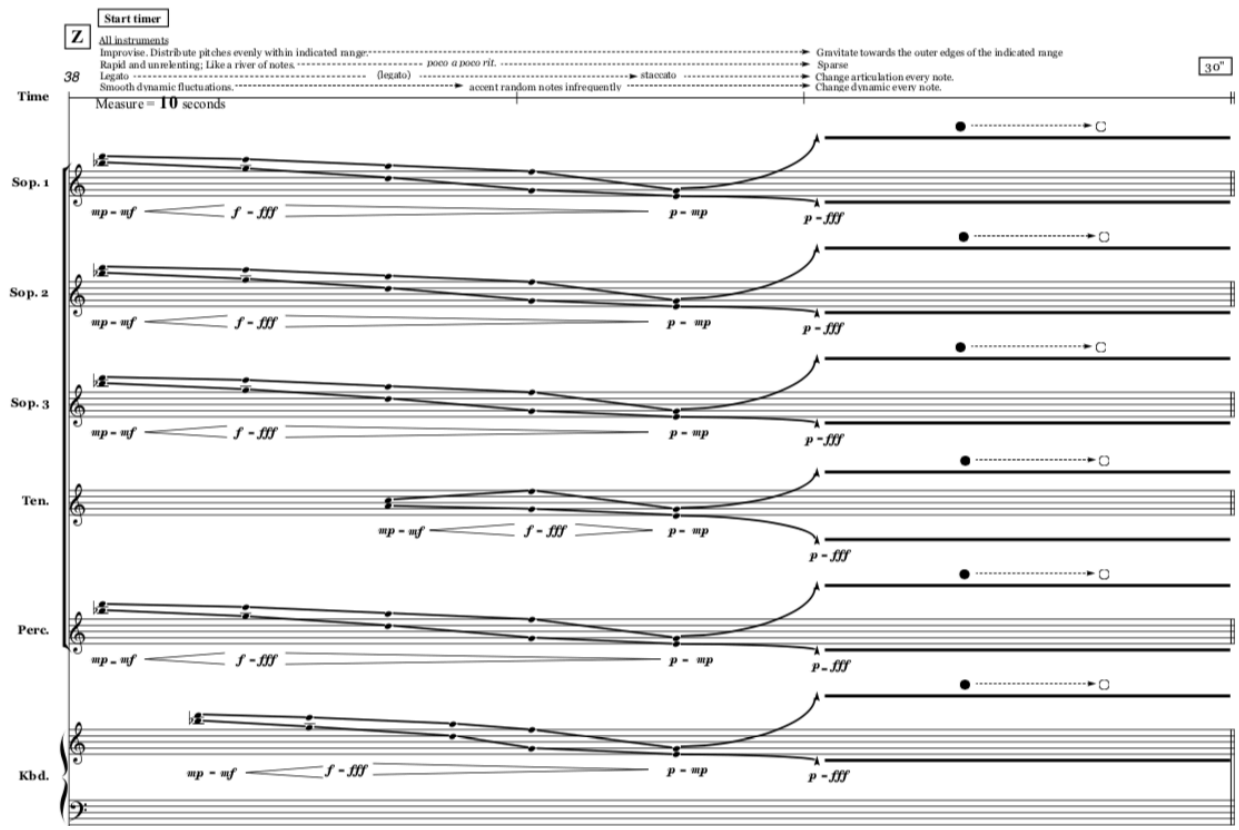

Figure 3-27 ***A Twisted Pair*** **mm. 38-40**

3.4 Evolutionary algorithms

The process of natural selection in nature relies on both large population sizes and large numbers of generations. Random deviations in the genetic code of any organism may be beneficial or detrimental to the organism as an individual; however, as these genes are replicated and passed to new generations, the environment in which the population exists causes desirable traits to be amplified and replicated more frequently while undesirable traits are attenuated. As a result, the population as a whole becomes increasingly well-suited to its environment.

Evolutionary algorithms loosely model the process of natural selection to solve optimization and search-based problems. Applications for evolutionary algorithms vary widely, encompassing both scientific and creative disciplines. Well-known examples of evolutionary computing include the evolved antennas for NASA's "Space Technology 5 mission" (Hornby et. al), Karl Sim's *Evolving Virtual Creatures* (1994), and Richard Dawkin's "weasel program" (46-50).

In music, composers and programmers have applied evolutionary algorithms to a variety of tasks including composition (Degazio 1997; Holbrook; Waschka), automated transcription (Reis et. al), timbre analysis (Horner et. al), improvisation (Biles), and sound synthesis (Magnus; Dahlstadt).

While planning *Through a Window*, I designed a small collection of genetic algorithms, a class of evolutionary algorithm, for transitioning between contrasting musical characters. Each algorithm receives a seed and target input from the user and evolves from the seed state to the target state over the course of multiple generations. The program outputs each generation so that the evolutionary process can be transcribed and performed. In performance, listeners hear the population of musical fragments slowly evolving in time.

I approached evolutionary algorithms in particular because I expected that the process of evolution would create a balance between continuity and randomness. Through evolution, I would be able to gradually transition between opposing musical states in a manner that felt focused but not overly predictable. The processes I wrote have clear origins and destinations but the specific route from one point to the other is not a straight line. Running the algorithm multiple times with the same seed and target would create similar but unique paths between the two states.

In the following sections, I provide a short overview of genetic algorithms, introduce some existing compositional examples, and detail how I used two specific algorithms to compose sections in *Through a Window*.

3.4.1 Rhythm Etudes

In this section, I use a simple application, *RhythmGen*, as a brief introduction to genetic algorithms. *RhythmGen* receives two rhythms from a user and slowly evolves from one rhythm to the other using the following steps:

1. Initialization

Evolution requires a population in order to function. In a genetic algorithm, an initial population can be generated randomly, pseudo-randomly, or can be supplied by the user. *RhythmGen* receives a single 'seed' rhythm from the user and repeatedly duplicates this rhythm to create a population of rhythms. Each population member, called a "chromosome," is also randomly altered in some slight manner to add diversity to the initial population (Figure 3-28). The population size will impact the speed at which *RhythmGen* converges on the target rhythm. Larger population size can create more diversity and generally converge more quickly. In practice, population sizes will be orders of magnitude larger than the six-chromosome population in Figure 3-28.

Rhythms are represented internally in *RhythmGen* as arrays of binary values where a "1" represents a note onset and a "0" represents a rest. The internal binary representation is called the "genotype" and the notated representation is called the "phenotype".

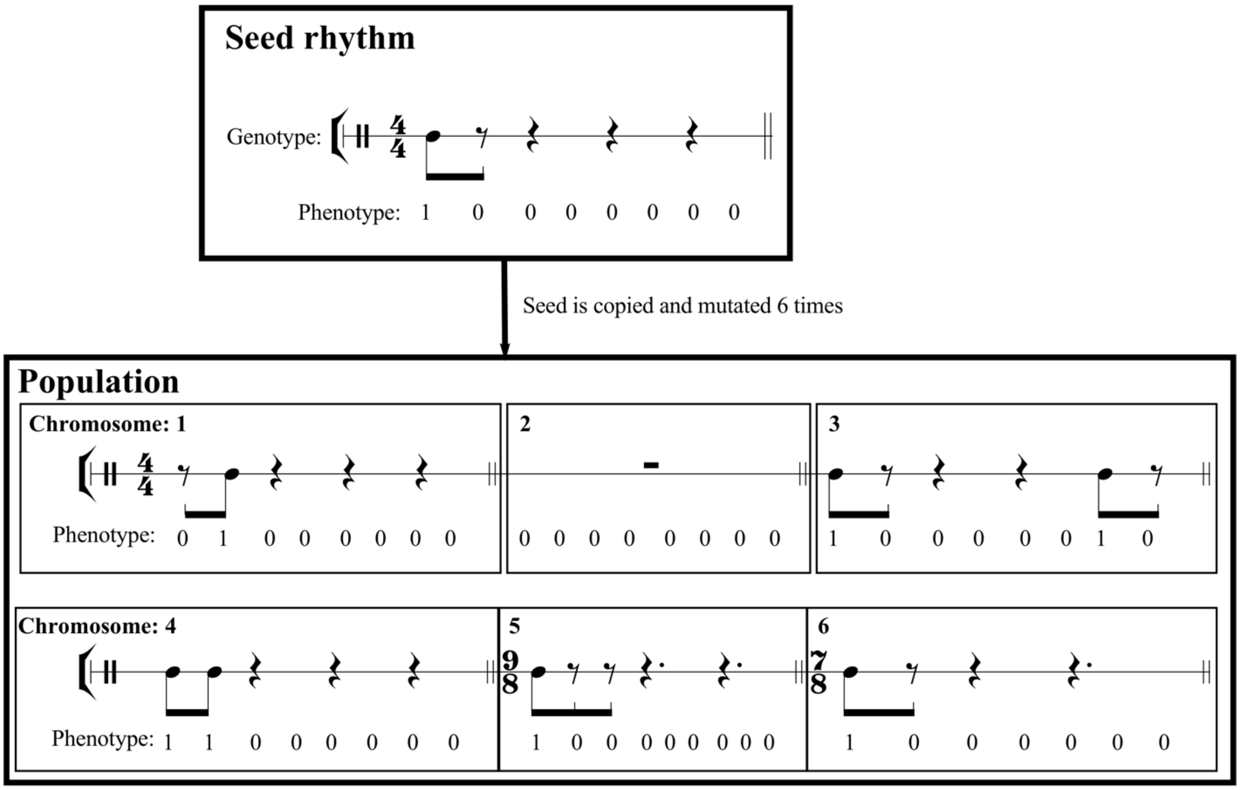

Figure 3-28 Initial population generated from the seed rhythm.

2. Fitness

After initialization, the fitness function determines how well each chromosome meets the criteria defined by the programmer and scores the chromosomes accordingly. In *RhythmGen*, the fitness of each chromosome is evaluated simply by tallying the number of binary digits that don't match the binary representation of the target rhythm (Figure 3-29). The chromosomes that return the lowest values are the fittest.

Target:	0	1	0	1	1	1	0	1	1	1	1
Chromosome 1:	0	1	0	0	0	0	0	0	0		

Figure 3-29 Fitness function. Chromosome 1 fitness = 4.

3. Selection and Crossover

The crossover function creates a new generation of hybrid 'children' by selecting chromosomes to act as 'parents' and breeding a 'child' by stitching together portions of data from both parents (Figure 3-30). A typical crossover function concatenates the beginning portion of parent 1 with the ending portion of parent 2. The crossover point is the point at which the concatenation occurs.

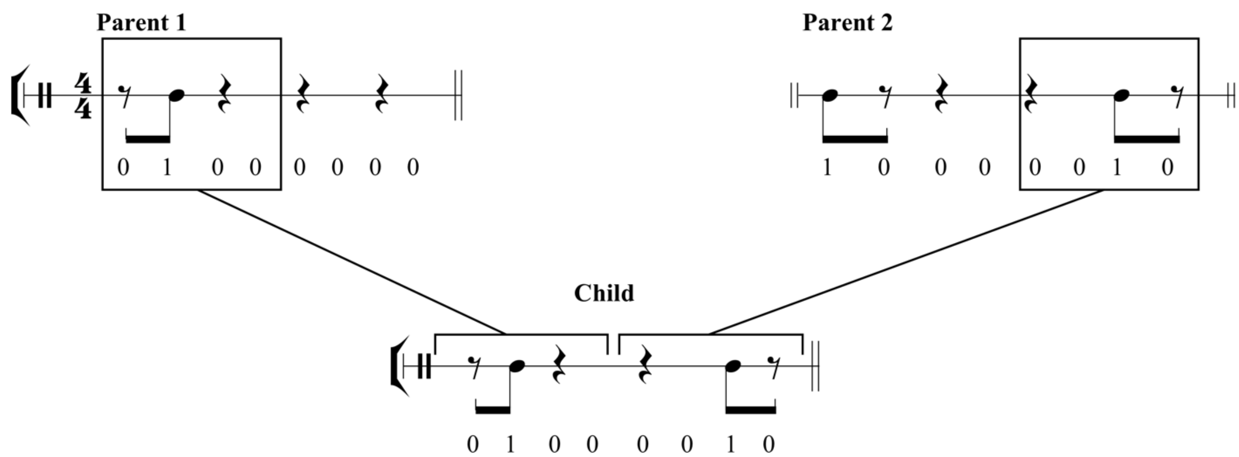

Figure 3-30 Crossover function. Crossover point = 50%.

RhythmGen selects chromosomes as parents based on a probability curve in which highly fit chromosomes are more likely to breed than unfit chromosomes. The number of children born each generation is set to half of the population size. Unfit chromosomes are removed and replaced by the offspring of the fitter chromosomes.

4. Mutation

Each child generated during the crossover function is subjected to random mutations, emulating the random mutations found in nature. Mutations create diversity within the population as a whole. *RhythmGen* randomly selects one of the four possible mutations listed in Figure 3-31.

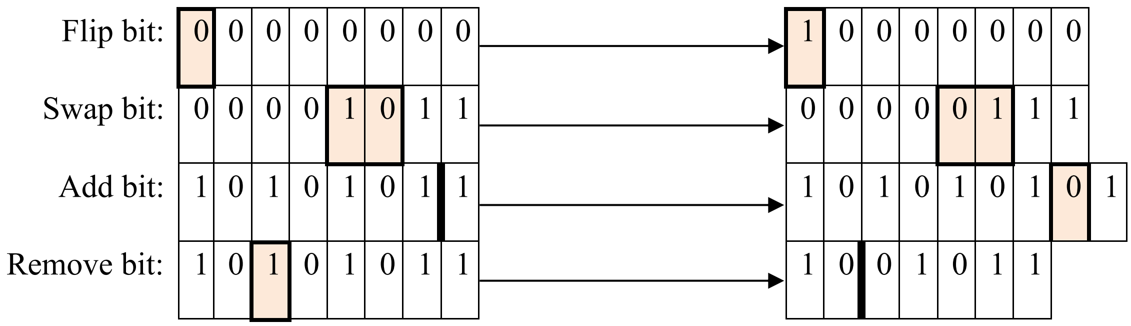

Figure 3-31 Types of mutations in *RhythmGen*

5. Termination

As the fitness, crossover, and mutation steps repeat, the population converges on the target rhythm. In more complex applications, a perfectly optimal solution may not exist. In these cases, it is helpful to look for solutions within a certain reasonable threshold of the desired outcome and to terminate the program if the population doesn't converge on a suitable solution after a certain number of generations.

Ken Fields, Neal Anderson, and I rehearsed *RhythmGen* over the network in November, 2017 using the Artsmesh software to connect our studios located at the Central Conservatory of Music Beijing, Indiana University Purdue, and the University of Calgary (Syneme 2017). In rehearsal, I created one instance of *RhythmGen* for each performer and provided the same seed and target rhythms to each instance. At the end of each generation, the three instances of *RhythmGen* would each output one of their most well-fit rhythms for the corresponding assigned performer. A Max patch displayed this rhythm for the performer using the Bach notation library (Agostini and Ghisi 2015). Since the instances of *RhythmGen* ran independently from one another, each performer received unique rhythms and followed a unique path between a shared origin and destination.

I used *RhythmGen* as a preliminary experiment to assess if the evolutionary process would be worth pursuing in *Through a Window*. Although *RhythmGen* is simplistic, the effect that it created when transitioning between states of low and high attack density was close enough to my expectations that I decided to develop the approach further to create a more nuanced musical process. This lead to *GritGen*, described in section 3.4.4, which I used to compose movement 1 from 1:12-3:10 and movement 3 from mm. 41-53.

3.4.2 Context in music

In *RhythmGen*, fitness is very easy to calculate because the correct target rhythm is supplied in advance. Fitness may be much more difficult to quantify in more complex situations involving multiple musical parameters. Decisions that could be made immediately and intuitively by a human composer about the musicality of a fragment may be quite complex to fully formalize in a way that can be quantified and coded.

Rather than creating a fixed, automatic fitness function, projects can instead adopt an interactive approach to fitness. In an interactive approach, a user manually ranks the chromosomes for fitness. The benefit of this approach is that it can allow musicians to apply their intuition directly. A disadvantage is that the process may be very time-consuming and will require a small population size.

Al Bile's *GenJam* program, short for Genetic Jammer, is a well-known example that uses an interactive approach to fitness to evolve a database of jazz riffs for improvising over a set of predefined chord changes (2007). Biles implements a voting mechanism for determining fitness.

For every generation, Biles auditions each population member and assigns it a binary fitness score. The chromosome is either good or bad. This approach allows Biles to apply his own musical judgement to the evolutionary process.

Rodney Waschka II applies evolutionary algorithms in several compositions ranging from solo works to orchestral compositions. Rather than approaching genetic algorithms as a tool to solve a particular musical problem, Waschka translates the evolutionary process itself into musical forms.

In these works, Waschka uses relatively small population sizes. Chromosomes consist of single measures of music and each chromosome is performed once every generation. The mutation function is intentionally sparse and simple. Waschka's approach uses a formal trajectory that leads from diversity to uniformity. There is no fitness function. Instead, the population members breed randomly. With each new generation, the population becomes increasingly 'inbred' until, over the course of several minutes it converges on a single musical motive and the work ends.

Waschka demonstrates how interdisciplinary thought can lead to unique insights in a creative discipline. Waschka does not use technology as a means of simplifying an existing problem or of composing in a pre-existing style, but instead applies the algorithm itself as a dynamic musical process.

Cristyn Magnus also adopts evolution as a musical process. However, rather than applying evolutionary algorithms at the symbolic level, Magnus applies the algorithm directly to groups of samples in the audio domain to create a population of sonic 'critters' (30). The critters exist within a "virtual ecology" that takes the form of a one-dimensional ring-shaped world. The critters have positions in this world as well as rules governing their movements. Critters are only able to breed with other critters that are within a certain proximity to themselves within the ring-world. Magnus assesses the fitness of the critters by comparing their similarity to stored waveforms (42). In sound installations, the waveforms can change over time or be supplied by participants using a microphone.

Waschka and Magnus use differing approaches for representing the evolutionary process. Waschka adopts a directed, goal-based process while Magnus encourages biodiversity and designs the algorithm to create variety within the population.

3.4.3 Evolutionary algorithms in *Through a Window*

Evolutionary Harmony

The final passage in movement 2 from mm. 302-349 introduces a series of gestures based on the harmonic progression presented in Figure 3-32. The bracketed harmonies served as start and end points in an evolutionary harmonic interpolation from a "seed" harmony to a "target" harmony. A program, titled *ChordBreeder*, uses an evolutionary algorithm to aid in the composition of the unbracketed intermediary chords.

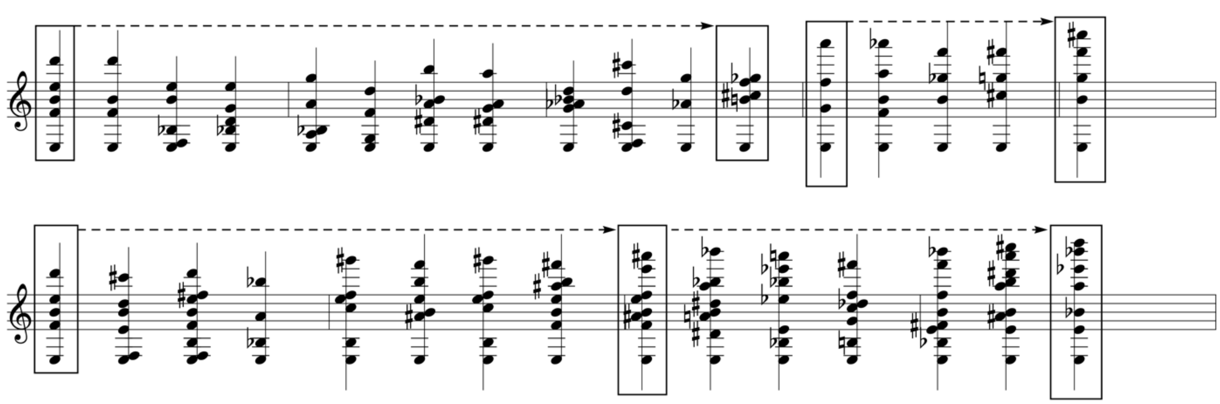

Figure 3-32, "Raw" evolved progression.

ChordBreeder is implemented in Max and uses a collection of analysis modules written in Javascript and C. The program receives a seed and target chord from the user and uses the analysis modules to assess fitness. Figure 3-33 shows *ChordBreeder's* user interface and Table 3-2 provides a brief summary of each analysis module.

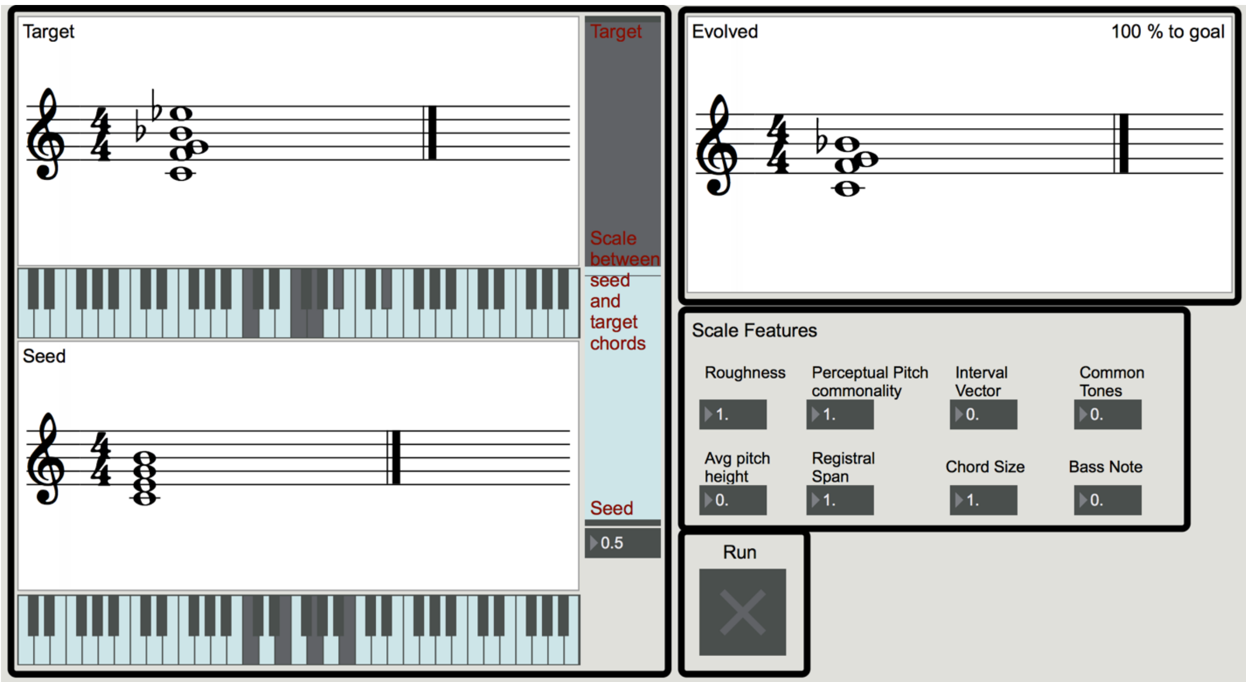

Figure 3-33 ChordBreeder interface.

Table 3-2 Fitness modules

Module	Definition
Sensory dissonance	Estimates the perceived roughness of a chord based on Sean Ferguson's sensory dissonance calculations (Ferguson, 23-28).
Perceptual pitch commonality	Estimates the perceived pitch commonalities between two chords based on Ferguson's methods (Ferguson, 29-37).
Interval vector	Returns the difference between the interval vectors of two chords.
Chord size	Returns the number of pitches.
Registral span	Returns the number of semitones between the lowest and highest pitches.
Average pitch height	Returns the average MIDI pitch value.
Bass note	Returns the semitone difference between the bass tones of two chords

In a *ChordBreeder* session, the user enters the two chords and sets weights for each analysis module so that certain modules will have more influence in determining the overall fitness of the candidate solutions than others. Once the weights are set, the evolutionary process begins. *ChordBreeder* repeatedly duplicates the seed chord, mutates each duplicate, and collects the mutated duplicates to create the initial population. After initialization, it sends the population to the analysis modules to assess fitness.

Before evolution begins, the target chord is sent to each of the analysis modules to find the desired output for each parameter. When the chromosomes are then sent to the fitness modules, each module analyzes the chromosome and returns the absolute-value difference between the chromosome's reading and the stored target reading. The differences from each module are summed and used to derive the chromosome's individual fitness score. Once the entire population is ranked for fitness, the fittest chromosomes are selected for breeding. The crossover function generates a new chromosome by taking a portion from the top of parent 1 and concatenating it with the complementary bottom portion from parent 2. The child is randomly mutated and then replaces one of the lesser fit chromosomes in the population. The evolutionary process loops until it either generates an optimal solution or is manually terminated. To compose the harmonic figures for mm. 302-349, I saved the top result from each generation, aiming for a kind of drunken harmonic walk from the seed chord towards the target chord.

Assessing fitness using analysis modules described above is much less efficient than a more traditional method such as measuring the difference between binary representations of the chromosomes and the target. However, even though this traditional method requires fewer generations and is less prone to becoming stuck on local maxima, it does not consider the perceptual similarity between the chords during the evolutionary process. Since my goal was to record the output of each generation in the process and use this data to compose a harmonic progression between the two chords, I chose to use methods of assessing fitness that resembled my own perception of harmonic similarity. The analysis modules extract several qualities that I find salient in my own listening practices such as registral span, interval content, and pitch density. However, even while using analysis modules that are capable of sorting chords into an order that resembles my own perception of similarity, I still found the evolutionary interpolation to be unsatisfactory. The output did not match the kind of drunken walk that I had expected. I concluded that my particular design for the evolutionary portion of the program was unmusical but that the analysis modules might still be effective for finding hybrid chords that combine elements of both the seed and target.

I revised *ChordBreeder* to include an interpolation slider. Instead of using the target chord to assess fitness, the user can set a percent value between the seed and target chords. An interpolation value of 0 will output the seed chord, a value of 100 will output the target chord, and a value of 50 will seek for hybrid chords halfway between the analysis outputs of both inputs. If chord 1 is three notes and chord 2 is five notes and the interpolation value is set to 50, then *Chordbreeder* will likely return four-note chords. Figure 3-34 provides a simple example in which only registral span and cardinality are used to assess the fitness of the chromosomes.

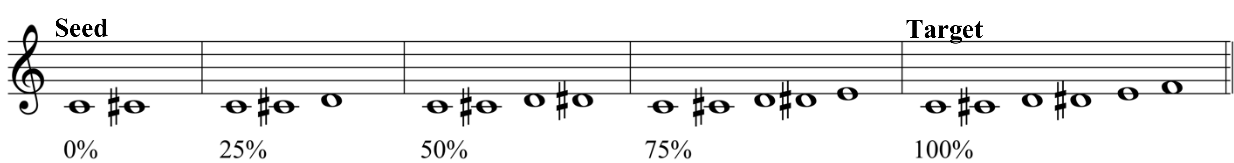

Figure 3-34 Optimized harmonies for different interpolation values using chord size and registral span to calculate fitness.

Instead of translating the evolutionary process to harmony as I had originally intended, I used the algorithm to find 'optimal' points between two chords and logged the top results for various interpolation values (Figure 3-35). Adjusting the weights of the analysis parameters can create very different harmonic results at given interpolation values. I tested various different weights for the analysis modules in the process of creating the harmonic progression in Figure 3-32. I intuitively decided to weight the sensory dissonance and interval vector modules more strongly than the other analysis modules because I felt this created a smooth transition between the seed and target sonorities. When I finally added the evolved sonorities to mm. 302-349 in movement 2, I decided to revise the sonorities to fit the texture of the section as a whole. Mm. 302-349 also contain a number of lower-frequency soundfiles. I octave-shifted many of the pitches in the evolved sonorities to keep the texture from becoming muddy when paired with these soundfiles.

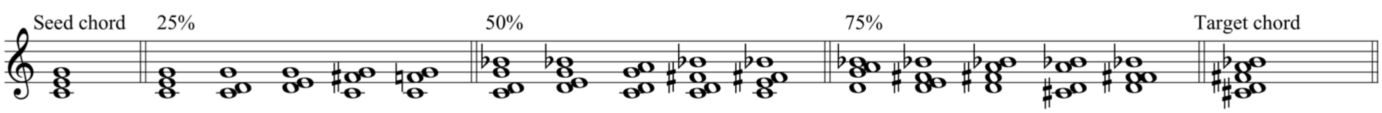

Figure 3-35 Top five most fit chromosomes at different interpolation values. Fitness is based on sensory dissonance, pitch commonality, chord size and registral span. All parameters weighted evenly.

Evolving Gestures

Unlike *ChordBreeder*, which uses evolutionary computing to solve an optimization problem, in movement 1 from 1:12-3:10 and movement 3 from mm. 41-53 I present the evolutionary model itself as a musical process. The two sections introduce populations of short musical fragments. The fragments mutate and breed to create new generations. Over the course of the sections, the fragments gradually evolve from single-notes to more complex gestures. The gradual transformation between these two states becomes the focus of the sections.

The two sections are nicknamed *Grit I* and *Grit II* and the algorithm itself is nicknamed *GritGen*. The name "grit" refers to the role of the mutation function in the evolutionary process. Mutations begin as small imperfections that accumulate and become magnified over time, leading to the emergence of complex shapes. While the mutation function involves chance and often produces results that are individually undesirable, it is an essential component in improving the population as a whole.

GritGen generates an evolving population of musical fragments. The population of fragments should seem to grow organically, resembling living organisms. This aesthetic choice informed *GritGen's* handling of musical time. *GritGen* defines the total length of each fragment as a value in seconds. It also calculates individual note durations as percent values of the fragment's total length rather than as metered beat values. The mutation function primarily uses floating point values when modifying note durations and as a result, *GritGen* naturally gravitates towards fluid, unmetered rhythms. This approach is well-suited to the networked performance environment as it does not call for a strict tempo or require tight coordination between the performers.

Figure 3-36 presents a short musical fragment and its representation in *GritGen.* The algorithm does not handle pitch directly but instead uses pitch height rankings to represent the melodic contour. These contours are mapped to pitch collections once the evolutionary process is complete.

Notated fragment:	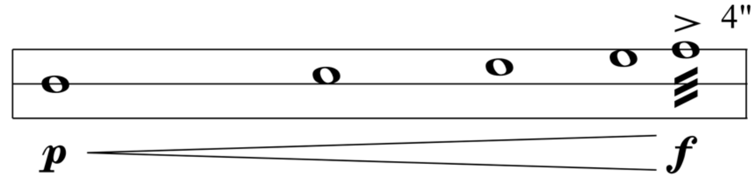				
Beats (% of total)	50	25	12.5	6.8	5.69
Pitch	0	1	2	3	4
Dynamic	Piano	→	→	→	Forte
Articulations					Accent
Other					Tremolo

Figure 3-36 The notated fragment (top) is represented as an array of note objects in *GritGen* (bottom). The program includes proportional durations for each note, a total duration for the fragment in seconds, contour rankings for each note, dynamics, articulations, and other elements such as tremolo or extended techniques.

The output from *GritGen* is relatively transparent in the score. I labelled each generation is and notated the majority each population for the musicians. The musicians freely select and perform individual fragments from the population whenever instructed to do so in their part. As the process continues, distinct gestures begin to emerge within the population. These gestures can proliferate and persist across subsequent generations. Figure 3-37 depicts a short gesture that emerged in *Grit I*, circled in blue. This gesture is not particularly well-fit on its own but emerged as a byproduct of the evolutionary process.

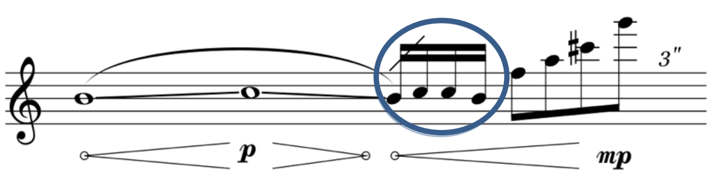

Figure 3-37 The circled gesture emerged as a byproduct of the evolutionary process and remains present throughout several generations.

Figure 3-38 diagrams the *GritGen* algorithm. First, *GritGen* generates an initial population from a single pre-composed ‘seed’ gesture. The program repeatedly copies the seed and applies mutations to each copy. The mutated copies form generation 1. In the example below, the population size is three. In *Grit I*, I set the population to 15 and in *Grit II*, I set the population size to 12. The population is then sorted by fitness, bred, and mutated to create a new generation. This new generation is handed back to the fitness function and the process continues in a loop. In both Grit I and Grit II, the process loops ten times. Figure 3-39 diagrams a portion of the ‘family tree’ for a fragment in *Grit I*, generation 6. Boxes indicate the segment taken from each parent for the following generation.

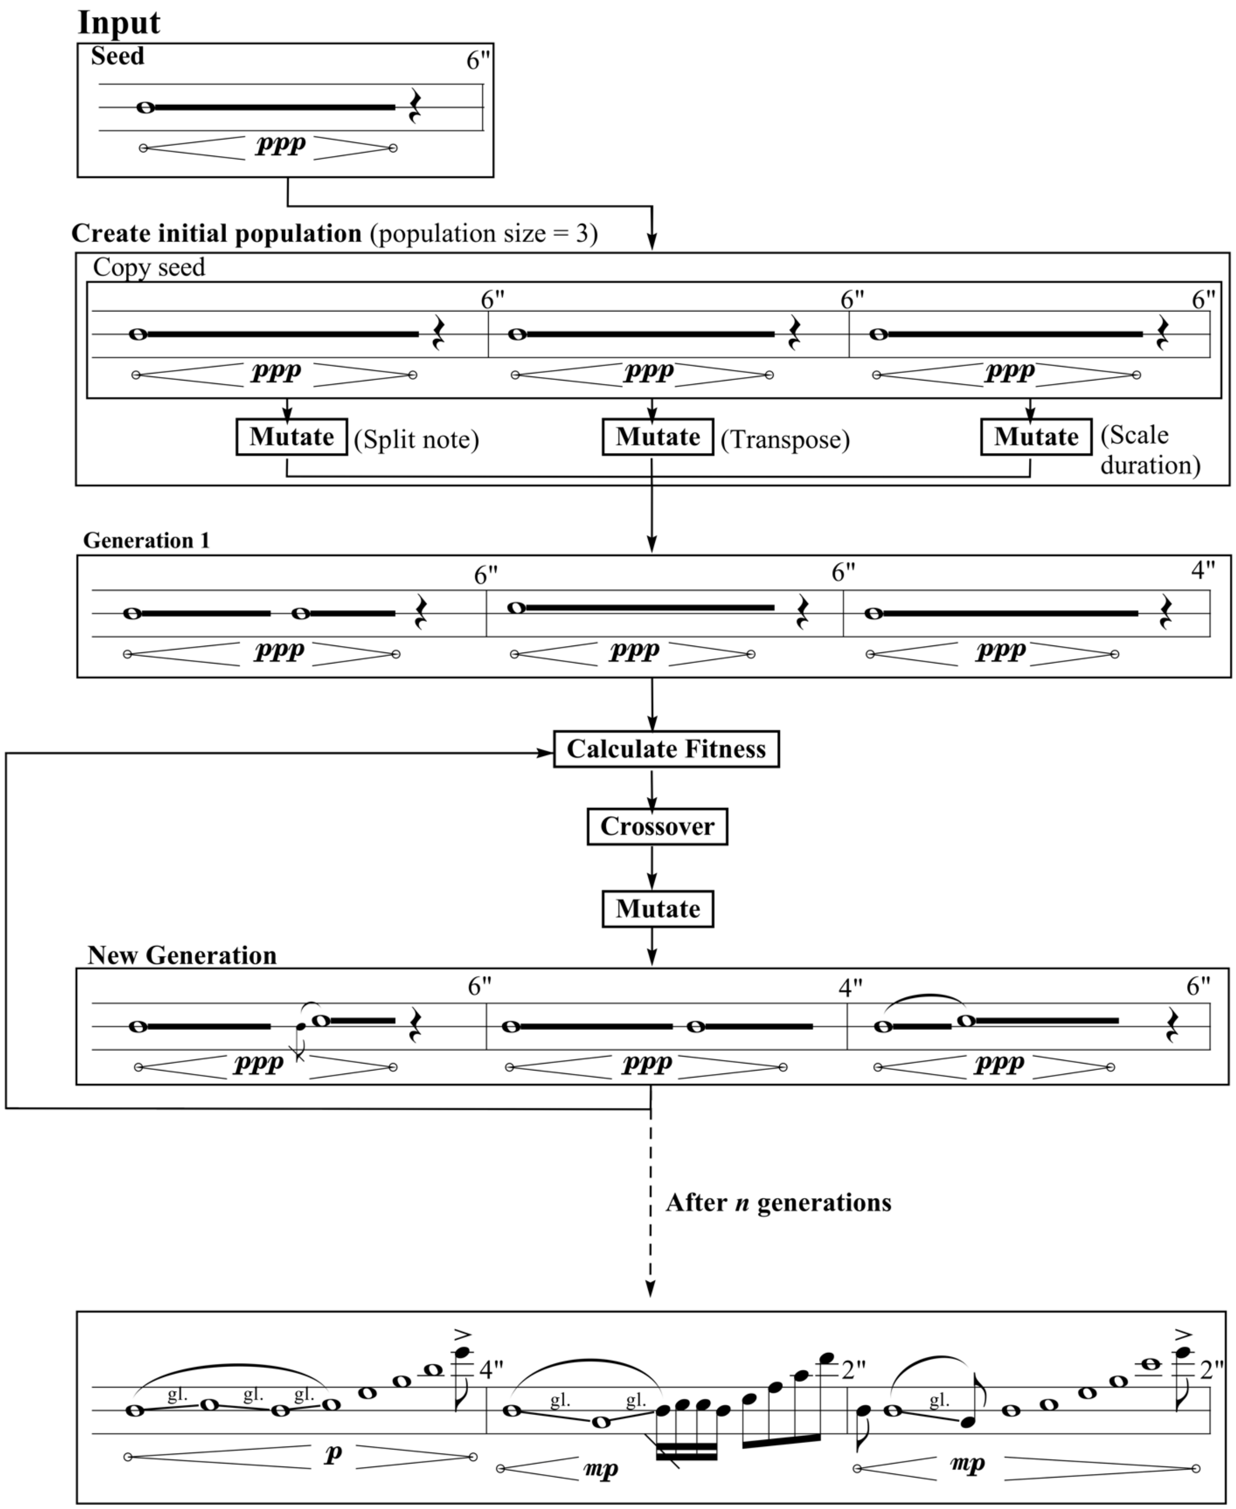

Figure 3-38 Diagram of the *GritGen* algorithm.

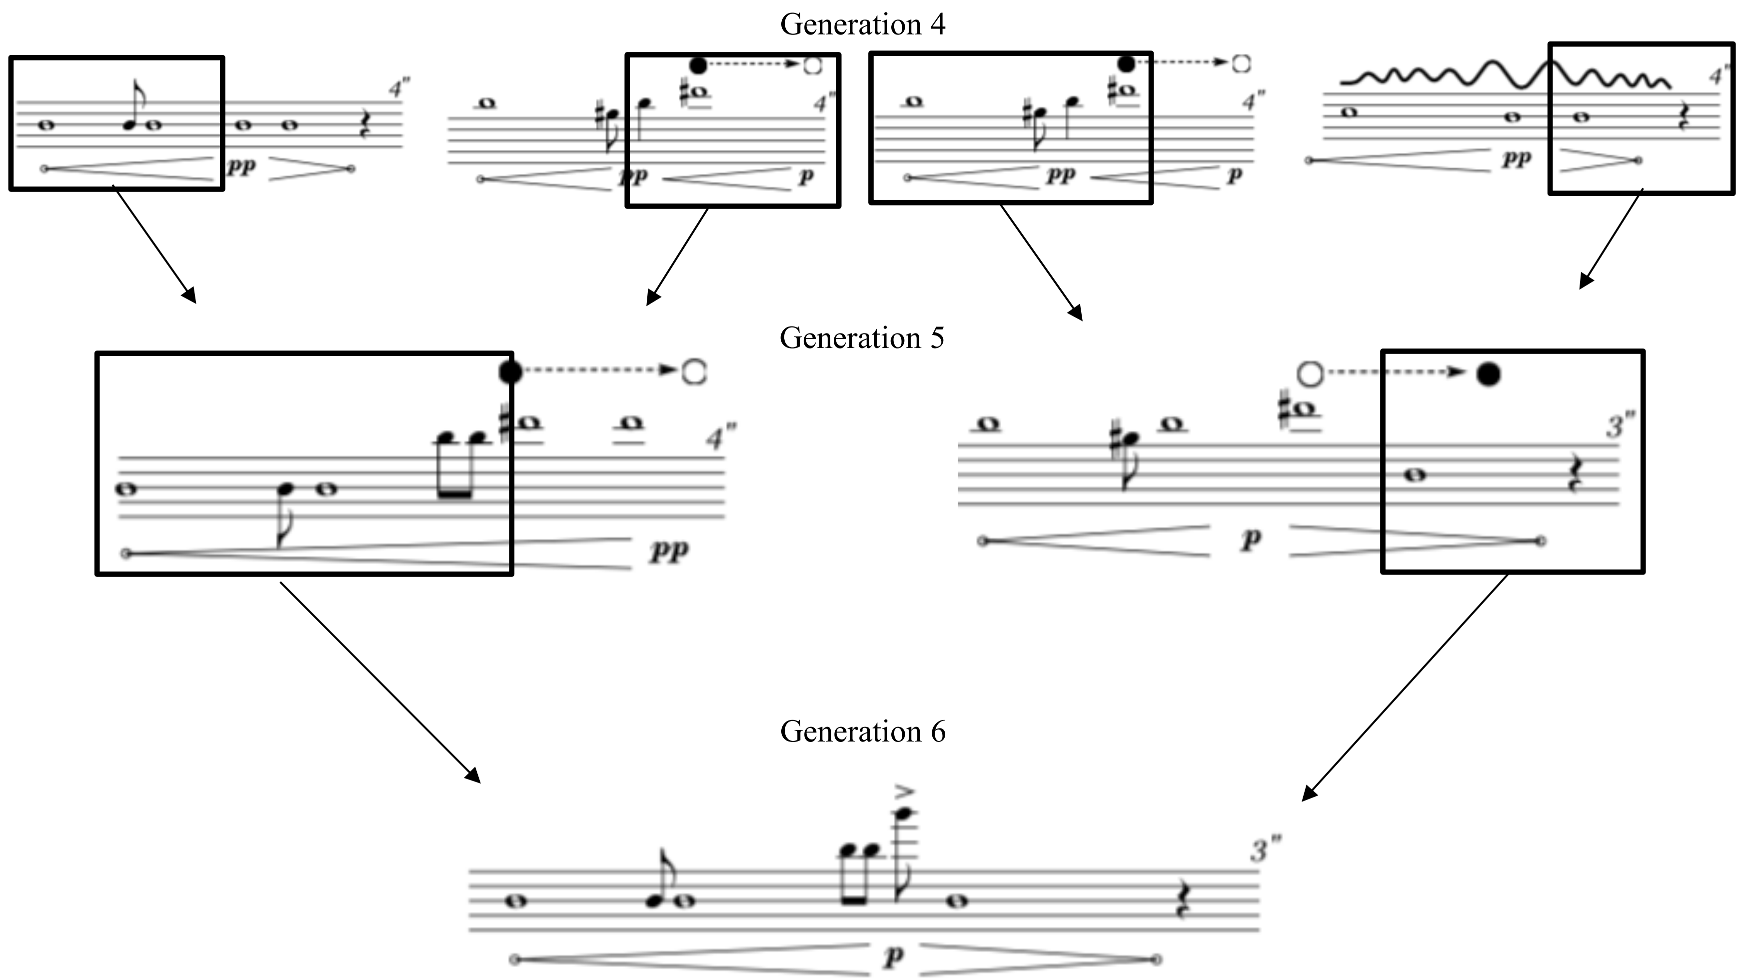

Figure 3-39 Excerpt of the family tree for a single fragment from Generation 6 in *Grit I*. Note that the mutations create some irregularity between each generation.

I composed simple, single-note seeds for *Grit I* and *Grit II* so that in early generations the listener will be more likely to perceive the mutations that *GritGen* applied to each population member (Figure 3-40). In later generations, the algorithm generates more complex fragments and the individual mutations may become less obvious.

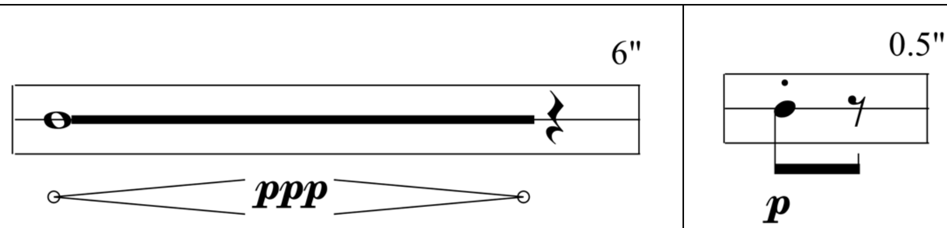

Figure 3-40, Seed fragments for Grit I (left) and Grit II (right).

Fitness

GritGen determines fitness using interactive selection (described in section 3.4.2). For each generation, I ranked each fragment based on how well its melodic contour matched a predetermined target contour (Figure 3-41). If two fragments possessed roughly equivalent contours, I instead ranked the fragments to encourage growth within the population, favouring the fragments with more attacks and more variety in pitch and rhythm. This form of selection allows me to rank the fragments in an intuitive manner rather than relying on a fixed calculation.

I wrote the algorithm on paper and used dice to perform the random selections. I kept the population sizes relatively small to help mitigate the amount of time it would take for me to evaluate the algorithm manually. The fitness bottleneck that sometimes arises from interactive selection was not an issue in *GritGen*. The target contours are simple enough that even with the small population size *GritGen* reached my desired target in ten generations.

I composed the seed and the target for both *Grit I* and *Grit II* to fit the musical context during their respective points in the composition as a whole. As a result, the two versions contrast one another. The overlapping, sustained character in *Grit I* connects with the overlapping balloon sounds of the previous section. *Grit I* evolves towards an ascending target contour that loosely reflects the ascending pitch contours that characterize much of movement 2.

I introduced *Grit II* at a moment of dissolution. At this point, the dense, chaotic series of interval cycles and swarms have dissolved into a sparse, playful dialogue. This playful dialogue slowly regains energy and evolves back into the wave-like contours that characterize the opening of the movement. The wave-like contour, which hints at the contour of the second recursive section at the beginning of movement 3, leads into the final recursive section at m. 54. The two *Grit* sections emerge from the same process and sound related to one another, but the contrasting musical parameters help keep them clearly distinct from one another.

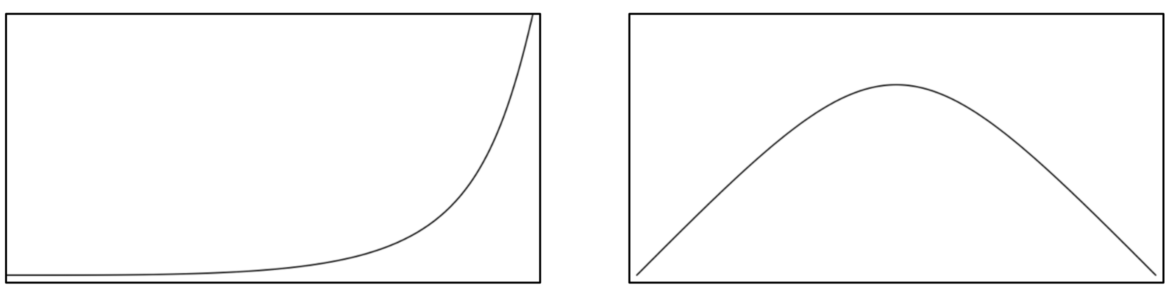

Figure 3-41 Target pitch contour for *Grit I* (left) and *Grit II* (right).

Crossover

After fitness is assessed, the ranked population is handed to the crossover function. The function chooses two parents based on a probability curve in which fit chromosomes are more likely to breed. If parent 1 has a longer duration than parent 2 then it may be possible for the crossover point to exceed the total length of parent 2. In this case, parent 1 is truncated at the crossover point. If the crossover point is positioned in the middle of a currently sounding note for both parents and the notes are different pitch values, then the notes will be split at the crossover point and articulated using either legato or a glissando. If these sounding notes are the same pitch for both parents, then the note will remain as a single onset. Since crossover can cause notes to be split, the function has the tendency to increase the overall attack density of the population.

Mutation

Mutation occurs once per generation for each population member. The mutation function involves three steps. First, it randomly chooses the position and breadth of the mutation. The function can modify a single note, a subset of the notes, or all of the notes. Next, the function randomly chooses a musical parameter to mutate and selects a mutation operation for this parameter (such as “split note”, “transpose”, or “scale duration”) as well as a mutation amount (for example, the number of semitones for a transposition).

Over multiple generations, the fitness, crossover, and mutation functions cause the population as a whole to gradually change in character. The musical parameters should appear to incrementally bend towards the target, like a plant twisting its leaves towards a light source.

3.5 Conclusion

I approached evolution, swarming, and recursion creatively to develop musical processes or textures that I could then juxtapose, vary, shape, and elaborate upon to eventually arrive at the finished material in *Through a Window*. My approach involved intuitive decisions and required repeated revisions to create processes that I felt were both engaging and fit within the context of the composition as a whole. Whenever a function didn’t match my expectations, it revealed an aspect of the algorithm that I had either overlooked or misunderstood. I could then revise my approach, furthering my understanding, to eventually find designs that generated processes or materials that better-suited my expectations. For example, I revised the evolutionary harmony generator to include the interpolation slider in order to create progressions that sounded coherent.

Adopting algorithms as musical processes felt very similar to my standard compositional practice. While composing, I generally develop families of harmonies, gestures, motives, rhythms, and processes through experimentation and repeated revisions. The algorithmically-inspired materials followed the same process of experimentation and repeated revisions to develop families of musical processes. I then further shaped the materials to bring out elements that I felt were particularly dramatic or otherwise musically engaging such as in mm. 230-236 described in section 3.2.4.

Chapter 4 Analysis

4.1 Connective fibres

In *Microsound*, Curtis Roads defines nine timescales for "dissecting" music, ranging from the infinite to the infinitesimal (3-4). In the following analysis, I consider the relationship between the macro- and meso-timescales. The macro-timescale spans the length of a single composition and the meso-timescale spans the length of musical phrases (3).

Roads uses the term "macroform" to describe the formal structure of a composition. The macroform is a "hierarchy of time scales" defined as follows:

> The uppermost level is the root symbol, representing the entire work. The root branches into a layer of macrostructure encapsulating the major parts of the piece. This second level is the form: the arrangement of the major sections of the piece. Below the level of form is a syntactic hierarchy of branches representing mesostructures that expand into the terminal level of sound objects. (12)

According to Roads, the macroform is the result of either a top-down or bottom-up process (12). The top-down approach describes macroforms that are dictated using preconceived formal plans, while the bottom-up approach describes macroforms that are the result of manipulations or elaborations on the meso-scale (13). Macroforms may also emerge from a combination of both approaches.

Through a Window combines aspects of both bottom-up and top-down approaches. The generative processes that I describe in Chapter 3 are examples of the bottom-up approach. From the top-down perspective, I structured the composition to gradually increase in intensity and complexity as the work progresses. The three recursive sections served as landmarks for me while composing and I consciously structured material to build towards or away from these sections. I didn't impose strict temporal proportions to the individual sections but instead let the processes develop until they reached a satisfying connection to the next section or concluded themselves in a natural manner. Although I have divided the work into sections based on different underlying processes, *Through a Window* should not sound like a series of contrasting concatenated sections, but should instead seem to grow and recede in an organic manner.

This section will consider some of the mesostructural characteristics that create consistency between adjacent sections. Throughout the analysis, I will reference sections by the titles used in Figure 4-1.

Figures 4-1 and 4-2 diagram *Through a Window*'s macroform. Figure 4-1 divides the work into sections corresponding to dominant underlying processes or conceptual themes. Figure 4-2 approximates the intensity levels that I aimed to evoke over the course of the composition. Here, I use the term "intensity" to describe a subjective quality that I perceive as a listener, likely created through the combination and manipulation of multiple musical parameters such as loudness, noisiness, register, attack density, and intervallic content. I suspect the perception of intensity is variable from person to person and likely also informed through experience. I use "intensity" and the intensity graphs as a means of sharing how I conceived of the formal shape of the composition. For example, Figure 4-2 shows the three movements exhibiting growth both individually and as a whole.

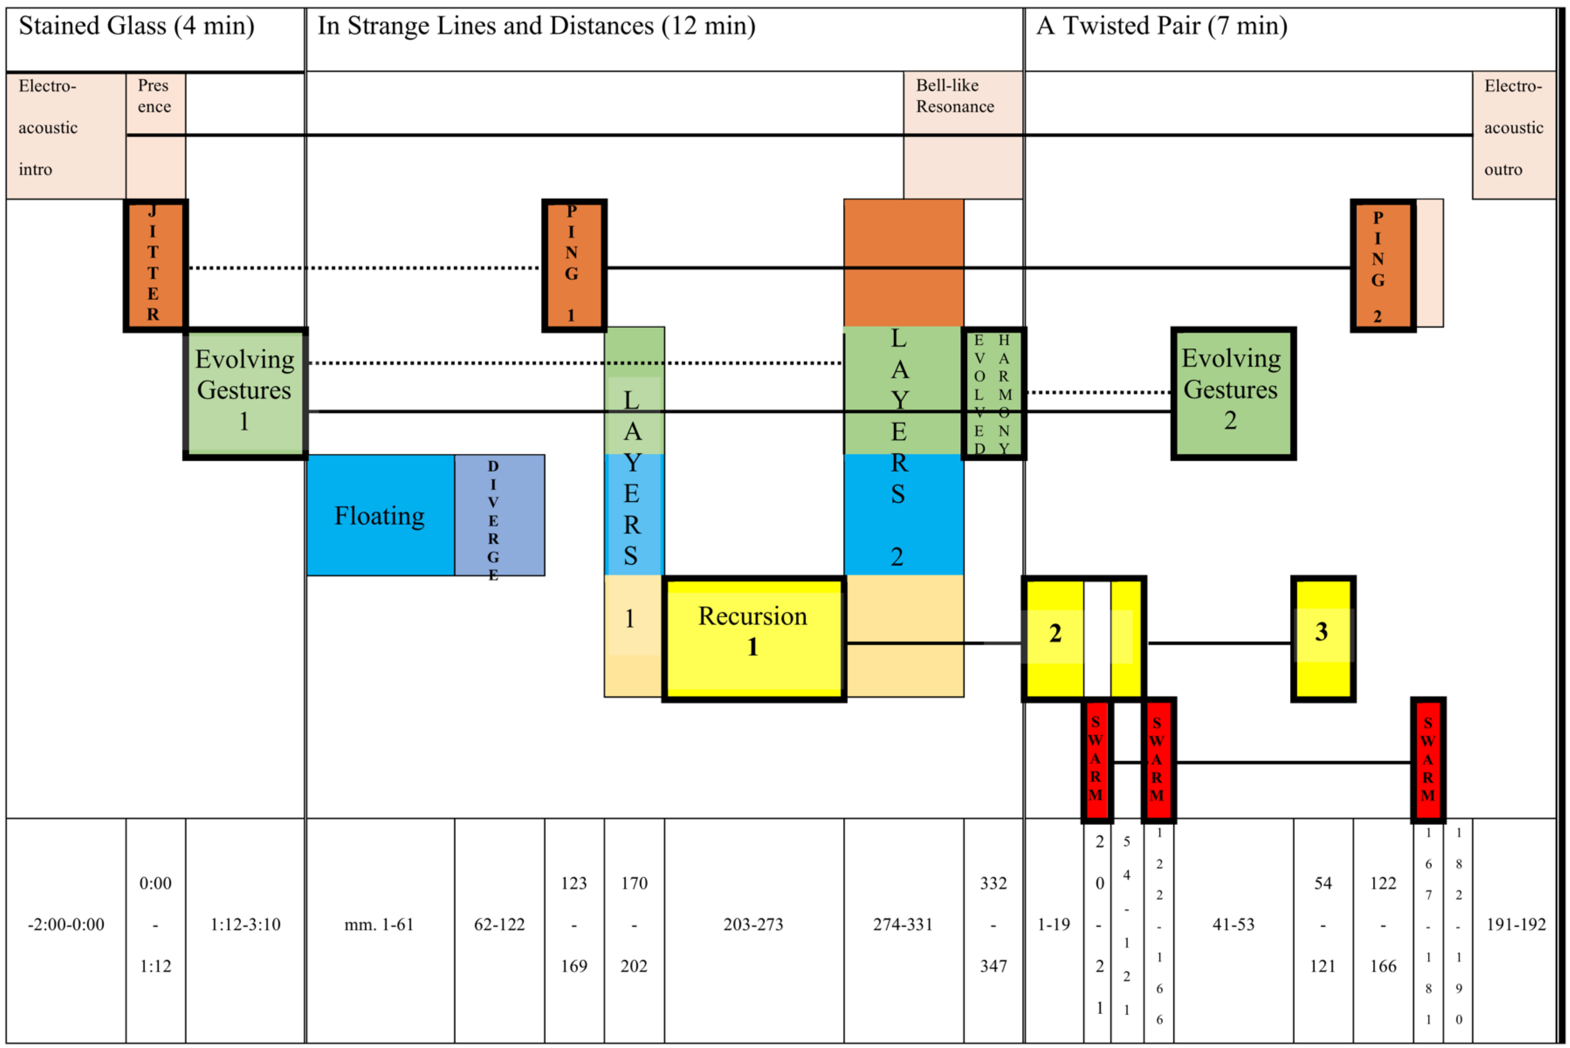

Figure 4-1 Sectional organization of processes in Through a Window

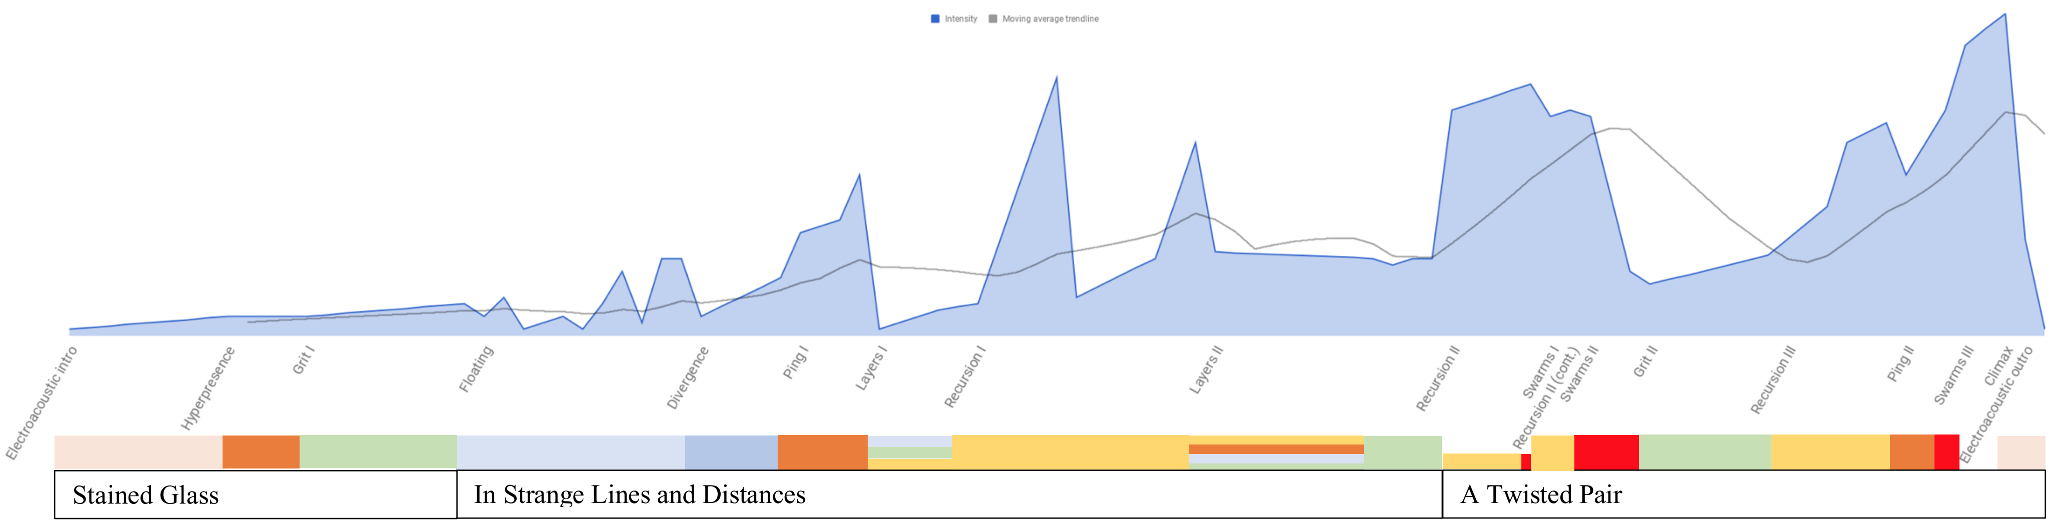

Figure 4-2 Approximate intended intensity throughout the composition. Sections are colour-coded to match Figure 4-1

The composition begins in a nascent state: a nebulous texture becomes increasingly well-defined, leading towards the opening of movement 2. The electroacoustic introduction uses layers of processed balloons to create an immersive texture somewhere between waves and hyper-present breathing. The instrumental musicians emerge from this electroacoustic texture, performing unpitched extended techniques, soft speech, and unconventional instruments such as deflating balloons. The unpitched sounds give way to soft, sustained, pitched tones in the Evolutionary Gestures I section which gradually evolves into an ascending contour setting up the opening melody of movement 2.

The opening of movement 2 shifts from the free unmetered textures of movement 1 to a metered and melodic texture with clearly-defined foreground and background elements. This creates an effect like a lens coming into focus. I composed the opening melodic material in movement 2 specifically to connect with the expanding, wave-like, rising gestures of Recursion I. The melody conceptually represents softly undulating water (Figure 4-3). The three pitches cycle and change duration in a fluid manner. The melody develops and ascends to higher pitches, representing further disturbances in the water. Over the course of the section, the registral span expands and the attack density increases until the floating material becomes the series of ascending gestures at mm. 40-44. These ascending gestures represent wakes in the water (Figure 4-4).

Figure 4-3 In Strange Lines and Distances, "floating" melody, mm. 1-2.

Figure 4-4 In Strange Lines and Distances, “wakes”, mm. 41-44.

The first 203 measures of movement 2 are primarily based on a (1,1,2)-cycle built on E. As described in section 3.2.1, the (1,1,2)-cycle generates three ‘islands’ of (012) sets related by T4. At first, the floating melody outlines (012) sets. However, as the melody develops, the intervals expand to create an augmented chord. The augmented version of the floating material persists beyond mm. 1-61 and adds a unifying element to the Divergence and Layers I sections. Figure 4-5 shows several of the pc transformations that I applied to the initial (012) set.

In the Divergence section, the lyrical and melodic quality of the ‘floating’ gesture becomes more playful and rhythmic. The section initially pairs the soprano 1 and percussion parts in rhythmic unison. However, the parts gradually diverge into two separate streams.

The wakes also resurface in later sections. The rising gestures at the conclusion to the Divergence and Ping sections are both examples of wakes. In the Ping I section, I initially composed the ascending runs as single disjointed notes, matching the raw ping data, but I soon revised the contour to embed “wakes” within the greater continuity of the piece. Figure 4-6 shows two examples of “wakes” occurring in the Ping I section.

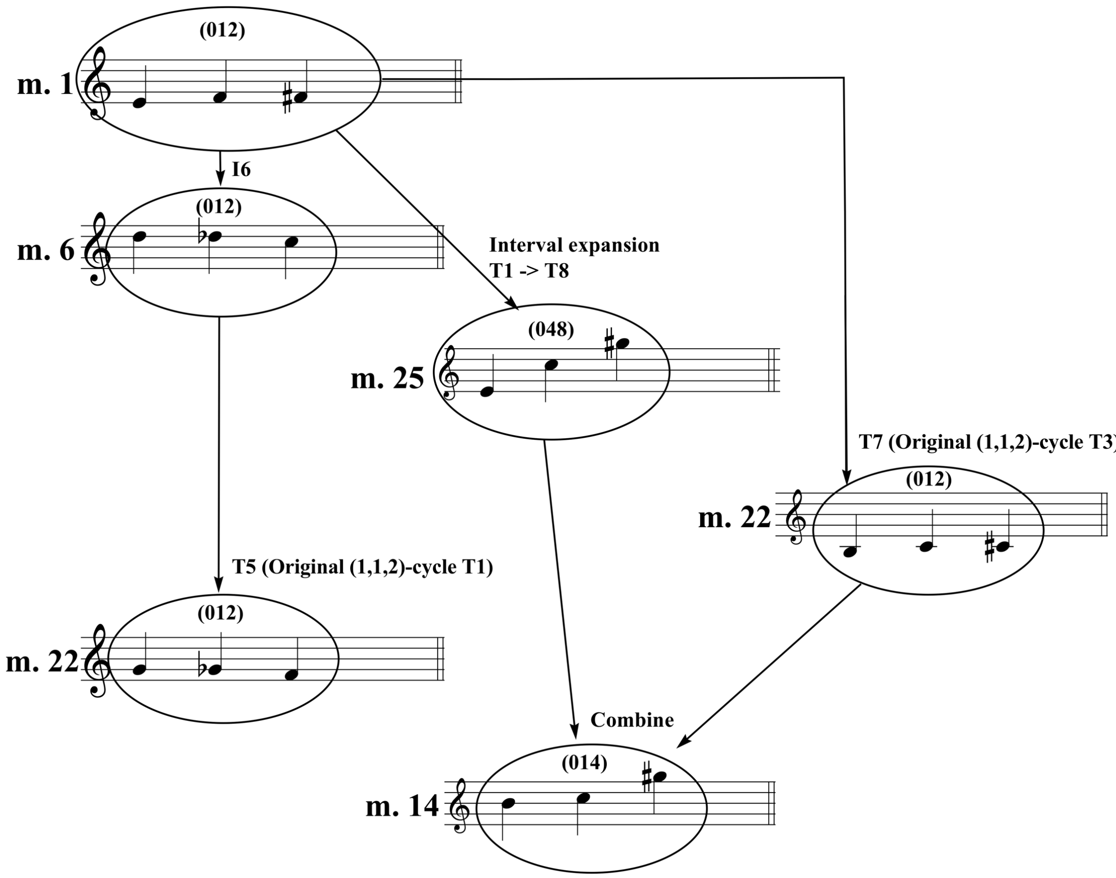

Figure 4-5 Pitch class transformations applied to the "floating" melody.

Figure 4-6 Wakes in Ping I (mm. 140-144)

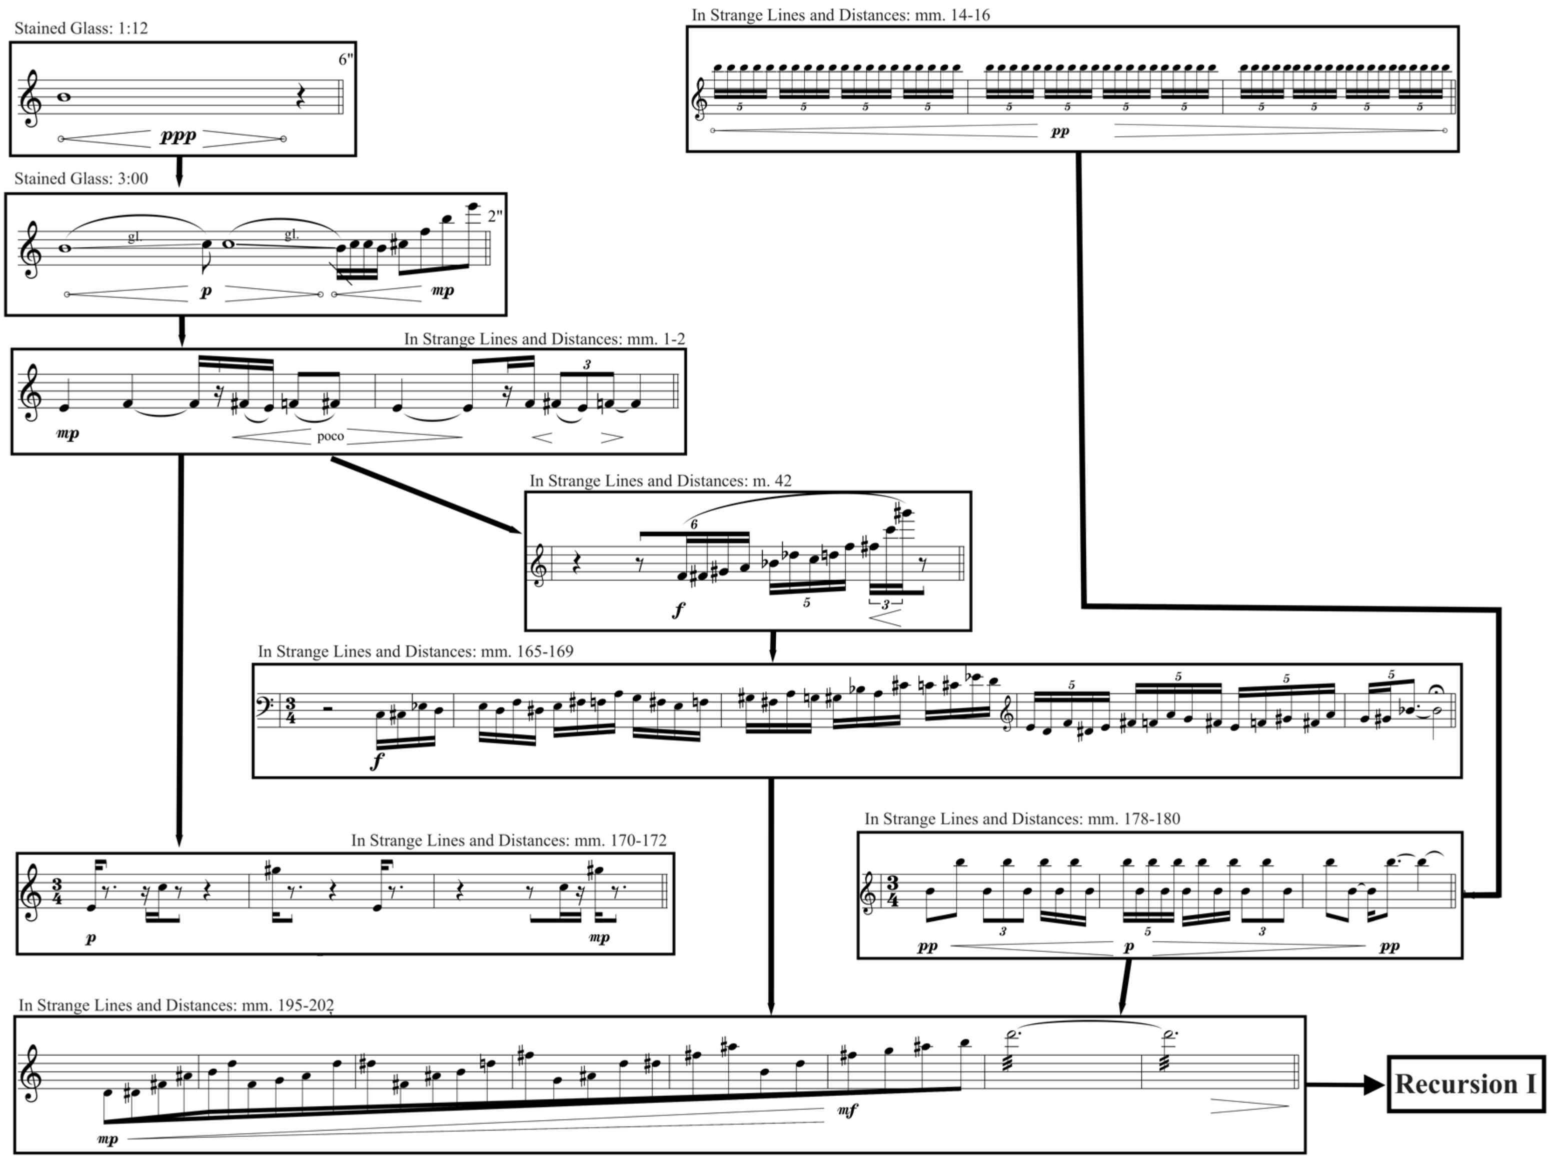

Figure 4-7 Mesostructural connections leading towards Recursion I in mm. 1-203

Figure 4-7 charts the mesostructural connections between sections, described above, in the development towards Recursion I. The excerpts are not comprehensive but rather highlight connections among the sections.

Figure 4-2 shows that related processes tend to express similar levels of intensity within the composition. The Recursion, Ping, and Swarm sections tend to be high intensity, while the

Evolutionary and Layer sections are low intensity. The Floating and Divergence sections are low to medium intensity. As the composition progresses, I introduce new processes roughly in order of intensity so that the intensity gradually increases over time. However, the intensity also ebbs when previous processes are revisited. Table 4-1 provides a brief summary of each section throughout the composition.

Table 4-1 Through a Window, summary of sections

I. *Stained Glass*		
Time:	Themes:	Context:
-2:00-0:00	- Hyperpresence - Anamorphology: Soundfiles	- Establishing the sound within the local node. - Establishing the electronics as a performer.
0:00-1:12	- Hyperpresence - Network as generator: Jitter becomes rhythm.	- Easing in the acoustic performers. - Exploring ambiguity and intimacy.
1:12-3:10	- Evolving fragments: Grit I - Anamorphology: Harmony	- Taking form. A nebulous texture becomes increasingly well-defined. - Establishing inter-nodal communication and creating liveness through musical dialogue. - Connecting to movement 2. Sustained pitches become rising gestures. - Low Kbd. notes foreshadow Kbd. clusters in Ping I, Ping II, Swarms III, and electroacoustic outro

II. *In Strange Lines and Distances*		
Time:	Themes:	Context:
mm. 1-13	Floating: - five subsections each build and rest, creating wave-like sectional structures that are microcosms of the larger form. - Pitch material and rising melody is an extended opening for Recursion I - Waves (continued) - Anamorphology: Harmony	Expository: 'Floating' material introduced.
14-21		Kbd. Solo I
22-47		- Building up foreground layers - Increasing attack density - Increasing register - Floating material develops into ascending melodic waves
48-51		Kbd. Solo 2
52-61		A more impactful and unified texture to resolve the Floating section and contrast the following section.
62-122	Divergence	- Sop. 1 presents expanded Floating material. - Percussion enters in sync with Sop. 1 and gradually diverges. - Change to a more rhythmic/playful texture. - Introduction of rhythmic 'shots' foreshadows melodic interjections in Ping I.
123-156	Ping I	- Increased tension - Rising figures connect to 'wakes' in the Floating section
157-169		- Increased number of layers and register. - Piano clusters foreshadow Swarms III
170-202	Layering of materials	- Moment of repose - Connects materials from several sections.
203-235	Recursion I	Ascending waves. Leading to a moment of maximum tension in Movement II
236-256		- Recession - Short deviation to slower material to create contrast
257-273		Second wave of high tension
274-331	Layering	- Suspension and recession of recursive material. - Intermingling of materials from all sections - Soundfiles connect to movement 1
332-347	Evolved harmony	Soundfiles act as a resonant reflection of the electroacoustic introduction.

III. *A Twisted Pair*		
Time:	Themes:	Context:
mm. 1-19	Recursion II	- Ascending and descending waves - Melodic interjections connect to Ping I - Building layers and attack density
20-21	Swarms I	- Material becomes a swarm of activity rather than independent lines.
22-37	Recursion II (continued)	- Culmination of Recursion II - Rhythmic unison (smeared when performed on the network) - Unison texture refers back to the culmination of the Floating section, mm. 52-61.
38-40	Swarms II	- Strongly contrasts the unison texture of the preceding section. However, due to the complexity of the material, the music will be perceived as a kind of sound mass rather than individual lines. - Resolves tension by interpolating from high activity to low activity.
41-53	Evolving fragments: Grit II	- Establishes a playful dialogue among the performers and computer improviser - Progressing towards an ascending and descending wave-like contour, similar to Recursion II - Gradually building momentum
54-121	Recursion III	- Recursive pitch material outlines a descending contour. - Building momentum. - Reintroduction of rhythmic shots - Concludes intervallic cycles
122-166	Ping II	- Revisits melodic materials from Ping I in a varied form
167-181	Swarms III	- Highly dissonant and impactful - Piano clusters refer back to preceding movements
182-190	Cluster and resonance	- Waves of sound, an exaggerated version of the Kbd. Onsets.
191-192	Outro	- A return to the opening - Calming down

4.2 Anamorphic perspectives on form

In section 2.3, I compared NMP to anamorphic art. In anamorphic art, the experience of a work is tied to the physical location of the viewer. This is also true in NMP where listeners experience the performance uniquely at each node. In Joshua Fineberg's *Guide to the Basic Concepts and Techniques of Spectral Music*, he describes anamorphosis as a musical formal approach:

> The idea behind this technique is to present a single object from different perspectives, which distort the object in various ways – sometimes even making it appear to be a different object altogether. In this way, one object can develop into a rich reservoir of musical and formal material that can sound very different in spite of its high degree of relatedness; creating very different and surprising effects without compromising the coherence of the musical material. (Fineberg 109)

Fineberg's comparison to anamorphosis is different from the one I make in chapter 2. The distorted projections occur across time rather than space. However, his description complements the formal approach I've taken in *Through a Window*. I composed multiple sections using the same underlying processes and distributed these sections throughout the composition. As the composition develops over time, processes return with new surface-level textures and characteristics reflecting the growth within the composition as a whole. These related sections sound like imperfect reflections of one another. The form therefore explores multiple realizations in time as well as space.

Figure 4-1 shows that the network sonification, recursion, swarms, and evolutionary processes are each presented as the dominant processes in three sections throughout the composition (outlined in bold). This mirrors the three movements, three sites, three algorithmic themes, and triple-interval cycles. The diagram also shows several connections across movements. The three movements are not independent from one another, but are intertwined, expressing different perspectives on similar concepts. The two Grit sections are an example of two different perspectives on a related process. The two sections begin with contrasting materials and progress towards differing musical targets but do so using the same algorithm. The three swarm sections create related textures but follow three contrasting pitch trajectories: Swarm I maintains a static center pitch, Swarm II descends, and Swarm III ascends. The two Ping sections both tend towards chromatic ascending intervals and present the same melodic material, but the textures draw inspiration from different ping sources (see section 2.5).

The recursive sections guided the formal trajectory of the composition. The three sections served as landmarks during the compositional process and I consciously arranged much of the surrounding material to develop to or from these sections. The recursive sections naturally exhibited a high degree of tension due to the continually expanding intervals from cycle to cycle. Recursion I creates a wave-like effect of crashing down and the receding. Recursion II creates an effect like being caught in a whirlpool. Recursion III resembles being carried down a river with a strong current. Recursion I and II both contain climactic moments, while Recursion III begins the motion towards the eventual climax of the composition in mm. 182-190.

The full recursive process from (1,1,1)-cycle to (e,e,e)-cycle is much too long to present in its entirety.[12] Instead, I selected short snapshots that I felt were representative of the total process. The excerpts I selected for Recursion I and Recursion III are bookends of the total

---

[12] Even with rotationally equivalent cycles filtered out, it takes 32 minutes to perform the cycle-of-cycles as 8th notes at 120 bpm.

interval expansion process; Recursion I begins with the (1,1,1)-cycle and progresses onwards to the (3,4,9)-cycle while Recursion III begins with the (8,8,9)-cycle and continues to the (e,e,e)-cycle which concludes the process.

Recursion I and III express opposing melodic contours. Since Recursion III uses very wide intervals, the octave-shifting boundary condition comes into play much more frequently and large ascending leaps tend to be mapped to smaller descending intervals. I was attracted to how the interval expansion process inverts itself when boundary limits are imposed: the opening ascending chromatic scale is mirrored by a concluding descending chromatic scale. Often in Recursion I, cycles may be perceived as descending gestures made up of ascending notes such as in Figure 4-8. In Recursion III, the opposite case may occur (Figure 4-9).

Figure 4-8 An interval cycle in Recursion I. The cycle is a descending series of gestures made up of ascending intervals.

Figure 4-9 An interval cycle in Recursion III (boxed). The cycle is an ascending series of gestures made up of descending intervals.

In Recursion II I used the "reflection" boundary condition to create a wave-like ascending and descending melodic contour, representing an intermediary step between Recursion I and III (Figure 4-10). Together, the three recursive sections represent the beginning, middle, and end of the total interval expansion process.

Figure 4-10 The reflection boundary condition creates an ascending and descending contour in Recursion II.

Sound-processing and soundfile playback help either differentiate or connect related sections. For example, the electronics differentiate the two Grit sections. In Grit I, pitch-shifting, delays, and spectral processing create a lush overlapping texture whereas in Grit II, short percussive soundfiles contribute to a playful dialogue within a sparse texture. In the case of Grit II, a pitch-detection module triggers a bank of samples that are transposed to match the pitches performed by a selected local performer. These samples are only heard in their respective local nodes so that the computer seems to mimic a single local performer. This mode of interaction in Grit II makes the electronics much more akin to that of an instrumental performer adding to an improvisatory dialogue, whereas in Grit I the electronics add resonance and function more like the sustain pedal on a piano.

In contrast to Grit, all three Swarm sections use the same sound-processing modules and this helps to create consistency among the sections. I use granulation and microtonal pitch-shifting modules to emphasize the sense of swarming and increase the overall intensity of the sections. The granulator records input from each performer into a circle buffer and uses a probability gate to limit the number of grains played by the granulator at any given time. The granulator also transposes each grain by a random microtonal amount, similar to the pitch-shifting modules.

Related processes often present similar approaches to layering and musical time. For example, most of the Evolutionary, Recursion, and Swarm sections use unmetered textures and rely on the stopwatches to help navigate time. These sections tend not to differentiate between foreground and background layers, but rather use the ensemble as a whole to create wave-like, cloud-like, or pointillistic textures. Two exceptions include the Recursion III section, which is metered and multi-layered and the Evolutionary Harmonies section, which includes foreground and background layers of soundfiles. In contrast, the Floating, Divergence, Ping, and Layers sections all present clear, metered materials with well-defined foreground and background elements.

When the more nebulous unmetered textures give way to metered materials, the effect is often like switching from an unfocused lens to a focused lens (or perhaps a nearly focused lens due to network latency). A clear example of this effect occurs at the transition between Swarms I and the culmination of Recursion II: Swarms I creates a highly complex cloud-like texture that suddenly changes into clear unison runs at m. 22. This is also similar to the culmination of the Floating section, which shifts from layered melodic lines in mm. 22-47 to culminate in a rhythmic unison melody from mm. 53-61. Another example occurs at the transition from the improvisatory textures in Grit II to the metered materials in Recursion III. In each of these cases, the shifts from unmetered materials to metered materials occur at prominent points in the formal structure creating stronger delineations between the sections.

4.3 Wave/form

When I generated the raw note data for Recursion I, I noticed a tendency for the material to cycle between layers of long overlapping gestures (Figure 4-11a), moments of long solo gestures (Figure 4-11b), and flurries of shorter fragments (Figure 4-11c). As a listener, I found cycling between these states to be hypnotic, like watching waves crashing onto a beach and then receding. Each crashing wave is different; some reach further in than others. The wave metaphor contributed to the composition's macrostructure.

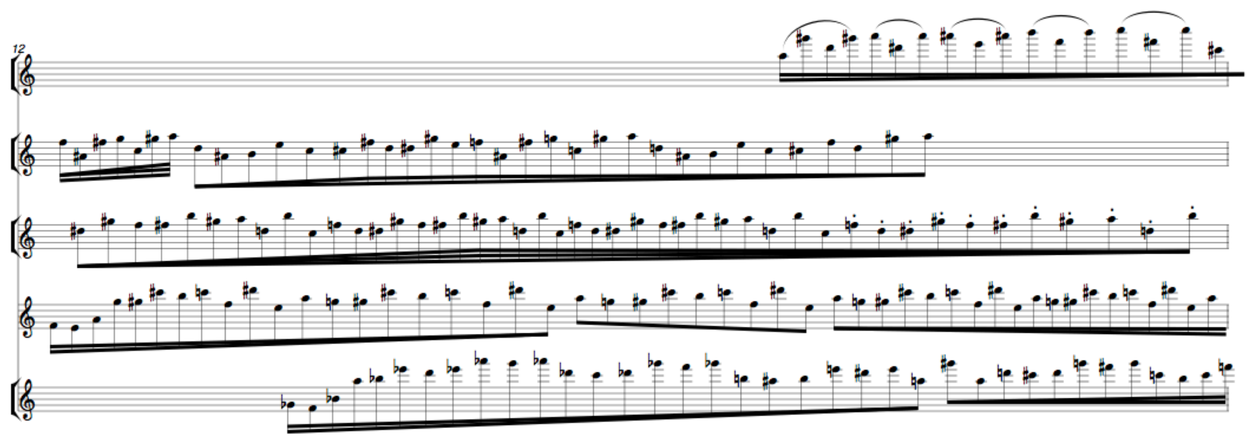

Figure 4-11a Overlapping interval series in the raw Recursion I note data. (Beaming added for clarity)

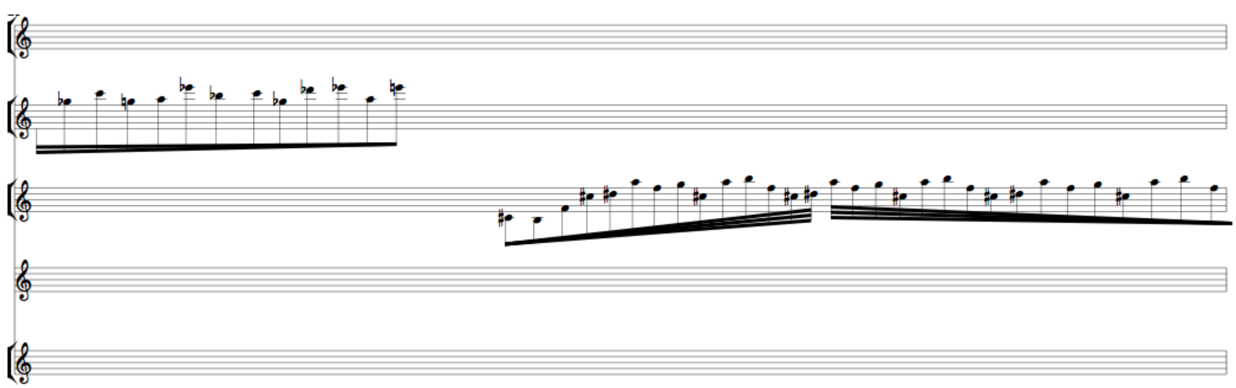

Figure 4-11b A solo interval cycle.

Figure 4-11c The boundary conditions prematurely cut off the cycle, creating a series of short fragments.

Many sections in movement 2, undergo a gradual build-up of energy followed by a sudden recession. For example, Figure 4-12 graphs the intended intensity of the Floating section. The section is divided into five subsections which alternate between high and low intensity. Each subsection grows in intensity before receding near the end. Figures 4-13 to 4-15 graph the attack density, sensory dissonance, and registral span over the course of the section. The graphs represent the average values for each measure in order to display general trends rather than moment-to-moment details. Figure 4-16 graphs the number of instruments performing foreground material over the course of the section. The attack density, sensory dissonance, and foreground instrument graphs outline a similar profile to my intended intensity. The graphs help to show how the surface-level texture contributes to the wave-like formal structure.

The wave-like shape allows me to easily embed catch-up points between sections (see section 2.1). Many sections end with a moment where the musicians can pause and regroup before a single musician begins the following section and resets the tempo. The five subsections in the Floating section demonstrate this pattern.

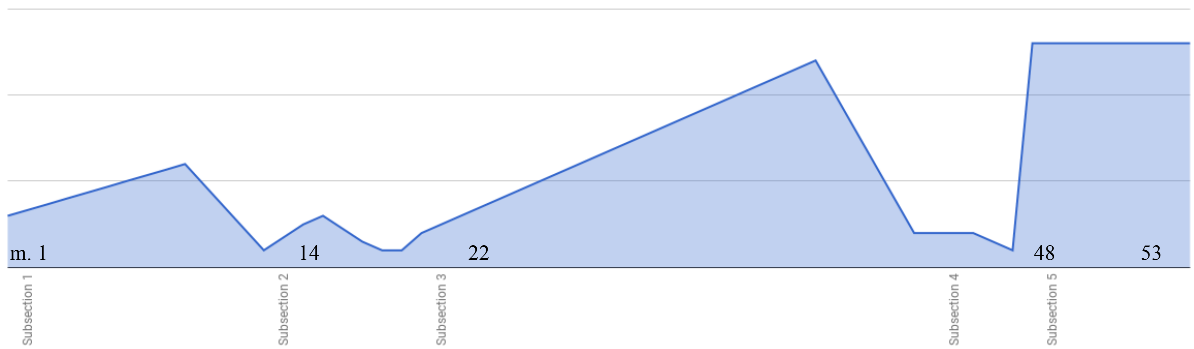

Figure 4-12 Floating (mm. 1-61), approximate intensity

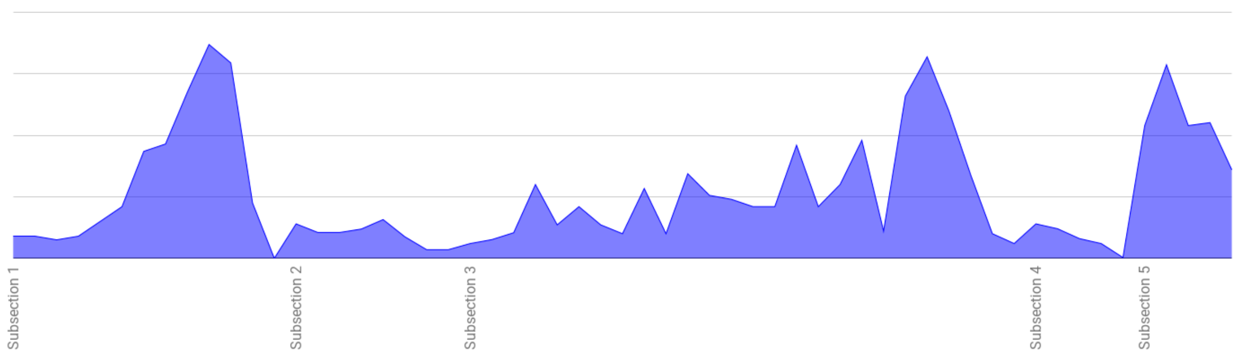

Figure 4-13 Floating (mm. 1-61), relative attack density

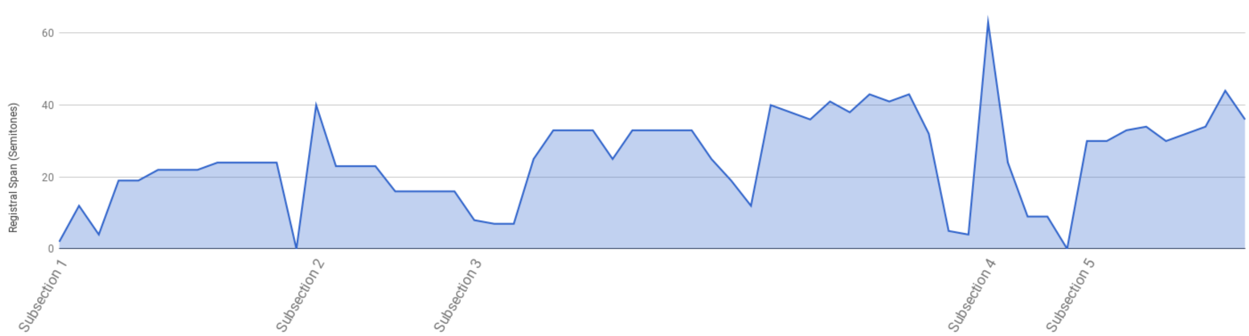

Figure 4-14 Floating (mm. 1-61), registral span (semitones)

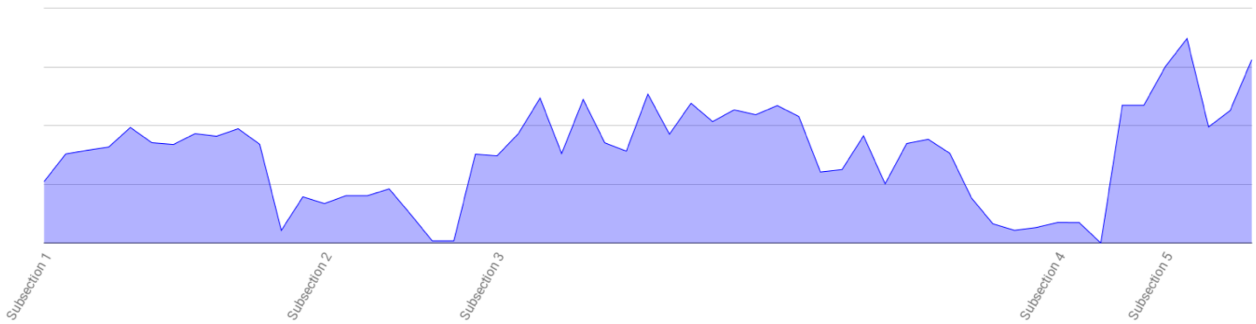

Figure 4-15 Floating (mm. 1-61), average sensory dissonance by measure.

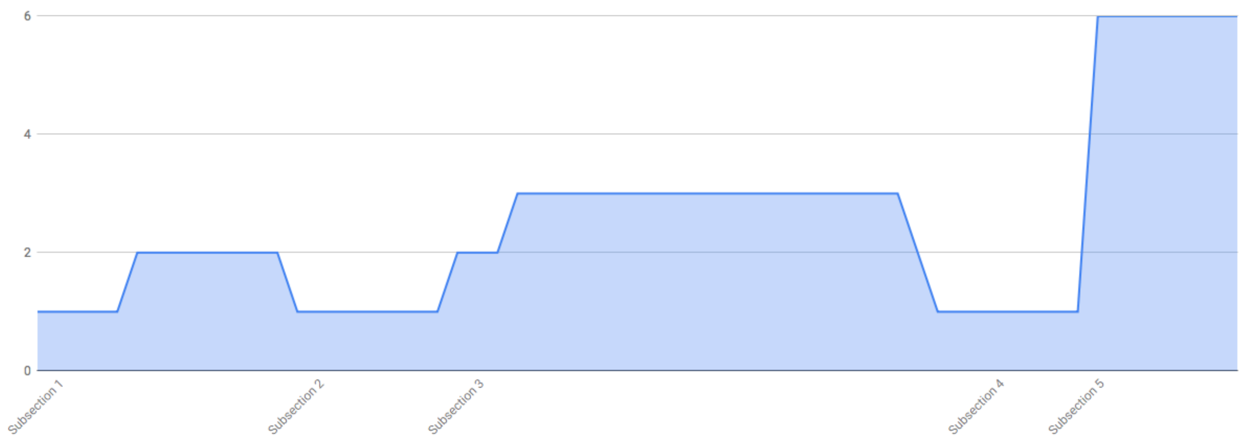

Figure 4-16 Floating (mm. 1-61), number of parts performing foreground melodic material.

The Floating section, as a whole, has a similar shape to its subsections; it builds up energy before quickly receding in intensity at the start of m. 62. The Floating section is a microcosm of the form of the movement itself: movement 2 alternates between sections of mounting and receding intensity representing waves crashing and then receding. The waves gradually increase in size until the climax at Recursion I. The wave pattern continues in movement 3 and the duration and intensity of the waves continue to increase. The final section of the work, from mm. 54-192, is a much longer wave building from mm. 54-181, before finally crashing down at mm. 182-190 and receding in mm. 191-192. The macroform is similar, but not identical, in shape to its component parts.

Figures 4-17 to 4-19 graph several musical parameters over the course of movements 2 and 3. I generated these graphs by exporting the NoteAbility Pro score files to the MusicXML format. I wrote a short script to parse the MusicXML files into lists of notes and measures. I used this note data to track different statistical characteristics. For each graph, I included average values for each measure as well a moving average of six measures to display general trends over time. Note that the attack density graph considers notes performed in unison between two instruments to count as two attacks even if they would be perceived as a single attack in performance. Also note that the sensory dissonance calculations respond to the attack density. When the attack density is high, the calculation includes more notes and is more likely to receive dissonant tones. The sensory dissonance and attack density charts show peaks in roughly equivalent positions to the 'intensity' graph in Figure 4-2. The contour graph measures the proportion of ascending vs. descending intervals for each measure. Positive values represent a tendency towards ascending intervals and negative values represent a tendency towards descending intervals. The contour graphs show a gradual meandering shift from ascending to descending intervals over the length of the two movements.

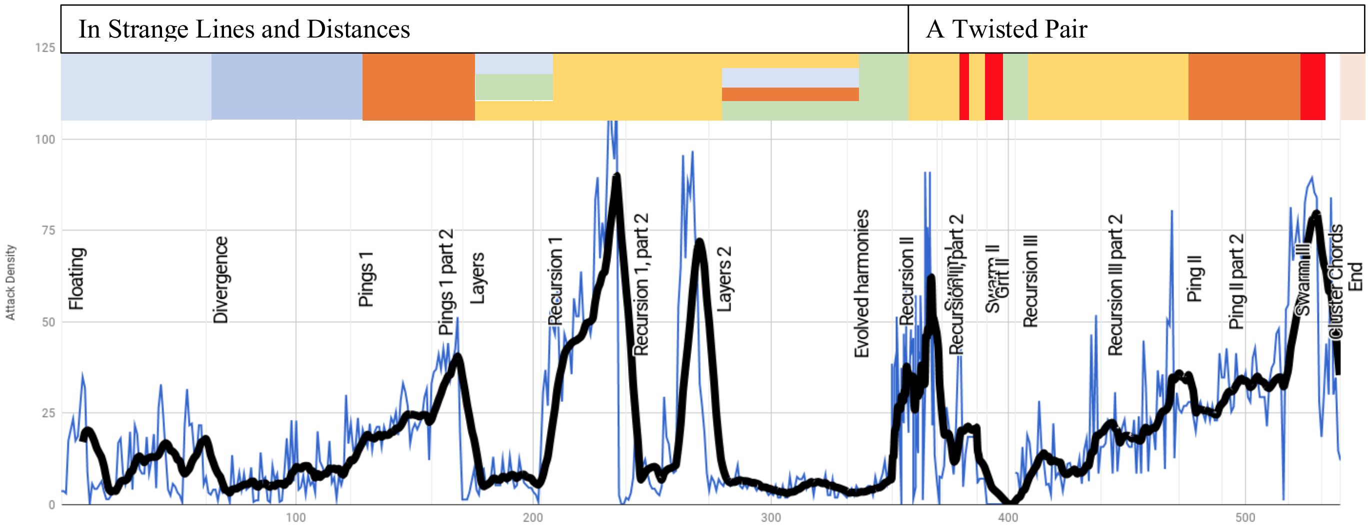

Figure 4-17 Movements 2-3, average attack density.

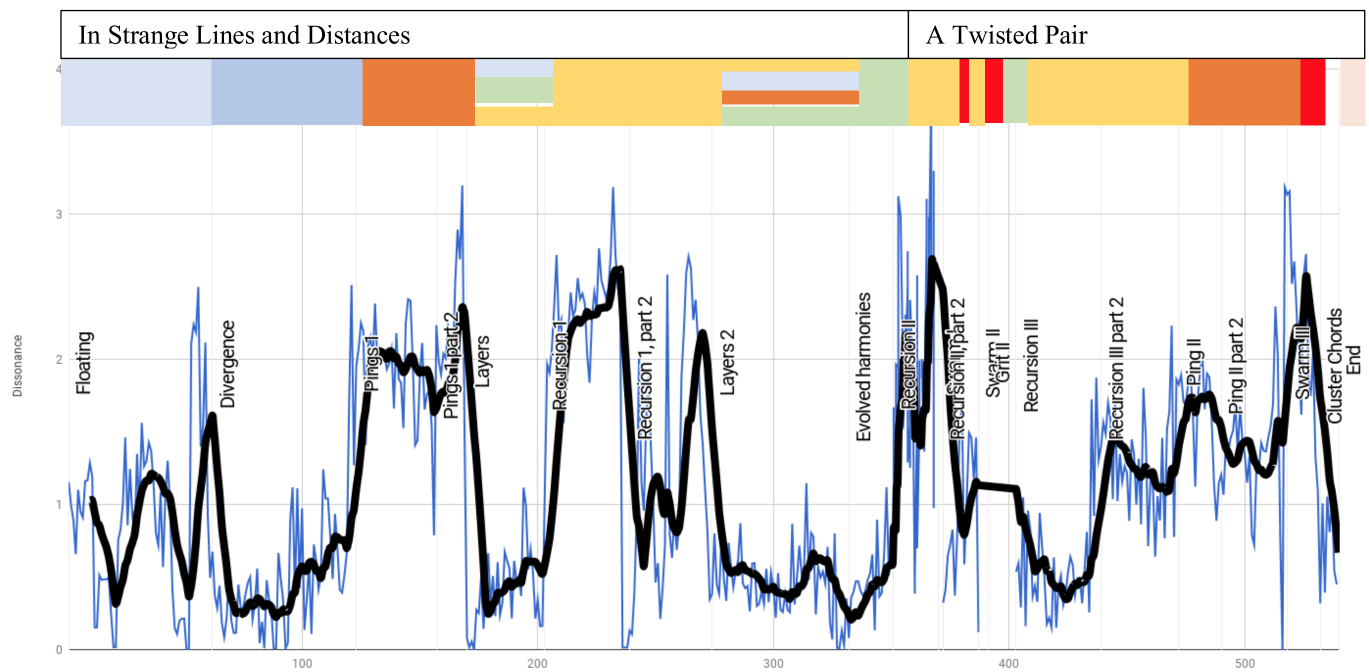

Figure 4-18 Movements 2-3, average sensory dissonance.

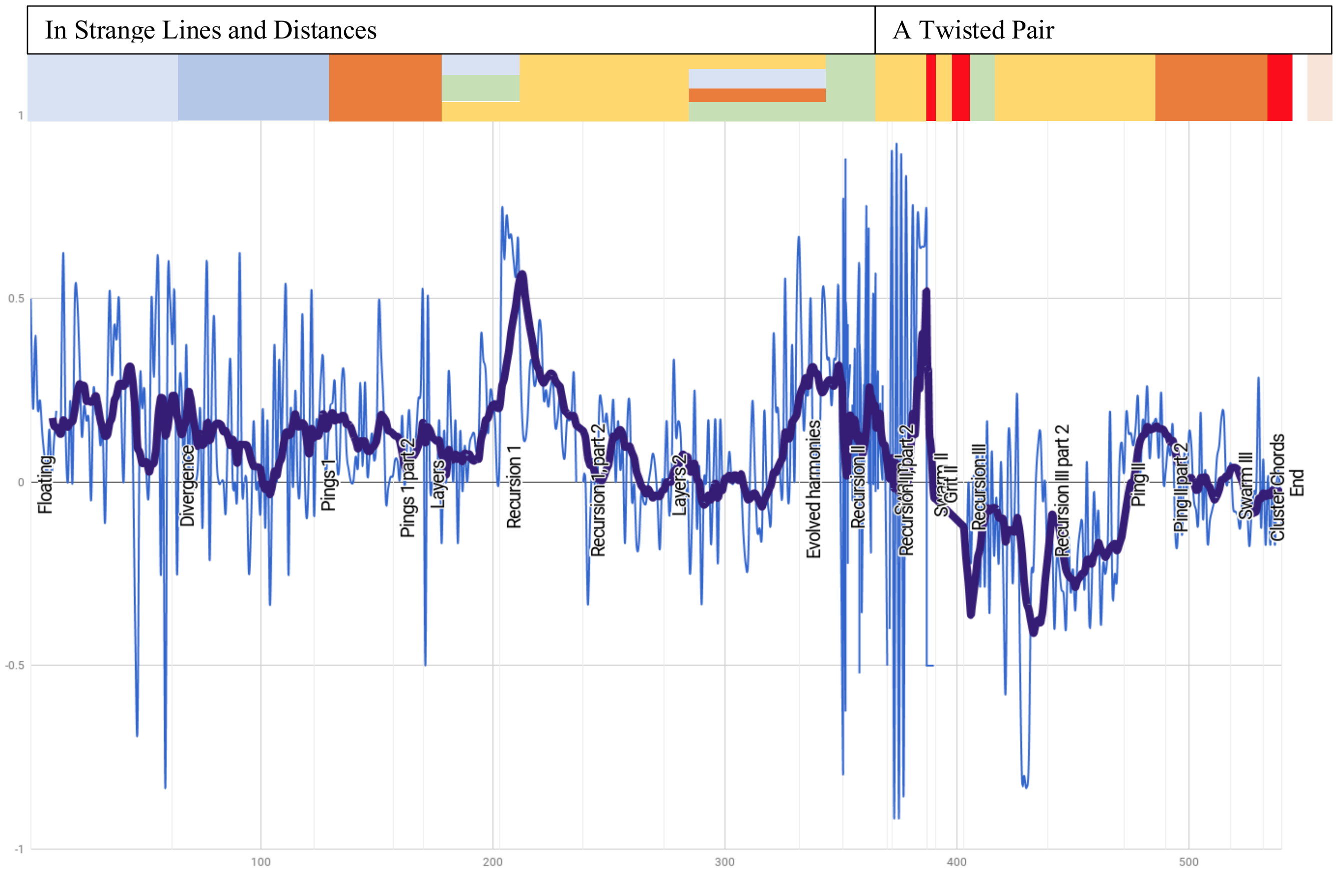

Figure 4-19 Movements 2-3, average contour direction. Positive values denote primarily ascending contours and negative values denote primarily descending contours.

By combining a gradual increase in intensity with a wave-like alternation of high and low intensity, I create a spiral-like macroform (Figures 4-19 and 4-20). The spiral can be viewed as an archetypal shape for the ebb and flow of intensity throughout the composition. However, although Figure 4-20 shows some similarities between the spiral shape and the macroform, the macroform is much less regular. This is due to the bottom-up aspects of composition. The materials and processes themselves help dictate their own proportions and intensities. The bottom-up and top-down approaches complement one another to create a more nuanced and detailed shape. The macroform emerged as a combination of both approaches.

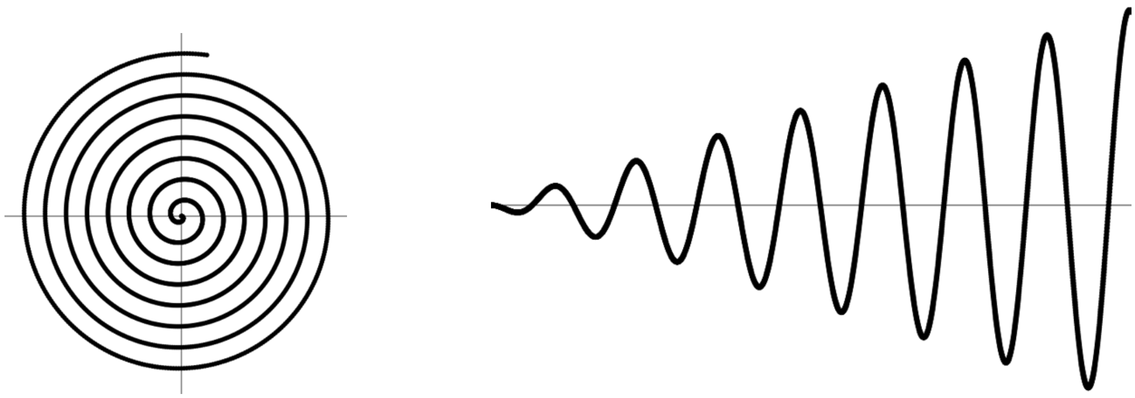

Figure 4-20 Spiral graphed over time is equivalent to a sine wave with an increasing amplitude.

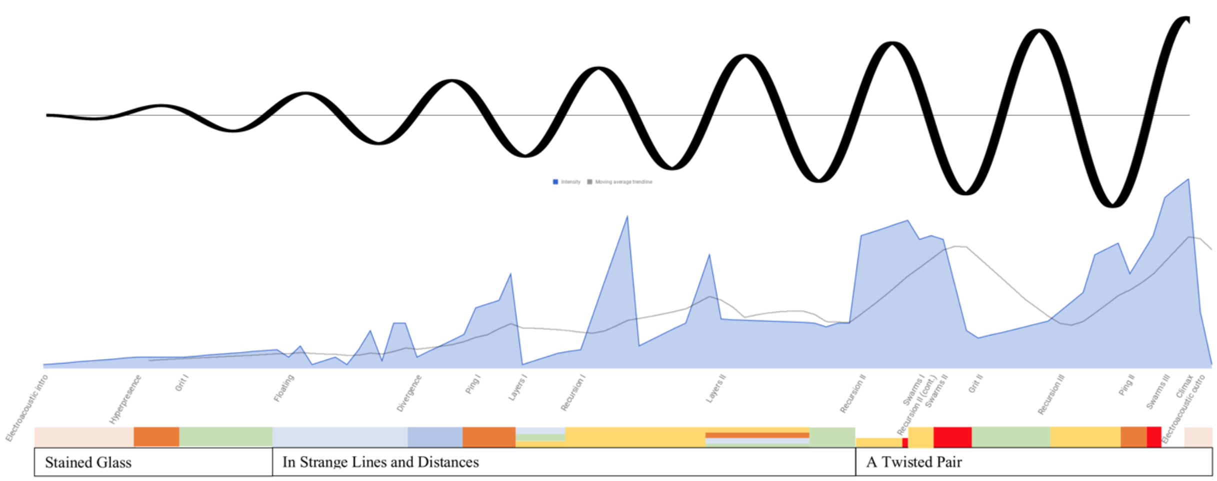

Figure 4-21 A spiral wave compared to the intended intensity graph.

Conclusion

Composing for a Distributed Ensemble

The networked performance environment strongly informed my compositional approach. I found that sound processing is a natural addition to a networked performance. Since the acoustic instruments are already mediated through the sound systems, it is a simple matter to add delays, feedback, reverberation, pitch-shifting, and other digital audio effects at each location. The audio effects do not need to be presented in parallel with the acoustic sounds and can therefore seem to extend beyond their typical instrumental ranges at remote nodes without the acoustic component of the performance breaking the illusion.

Sample banks and digital controllers are both very well-suited to NMP. Soundfiles can be stored on each computer individually and cued from any location by sending very small pieces of data over the network. By including multiple banks of soundfiles and mapping the cues to different banks at each location, I can create textures that are similar but also unique to each location.

I find that NMP creates a unique acoustic environment. I enjoy listening to performances mediated through sound systems and made to sound larger-than-life, mysterious, or otherwise coloured by the technology. In *Through a Window*, I used close microphone placements to create a sense of closeness between the remote sites. Although this element is not often actively manipulated and does not necessarily inform the macroform of the composition, it shapes the sound of the work itself and is thus an important aspect in the sonic identity of the composition. Even though the work's variable instrumentation means that timbres will be different from performance to performance, the mediated environment enables the computer musicians to balance the sound and creates an effect that will sound close and intimate.

The network is a challenging environment to compose for largely due to the latency in transmitting audio between sites. In response, I frequently composed fluid textures and used proportional notation to help accommodate the latency and help the musicians coordinate in an unfamiliar performance setting. The stopwatches, anchor lines, cues, and catch-up points were essential to help the musicians perform together within this setting. The work demands a very high degree of technical proficiency from the musicians and I was very pleased to hear the musicians rise to the challenge during the première.

Regarding networked performance practice, the most persistent challenge was communicating clearly and efficiently between the locations in rehearsal. It was often unclear who was speaking and whether everyone was being heard clearly in the remote nodes. It was also easy to forget to speak directly into the microphones. Often, performers needed to repeat questions or directions several times and this slowed down the rehearsal. Ideally, a separate room microphone should be used at each site to allow everyone to speak freely in rehearsal. This microphone must be muted and unmuted diligently during rehearsal in order to maintain the correct sound quality and to minimize the potential for feedback while the musicians are playing. A video chat connecting the sites is also very helpful to show who is speaking in the remote nodes and to help communicate body language as well as voice.

Musically, the most challenging sections of the composition to perform over the network were the sections that required tight synchronization and also contained short melodic lines that bounced between the instruments such as in movement 2 from mm. 123-156 and movement 3 from mm. 54-85. The latency and lack of visual communication created a significant challenge for the ensemble during these sections. The anchor line technique helps to mitigate this challenge but the section still demands a very high degree of precision from the ensemble.

As a computer musician, I also found it challenging to ensure that the ensemble could reliably hear the anchor lines during these sections. I was located in the middle of the hall rather than onstage and perceived the balance differently than the ensemble. The anchor lines would sometimes be audible from my position in the center of the hall, but difficult to perceive from the stage. A simple solution is to place anchor lines and cues in the front monitors where they sound louder to the ensemble than to myself.

I was also surprised in rehearsal by how dramatically spatialization influenced the tempo. The ensemble had a pronounced tendency to slow down when I placed the anchor lines in the rear speakers. As I moved the anchor lines nearer to the stage, the performance became more manageable because there was less acoustic delay between the remote anchor lines and the local musicians.

Moving forward, I am excited by the possibilities for community engagement that the network offers. NMP ensembles and communities such as the Global Loop Orchestra, NowNet Arts, and the Ethernet Orchestra have emerged in recent years and I hope that the NMP community continues to grow (Fields 2017; “NowNet Arts” 2017; Mills 2010). I plan to continue composing NMP works for acoustic and electronic instruments and to continue developing tools and documentation to facilitate networked performances.

In particular, I first plan to pursue a series of networked poetry readings. In *Through a Window,* I found that the intimacy created through close microphone placement was particularly effective when it involved human voice. A poetry reading involving multiple distributed speakers and performed in multi-channel environments will heighten the sense of intimacy that I explored in *Through a Window*. Next, I plan to explore the potential of the network as a tool for online database access. This may involve compositions that access online databases of crowdsupplied recordings or that facilitate audience participation through social media. Thirdly, I hope to explore the possibilities of the network to facilitate interactions between participants on private home networks through a shared virtual environment. These environments may be simple game-like compositions, perhaps using shared step-sequencer-like interfaces or perhaps involving more open-ended virtual spaces in which each user creates sound by using an avatar to interact with a three-dimensional environment, similar to Robert Hamilton's *Echo::Canyon* (2013). Finally, I plan to compose open-form works where each musician is challenged to listen and respond to audio cues from their remote collaborators based on a set of rules or musical fragments. The trajectory of the compositions will change based on the choices made by the musicians throughout the performance, emphasizing the themes of collaboration and communication.

Composing algorithms

Given the systematic nature of many elements within this document, I would like to emphasize the role of creativity and musicality in shaping my approach to designing the algorithms and processes in *Through a Window*. While describing his GENDY software, Iannis Xenakis argues that compositional choice is a crucial factor in designing musical systems:

> I am always trying to develop a program that can create the continuity of an entire piece. This is a struggle, because there are always parts that you prefer over others. So you have to change them, to stop the process, start some other one, and then put these two different ones together. This can be taken very far. As I move toward multiple voices, the problem becomes even more complicated. I have introduced probabilistic controls in GENDY, within certain limits. When I say limits, I do not mean just any kind – these limits must yield interesting results to the ear. Who says it is interesting? I do. This is where compositional choice enters in. (Robindoré 13-14)

This has been my experience in *Through a Window*. When investigating different algorithmic concepts, I created aural impressions of how I expected the concepts would sound when translated to music. As I implemented the algorithms or adapted them to music, the generated material often fell short of my expectations and I would revise and rework the design in order to better capture the texture or quality that I hoped to convey. This provided valuable insight into my own compositional process. When the algorithms fell short, it revealed elements that I had taken for granted, overlooked, or misunderstood either about the function of the algorithm or its mapping to music.

I didn't attempt to create a generalized system for composing music or model a specific style of music but instead attempted to translate conceptually beautiful ideas into the domain of music. In this composition, the algorithms are not a tool for solving a particular problem but are themselves musical.

I have also not attempted to generate an entire composition from a single algorithm as suggested by Xenakis. Instead, I combine and juxtapose many different algorithmic processes, algorithmically-inspired processes, and non-algorithmic materials in a way that I hope creates a dynamic and engaging composition. The different algorithms have different functions within the piece and I utilize them at opportune moments to create the spiral-like flux from low to high intensity and back.

Throughout the work, algorithmic and non-algorithmic approaches overlap and intertwine, helping to form the musical language of the composition. I believe I have integrated evolution, recursion, and swarming into my compositional toolbox and I intend to continue to develop, experiment, and refine these approaches in future work.

Bibliography

Agostini, Andrea, and Daniele Ghisi. "A Max Library for Musical Notation and Computer-Aided Composition." *Computer Music Journal* 39, no. 2 (2015): 11-27.

Ascott, Roy. "Is There Love in the Telematic Embrace?" *Art Journal* 49, no. 3 (1990): 241-247.

Bacon, Francis. *The New Atlantis*. London: Printed for John Crooke at the Signe of the Ship in St. Pauls Church-yard, 1660. Accessed July 2, 2018. Early English Books Online.

Barbosa, Alvaro. "Displaced Soundscapes: A Survey of Network Systems for Music and Sonic Art Creation." *Leonardo Music Journal* 13 (2003): 53-59.

Beyst, Stefan. "Luigi Nono's Quando Stanno Morendo. Diario Polacco 2. Cries, Whispers and Celestial Voices." Stefan Beyst. 2004. Accessed July 1, 2018. http://d-sites.net/english/nonopolacco.html.

Bhumber, Kiran. "Raula." *Kiran Bhumber*. 2017. Accessed July 1, 2018. http://www.kiranbhumber.com/raula.

Biles, John. "GenJam in Perspective: A Tentative Taxonomy for GA Music and Art Systems." *Leonardo 36* no. 1 (2003): 43-45.

———. "Evolutionary Computation for Musical Tasks." In *Evolutionary Computer Music*, edited by Eduardo Miranda and John Biles, 28-51. Berlin: Springer, 2007.

———. "Improvising with Genetic Algorithms: GenJam." In *Evolutionary Computer Music*, edited by Eduardo Miranda and John Biles, 137-169. Berlin: Springer, 2007.

Bischoff, John, Rich Gold, and Jim Horton. "Music for an Interactive Network of Microcomputers." *Computer Music Journal* 2, no. 3 (1978): 24-29.

Blackwell, Tim, and Michael Young. "Swarm Granulator." In *Proceedings of the European Workshop on Evolutionary Music and Art*, Coimbra, Portugal. 2004.

Blackwell, Tim. "Swarming and Music." In *Evolutionary Computer Music*, edited by Eduardo Miranda and John Biles, 194-217. Berlin: Springer, 2007.

Böhm, Volker. *Vb.stretch~*. 2013. Accessed August 11, 2018. http://vboehm.net/downloads

Boulez, Pierre. *Répons*. London: Universal Edition, 1985.

———. *Dialogue de L'Ombre Double*. Wein: Universal, 1984.

Braasch, Jonas. "Virtual Cistern Concert." Presented at *NowNet Arts Conference*, Stony Brook University and Rensselaer Polytechnic Institute, Stony Brook and Troy, New York. 2018.

Cáceres, Juan-Pablo, and Chris Chafe. "JackTrip: Under the Hood of an Engine for Network Audio." *Journal of New Music Research* 39, no. 3 (2010): 183–187.

Cáceres, Juan-Pablo, and Alain Renaud. "Playing the Network : The Use of Time Delays as Musical Devices." In *Proceedings of the* 2008 *International Computer Music Conference*, 24-29, 2008.

Cáceres, Juan-Pablo, Alain Renaud, and Justin Yang. *Net:Disturbances*. 2008. Accessed July 10, 2018. https://ccrma.stanford.edu/~jcaceres/music/.

"CANARIE Network." *CANARIE*. Accessed June 20, 2018. https://www.canarie.ca/network/

Carnegie Hall. "Kaija Saariaho on Composing Using a Computer Today." *Youtube.* 2012. Accessed July 03, 2018. https://www.youtube.com/watch?v=pP6BYRjHxgw.

Carôt, Alexander, and Christian Werner. "Network Music Performance – Problems, Approaches and Perspectives." *Music in the Global Village Conference*, Hungary, Budapest. 2007.

Cayko, Ethan. *Rhythmic Topologies and the Manifold Nature of Network Music Performance*. Master's Thesis, University of Calgary, 2016.

Chafe, Chris. "Distributed Internet Reverberation for Audio Collaboration" In *Proceedings of the AES 24th International Conference*, Banff, 2003.

———. "Living with Net Lag." In *Proceedings of the AES 43rd International Conference*, Pohang, 2011.

———. "Ping." *Chris Chafe*. 2001. Accessed May 22, 2018. http://chrischafe.net/portfolio/ping/

———. "Tapping into the Internet as an Acoustical/Musical Medium," *Contemporary Music Review* 28, no. 4 (2009): 413-420.

———. "Tomato Quintet." *Chris Chafe*. 2007. Accessed May 17, 2018. http://chrischafe.net/tomato-quintet/

Chafe, Chris, Juan-Pablo Caceres, and Michael Gurevich. "Effect of temporal separation on synchronization in rhythmic performance." *Perception* 39, no. 7 (2010): 982-992.

Clarke, Michael, and Deborah Clarke. "Anamorphosis." *The Concise Oxford Dictionary of Art Terms.* 2nd ed. Oxford University Press, 2010. Accessed May 19, 2018.

Cockos Incorporated. "NINJAM," Accessed May 2016. http://www.cockos.com/ninjam/

Cycling '74. "Max." Cycling '74, Version 7.3.5 (2018). Accessed July 12, 2018. https://cycling74.com/.

Dawkins, Richard. *The Blind Watchmaker*. London: Penguin, 1986.

Deal, Scott. *Goldstream Variations*. 2012. Accessed May 22, 2018. http://tavellab.net/product/goldstream-variations/

Degazio, Bruno. "A Particle System for Musical Composition." *EContact!*, 2013. Accessed February 2, 2016. http://www.econtact.ca/

———. *On Growth and Form*. Canadian Music Centre, 1988.

———. "On Growth and Form." In *Proceedings of the 1992 International Computer Music Conference* compact disc, San Jose, California. ICMA. 1992.

———. “New Software Composition Tools.” In *Proceedings of the Fourth Biennial Arts and Technology Symposium*, New London, Connecticut, 1993.

———. “The Evolution of Musical Organisms.” *Leonardo Music Journal*, Vol 7. (1997): 27-33.

Du Moncel, Theodore. *The Telephone, the Microphone, and the Phonograph*. New York: Harper, 1879.

Emmerson, Simon. "‘Losing Touch?’: The Human Performer and Electronics." In *Music, Electronic Media and Culture*, 194-216. Burlington: Ashgate, 2000.

Ernst, David. *The Evolution of Electronic Music*. New York: Schirmer, 1977.

Ferguson, Sean. *Concerto for Piano and Orchestra*. PhD diss., McGill University, 2000.

Fields, Kenneth. "GLOrk: Global Loop Orchestra." *Facebook*. September 2017. Accessed July 9, 2018. https://www.facebook.com/globalloop/.

———. "Syneme: Live." *Organised Sound* 17, no. 01 (2012): 86-95.

Fineberg, Joshua. “Guide to the Basic Concepts and Techniques of Spectral Music.” *Contemporary Music Review* 19, no. 2 (2000): 81-113.

Föllmer, Golo. "Electronic, Aesthetic and Social Factors in Net Music." *Organised Sound* 10, no. 3 (2005): 185-192.

Freeman, Jason, Kristjan Varnik, C. Ramakrishnan, Max Neuhaus, Phil Burk, and David Birchfield. “Auracle: A Voice-controlled, Networked Sound Instrument.” *Organised Sound* 10, no. 3 (2005): 221-31.

"Global Partners." *CANARIE*. Accessed June 20, 2018. https://www.canarie.ca/partners/global/

Goodall, Jane. *Through a Window: My Thirty Years with the Chimpanzees of Gombe.* Boston: Houghton Mifflin Company, 1990.

Gresham-Lancaster, Scot. “The Aesthetics and History of the Hub: The Effects of Changing

Technology on Network Computer Music." *Leonardo Music Journal* 8, no. 1 (1998): 39–44.

Hajdu, Georg, and Nick Didkovsky. "On the Evolution of Music Notation in Network Music Environments." *Contemporary Music Review* 28, no. 4 (2009): 395-407.

Hamel Keith. *Full Circle for Trombone and Interactive Electronics*. 2014. Accessed July 2, 2018. http://debussy.music.ubc.ca/Keith/scores.html.

———. "NoteAbility Pro." Opus 1 Music, Version 2.653 (2018). Accessed July 12, 2018. http://debussy.music.ubc.ca/NoteAbility

Hamel, Keith, and Bob Pritchard. *UBC Max/MSP/Jitter Toolbox*. Version 0.99. 2016. Accessed August 11, 2018. http://www.opusonemusic.net/muset/toolbox.html.

Hamilton, Rob, and Chris Platz. *Echo::Canyon*. 2013. Accessed July 10, 2018. https://ccrma.stanford.edu/~rob/portfolio/echocanyon.

Harvey, Jonathan. *Bhakti for 15 Instrumentalists and Quadraphonic Tape (1982)*. London: Faber, 1989.

———. *Speakings*. London: Faber Music, 2008.

Hexler. "TouchOSC." Hexler.net, Version 1.9.10 (2018). Accessed July 12, 2018. https://hexler.net/software/touchosc

Holbrook, Geoffrey. *Optimizing Future Perfect: A Model for Composition with Genetic Algorithms*. PhD diss., Columbia University, 2015.

Hope, Cat, and Lindsey Vickery. "The Decibel Scoreplayer – a Digital Tool for Reading Graphic Notation." *TENOR 2015: International Conference on Technologies for Music Notation and Representation*, 59-69, Paris. 2015.

Hornby, Gregory, Jason Lohn, and Derek Linden. "Computer-Automated Evolution of an X-

Band Antenna for NASAs Space Technology 5 Mission." *Evolutionary Computation* 19, no. 1 (2011): 1-23.

Horner, Andrew. “Envelope matching with genetic algorithms.” *Journal of New Music Research*, no. 24 (1995): 318–341.

Horner, Andrew, James Beauchamp, and Lippold Haken. "Machine Tongues XVI: Genetic Algorithms and Their Application to FM Matching Synthesis." *Computer Music Journal* 17, no. 4 (1993): 17-29.

Hugill, Andrew. “The Origins of Electronic Music.” Chapter. In *The Cambridge Companion to Electronic Music*, edited by Nick Collins and Julio d'Escrivan, 2nd ed., 7–24. Cambridge Companions to Music. Cambridge: Cambridge University Press, 2017.

"Israel Starling Murmuration Captured on Video." BBC News. February 02, 2015. Accessed June 23, 2018. https://www.bbc.com/news/av/science-environment-31105433/israel-starling-murmuration-captured-on-video

Jacob, Bruce. “Composition with genetic algorithms.” In *Proceedings of the International Computer Music Conference*, 452–455, Banff. 1995.

Jacob, Christian, and Gerald Hushlak. "Design Through Evolutionary Swarms." *Swarm Art*. Accessed May 25, 2018. http://swarmart.com/SwarmArtNew/SwarmArt.com/Home of SwarmArt.html.

Kane, Brian. “Aesthetic Problems of Net Music.” In *Proceedings of Spark 2007*. University of Minnesota, 2007.

Kim-Boyle, David. “Network Musics - Play, Engagement and the Democratization of Performance.” *Contemporary Music Review* 28, no. 4 (2009): 363-375.

———. "Sound Spatialization with Particle Systems." In *Proceedings of the 8th International Conference on Digital Audio Effects*, 65-68, Madrid. 2005.

———. "Spectral and Granular Spatialization with Boids." *In Proceedings of the 2006 International Computer Music Conference*, 139-42, New Orleans. 2006.

Landy, Leigh. *Understanding the Art of Sound Organization*. Cambridge, Mass.: MIT Press, 2007.

Lerdahl, Fred, and Ray Jackendoff. *A Generative Theory of Tonal Music*. Cambridge: MIT Press, 1996.

Levy, Benjamin. *Bc.yinstats*. WoMax. 2009. Accessed August 11, 2018. http://recherche.ircam.fr/equipes/repmus/WoMax.

Lewis, George E. "Too Many Notes: Computers, Complexity and Culture in Voyager." *Leonardo Music Journal* 10 (2000): 33-39.

Lutoslawski, Witold. *Chain 1.* London: Chester Music, 1983.

———. *String Quartet (1964): Study Score*. London: Chester, 1970.

Magnus, Cristyn. *Evolutionary Sound: A Non-Symbolic Approach to Creating Sonic Art With Genetic Algorithms*. PhD diss., University of California, San Diego, 2010.

Malloch, Joseph, and Marcelo M. Wanderley. "The T-Stick: From Musical Interface to Musical Instrument." In *Proceedings of the International Conference on New Interfaces for Musical Expression*, 66-69, New York. 2007.

Manning, Peter. *Electronic and Computer Music*. 4th ed. Oxford: Oxford University Press, 2013.

Manoury, Philippe. *Jupiter.* Paris: Editions Durand, 1992.

McConville, Brendan. "A 'Simple Composition' of Charles Wuorinen: Isomorphism, Self-Similarity, and Nesting in Cello Variations." *Theory and Practice* 36, no. 1 (2011): 151-

177.

Mills, Roger. “About.” *Ethernet Orchestra*. September 26, 2010. Accessed July 9, 2018. http://ethernetorchestra.netpraxis.net/info/?page_id=2.

Miranda, Eduardo, and John Biles. *Evolutionary Computer Music*. London: Springer, 2007.

Morris, Robert. “New Directions in the Theory and Analysis of Musical Contour.” *Music Theory Spectrum* 15, no. 2 (1993): 205-28.

"Network." *Cybera*. 2016. Accessed June 20, 2018. https://www.cybera.ca/network/network-index/

Neuhaus, Max, Jason Freeman, Kristjan Varnik, C. Ramakrishnan, Phil Burk, and David Birchfield. *Auracle*. 2005. Accessed October 19, 2014. http://auracle.org

Nobles, Jordan. "Open Score Collection." *Jordan Nobles - Composer*. 2017. Accessed May 22, 2018. https://jordannobles.com/music/open/

Noll, Udo. "Radio Aporee ::: Projects." *Radio Aporee*. 2006. Accessed May 22, 2018. https://aporee.org/aporee.html

Nono, Luigi. *Quando Stanno Morendo: Diario Polacco N. 2.* Milano: Ricordi, *1982*.

“NowNet Arts.” *NowNet Arts.* 2017. Accessed July 9, 2018. https://nownetarts.org/about.

Oda, Reid. *Tools and Techniques for Rhythmic Synchronization in Networked Musical Performance*. PhD diss., Princeton University, 2017.

Ogborn, David. 2012. “Composing for a Networked, Pulse-Based, Laptop Orchestra.” *Organised Sound* 17, no. 1: 56–61.

Oliveros, Pauline. “Networked Music: Low and High Tech.” *Contemporary Music Review* 28 no. 4 (2009): 433-435.

Orosz, István. "Art with a Double Meaning." In *Proceedings of Bridges Conference* *on Art and Mathematics*, 239-246, Hungary, Pécs. 2010.

Parabola Editors. "Through a Window, by Jane Goodall." *Parabola*. 2015. Accessed July 1, 2018. https://parabola.org/2015/12/29/through-a-window-by-jane-goodall/.

Perkis, Tim. "Wax Lips." Hub. 1991. Accessed August 08, 2018. http://www.o-art.org/history/08_Computer/Hub/hubspecs.html.

Pressnitzer, Daniel, and Stephen Mcadams. "Acoustics, Psychoacoustics and Spectral Music." *Contemporary Music Review* 19, no. 2 (2000): 33-59.

Puckette, Miller. "Not Being There." *Contemporary Music Review* 28 no. 4 (2009): 409–412.

Rebelo, Pedro. "Dramaturgy in the Network." Contemporary Music Review 28, no. 4 (2009): 387–393.

———. *Netrooms: The Long Feedback.* 2008. Accessed July 10, 2018. https://netrooms.wordpress.com/.

Reich, Steve. *Piano Phase: For Two Pianos or Two Marimbas (1967)*. Wien: Universal Ed., 2010.

Reis, Gustavo Miguel Jorge dos, and Francisco Fernandez de Vega. "Electronic Synthesis Using Genetic Algorithms for Automatic Music Transcription." 1959-1966. In *Proceedings of the Genetic and Evolutionary Computation Conference*, London. 2007.

"Rettangoli Gialli Concentrici Senza Angoli Al Suolo by Felice Varini." *My Favorite Arts*. Accessed May 22, 2018. https://theartstack.com/artist/felice-varini/rettangoli-gialli-concen

Reynolds, Craig. "Flocks, Herds and Schools: A Distributed Behavioral Model." *Computer Graphics* 21, no. 4 (1987): 25-34.

Roads, Curtis. *Microsound*. Cambridge, Mass.: The MIT Press, 2001.

———. *The Computer Music Tutorial*. Cambridge, Massachusetts: The MIT Press, 1996.

Robindoré, Brigitte, and Iannis Xenakis. "Eskhate Ereuna: Extending the Limits of Musical Thought: Comments on and by Iannis Xenakis." *Computer Music Journal* 20, no. 4 (1996): 11-16.

Rose, Francois. "Introduction to the Pitch Organization of French Spectral Music." *Perspectives of New Music* 34, no. 2 (2000): 6-39.

Rottondi, Cristina, Chris Chafe, Claudio Allocchio, and Augusto Sarti. "An Overview on Networked Music Performance Technologies." *IEEE Access* 4 (2016): 8823-843.

Saariaho, Kaija. *Lonh.* London: Chester Music, 1996.

———. "Timbre and Harmony: Interpolations of Timbral Structures." *Contemporary Music Review* 2, no. 1 (1987): 93-133.

Schacher, Jan, Daniel Bisig, and Martin Neukom. "Composing with Swarm Algorithms – Creating Interactive Audio-Visual Pieces Using Flocking Behavior." In *Proceedings of the 2007 International Computer Music Conference,* 100-107, Copenhagen, 2007.

Schafer, Raymond Murray. *Music for Wilderness Lake: for 12 Trombones.* Ontario: Arcana Editions, 1981.

Schnell, Norbert. *Yin~*. Ircam IMTR. Accessed August 11, 2018. http://imtr.ircam.fr/imtr/Max/MSP_externals

Settel, Zack. *Punjar for Soprano Sax and Electronics*. 1997. Accessed July 2, 2018. http://sheefa.net/zack/portfolio/Punjar/Punjar_score.pdf.

Simoni, Mary. "The Analysis of Electronic Music." In *The Cambridge Companion to Electronic Music*, by Nicolas Collins and Escrivan Rincón Julio D, 274-91. 2nd ed. Cambridge: Cambridge University Press, 2017.

Sims, Karl. "Evolved Virtual Creatures." Karl Sims Home Page. Accessed May 17, 2018. http://www.karlsims.com/evolved-virtual-creatures.html

Sivuoja-Gunaratnam, Anne. "Miniatures and Tensions: Phenomenological Reverberations in and Around Kaija Saariaho's Lichtbogen." *Intersections* 25, no. 1-2 (2005): 44-66.

Smith, Benjamin. *Telematic Composition*. PhD diss., University of Illinois at Urbana-Champaign. 2011.

Stewart, D. Andrew. *Catching Air and the Superman*. PhD diss., McGill University, 2010.

Stockhausen, Karlheinz. *Klavierstucke XI.* Vienna: Universal Edition, 1956.

Syneme. "Artsmesh." Artsmesh Network Music. Version 2.3, (2017). Accessed July 12, 2018. https://www.artsmesh.com

Tanaka, Atau, and Petra Gemeinboeck. *Net_Dérive*. Atau Tanaka. 2006. Accessed October 19, 2014. http://www.ataut.net/site/Net-Derive

Tanenbaum, Andrew, and David Wetherall. *Computer Networks*. Harlow: Pearsonn, 2014.

Truax, Barry. *Acoustic Communication*. 2nd ed. Westport, Conn.: Ablex, 2001.

Tutschku, Hans. *Still Air I.* 2013.

Varini, Felice. "I Am a Painter." Interview by Gil Dekel. *Poetic Mind*. October 2, 2014. Accessed May 22, 2018. http://www.poeticmind.co.uk/interviews-1/i-am-a-painter

———. "Rettangoli Gialli Concentrici Senza Angoli Al Suolo." *Felice Varini*. 1997. Accessed May 22, 2018. http://www.varini.org/varini/02indc/20indcd97.html

Waschka, Rodney. "Composing with Genetic Algorithms: GenDash." In *Evolutionary Computer Music*, edited by Eduardo Miranda and John Biles, 117-136. Berlin: Springer, 2007.

Weinberg, Gil. "Interconnected Musical Networks: Toward a Theoretical Framework." *Computer Music Journal 29,* no. 2 (2005): 23-39.

Wilson, Scott. *Flame*. Canadian Music Centre, 2006.

Wuorinen, Charles. *Cello Variations: For Cello Alone: (1970).* New York: C.F. Peters, 1976.

APPENDIX A: SCORE

Through a Window

for 4-6 distributed performers and electronics

3 nodes

Naithan Bosse

(2017)

Through A Window

4-6 distributed instruments and electronics
3 locations
Year: 2017
Duration: ~23 minutes

Movements

I. Stained Glass and Copper Wire 4 minutes
II. In Strange Lines and Distances 12 minutes
III. A Twisted Pair 7 minutes

Instrumentation

Sustaining soprano instrument 1
Ideal range: C4-C7.
Minimum required range: C4-G6.

Sustaining soprano instrument 2
Ideal range: F3-G6
Minimum required range: G3-G6.

Sustaining soprano instrument 3 **(optional)**
Range: C4-G6

Sustaining tenor instrument **(optional)**
Ideal range: C2-Bb5
Minimum required range: E2-C#5.

Percussion
- Mallet instrument
  Range: F3-F5
- Unpitched percussion instrument(s)

Keyboard (acoustic or electronic)

- Additionally, each performer must have 1 party balloon and 1 mobile phone.
- When possible, soprano 1, percussion, and keyboard should be distributed to different locations.
- Several versions of each part have been created to accommodate a variety of instrumental ranges and transpositions. Choose the version that most closely matches your instrument.
- If the optional instruments are not present, several sections should be performed by other instruments. See the cue list below.

Performance Instructions

Latency
The amount of time needed to send sound from one location to another over a network can be significant enough that synchronizing with remote musicians becomes impossible. With such a delay, performing in perfect alignment in your own location will sound terribly behind the beat in the remote locations. Typically everyone will unintentionally slow down to compensate for the delay. I have included the following directions in the score to help navigate performance over long distances and to help mitigate the effects of the delay.

Anchor: 

The musician acting as anchor is responsible for maintaining the tempo. The remote musicians align to the anchor musician's performance in their respective locations. As a result, the anchor musician will perceive the remote musicians as lagging behind the written score. It is the responsibility of the anchor musician to maintain the performance tempo without regard for the lagging remote musicians. If you find yourself decelerating while performing as anchor, you can compensate by performing with a slight and continuous accelerando. It is the responsibility of the non-anchor musicians to adapt to the acceleration in their own locations. Anchor sections are notated as cues in all parts.

Cue:
Since audio is transmitted more quickly than video over the network, visual conducting is ineffective. Several cues are embedded in the score (and parts) to help provide aural landmarks for all remote musicians. Cues should always be performed incisively and rigidly.

Clock-time
Several extended sections use stopwatches to measure time. A custom stopwatch mobile app allows computer 1 to remotely start, stop, and reset the stopwatch settings for all performers. The clock app should be open for the full performance. The clock rate can be increased or decreased to allow for slower rehearsal or faster performance while still displaying the correct clock-times.

Notation Legend

Accidentals:

In clock-time sections: *A la note*.
Accidentals modify only a single note plus any immediately following notes of the same pitch. (The figure is performed C#, B, C in clock-time sections)

In metered sections: Standard notation.
Modified pitches remain modified for the duration of the bar. (The figure is performed C#, B, C# in metered sections).

Start timer: Computer 1 starts the stopwatch app.

Time-code: Describes the running stopwatch time at the end of the current system and the start of the following system. **1' 12"**

Duration bar: Hold note for length of the bar

Mobile: Repeatedly (or continuously) perform the boxed music segment for the length of time indicated by the duration bar.

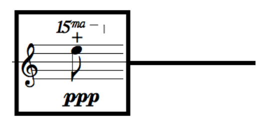

Jitter cue: Perform the contents of the 'jitter' mobile anytime the jitter button flashes in the phone app.

Recite: Recite the indicated text in a near-whisper. Stand close to the mic.

White noise: This can be interpreted in a variety of ways depending on your instrument. For example, a bowed string instrument may choose to interpret the direction by performing an extreme sul. pont. or by bowing the body of the instrument. A wind instrument may interpret the direction by performing a breath tone.

Dash arrow: Gradually modulate from one playing style to another. Staccato ------------------------------► Legato

Pulsed crescendo/diminuendo: Perform several small crescendi/diminuendi embedded inside a larger-scale crescendo/diminuendo.

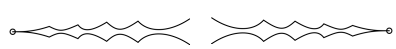

Slowly deflate balloon: Use the air stream to create a sound like ocean waves. Avoid directing the air stream directly into the microphone as this will create an overly bassy sound. Instead, direct the air-stream into the mic at an oblique angle.

Unpitched: Percussive. This can be interpreted in several ways depending on your instrument. A wind instrument may use key clicks while a bowed string instrument may pluck a muted string.

Muted: If performing on an acoustic piano, reach inside the body of the instrument and mute the string with your palm. Avoid exciting ringing harmonics. 

Pitched:  Partially pitched:  Unpitched: 

Similar to the 'white noise' note-head, the pitched/partially-pitched/unpitched directions can be interpreted creatively based on what is appropriate for your instrument.

Feathered beams:

Accelerate 

Decelerate 

Highest possible pitch on your instrument: 

Pitch boundaries: Improvise within the indicated pitch ranges.

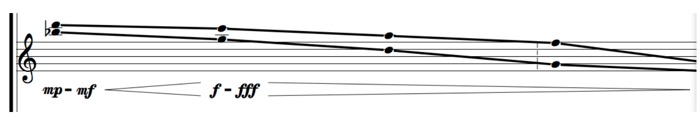

Chromatic Cluster: Perform all pitches between the outer pitches (inclusive).
(C1-Bb2 in the figure)

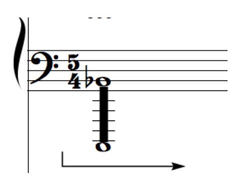

White-note cluster: Signified by the natural sign above the cluster, perform all white-notes between the outer pitches (inclusive).
(Treble clef E5-F6 in the figure)

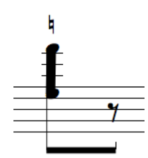

<u>Movement 1, pp. 3-9 and movement 3, pp. 61-66.</u>
Each time you are instructed to perform a cell, select and perform one fragment from the available pool of cells notated above the staff as a numbered "generation." The cells are notated proportionally with the duration specified in seconds at the end of each cell.
(The figure shows a 4 second long cell from generation 3).

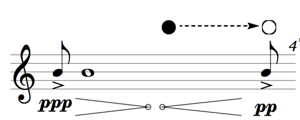

Electronics

Hardware requirements

- 1 microphone per performer. (DPA mics are preferred)
- 1 phone per performer (including computer performers)
- At each location:
  - 1 computer
  - 8-channel audio system (4 channels is also possible)
  - High-speed wired internet connection (fiber-optic).

Software requirements

- Max performance patches (available at www.naithanbosse.com/ThroughAWindow)
  - TAW_node1.maxpat
  - TAW_node2.maxpat
  - TAW_node3.maxpat
  - nb.toolbox.zip
- MaxComm/Mira (available on the Apple App Store)
  - TAW.maxcomm
- Skype
- Jack Audio Connection Kit (with Qjackctrl)
- Select one of the following and follow the corresponding setup instructions (below).
  - Jacktrip
  - ArtsMesh

Speaker configuration (stereo pairs)

8-channels	4 channels
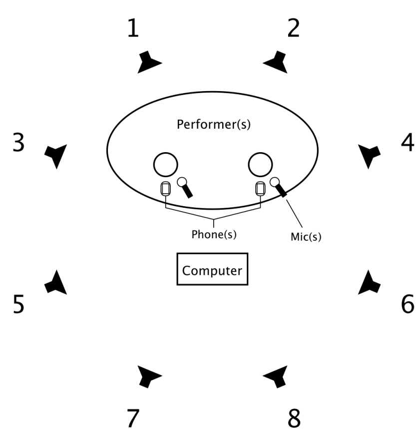	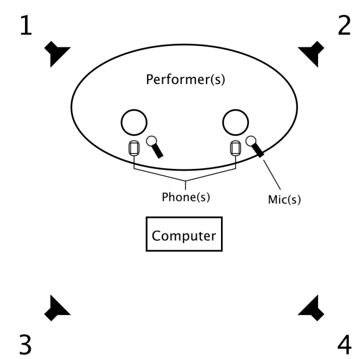

Audio Setup – Jacktrip Version

1. Make sure that Max is closed and you are connected to the correct network.
2. Open QJackCtrl
3. Click "Setup"

4. All nodes agree upon a Sample Rate and Frames/Period size. (Sample Rate: ________ Frames: ______)
5. Set "Driver" to coreaudio and click OK

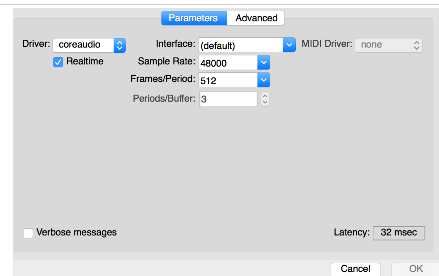

6. In the main window, click Start to activate the Jack server.
7. Open Max and follow the setup instructions.
8. Start a Skype call with the remote locations.
9. In QJackCtrl, open the Connect window.
10. When connecting to a remote node, Jacktrip will automatically cross connect your system input/output with the remote location. Make sure your audio levels are set low or even muted to avoid feedback.

11. Node 1: In terminal, type
"jacktrip -s -n2 -b24 -r3 --clientname node2"
Wait until node 2 is ready and hit enter.
Click "Disconnect all" in the QJackCtrl connect window.

11. Node 2: In terminal, type
"jacktrip -c [enter node1 ip here] -n2 -b24 -r3 --clientname node1"
Wait until node 1 is ready and hit enter.
Click "Disconnect all" in the QJackCtrl connect window.

12. Node 1: Open a new terminal window and type
"jacktrip -s -n2 -b24 -r3 --clientname node3 -o10"
Wait until node 3 is ready and hit enter.
Click "Disconnect all" in the QJackCtrl connect window.

12. Node 3: In terminal, type
"jacktrip -c [node1 ip here] -n2 -b24 -r3 --clientname node1 -o10"
Wait until node 1 is ready and hit enter.
Click "Disconnect all" in the QJackCtrl connect window.

13. Node 2: Open a new terminal window and type
"jacktrip -s -n2 -b24 -r3 --clientname node3 -o20"
Wait until node 3 is ready and hit enter.
Click "Disconnect all" in the QJackCtrl connect window.

13. Node 3: In terminal, type
"jacktrip -c [node1 ip here] -n2 -b24 -r3 --clientname node2 -o20"
Wait until node 2 is ready and hit enter.
Click "Disconnect all" in the QJackCtrl connect window.

14. Create the following audio connections in the Connect window in QjackCtrl by highlighting the desired inputs and outputs and clicking the connect button.

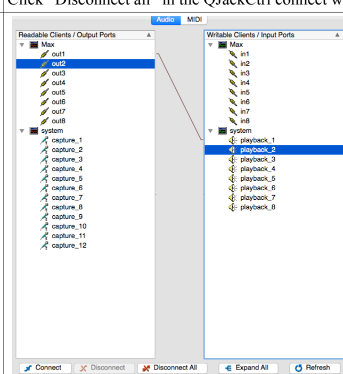

Node 1:	Node 2:	Node 3:
System receive 1-n => Max send 1-n	System receive 1-n => Max send 1-n	System receive 1-n => Max send 1-n
Max receive 1-8 => System send 1-8	Max receive 1-8 => System send 1-8	Max receive 1-8 => System send 1-8
Max receive 9-10 => Node 2 send 1-2	Max receive 9-10 => Node 1 send 1-2	Max receive 9-10 => Node 1 send 1-2
Max receive 11-12 => Node 3 send 1-2	Max receive 11-12 => Node 3 send 1-2	Max receive 11-12 => Node 2 send 1-2
Node 2 receive 1-2 => Max send 5-6	Node 1 receive 1-2 => Max send 5-6	Node 1 receive 1-2 => Max send 5-6
Node 3 receive 1-2 => Max send 7-8	Node 3 receive 1-2 => Max send 7-8	Node 2 receive 1-2 => Max send 7-8

15. Follow the instructions in Max to perform sound-check.
16. Mute Skype during performance.

Cue list

If either or both of the optional instrumental parts are missing, then the following excerpts should be covered by other instruments according to the list below. These cues are also notated and labelled in the appropriate parts.

A Copper Wire

Missing Instrument...	...at excerpt...	...is covered by
Soprano 3	0:00-0:24	Keyboard
Tenor	1:12-1:36	Computer 1

In Strange Lines and Distances

Missing Instrument...	...at excerpt...	...is covered by
Soprano 3	mm. 5-12	Computer 1
Tenor	mm. 6-13	Soprano 2
Tenor	mm. 22-33	Soprano 2
Soprano 3	m. 139	Percussion
Tenor	mm. 140-148	Percussion (transposed)
Soprano 3	m. 150	Keyboard
Soprano 3/Tenor	mm. 157-169	Computer 1 (change processing)
Tenor	m. 203	Keyboard
Soprano 3	m. 204	Percussion
Tenor	m. 206	Percussion/Soprano 2/Keyboard
Soprano 3	mm. 207-209	Percussion
Tenor	mm. 211-213	Computer 1
Tenor	mm. 215-219	Keyboard
Soprano 3	mm. 215-219	Soprano 2
Soprano 3	mm. 221-222	Soprano 1
Soprano 3	mm. 223-224	Soprano 2
Soprano 3	m. 233	Percussion
Soprano 3	mm. 238-268	Computer 2
Tenor	mm. 231-258	Computer 1
Tenor	m. 260	Keyboard
Tenor	mm. 264-269	Soprano 2 (transposed)
Tenor	mm. 273-274	Percussion (transposed)
Tenor	mm. 277-279	Soprano 1 (transposed)

A Twisted Pair

Missing Instrument...	...at excerpt...	...is covered by
Soprano 3	m. 2	Soprano 2
Soprano 3	m. 6	Percussion
Soprano 3	mm. 7-10	Keyboard
Tenor	mm. 16-19	Percussion
Soprano 3	mm. 40-71	Computer 1
Tenor	mm. 40-71	Computer 2
Soprano 3/Tenor	mm. 75-103	Soprano 2/Computer 1
Tenor	m. 129	Percussion
Tenor	mm. 133-134	Keyboard
Tenor	mm. 137-140	Keyboard
Tenor	mm. 141-146	Soprano 2
Tenor	mm. 149-152	Soprano 1

I. Stained Glass

Naithan Bosse

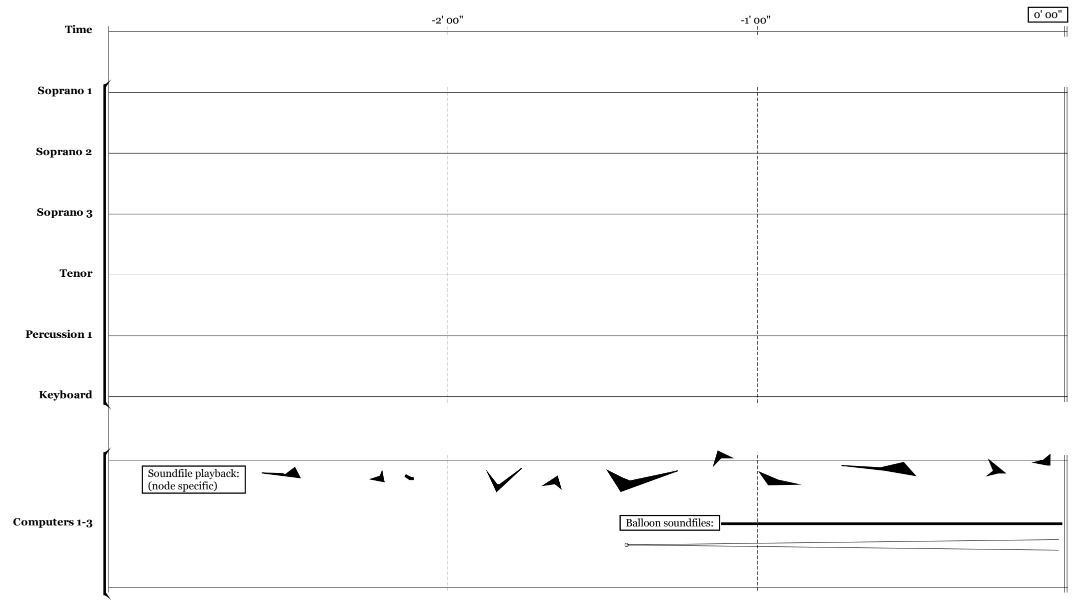

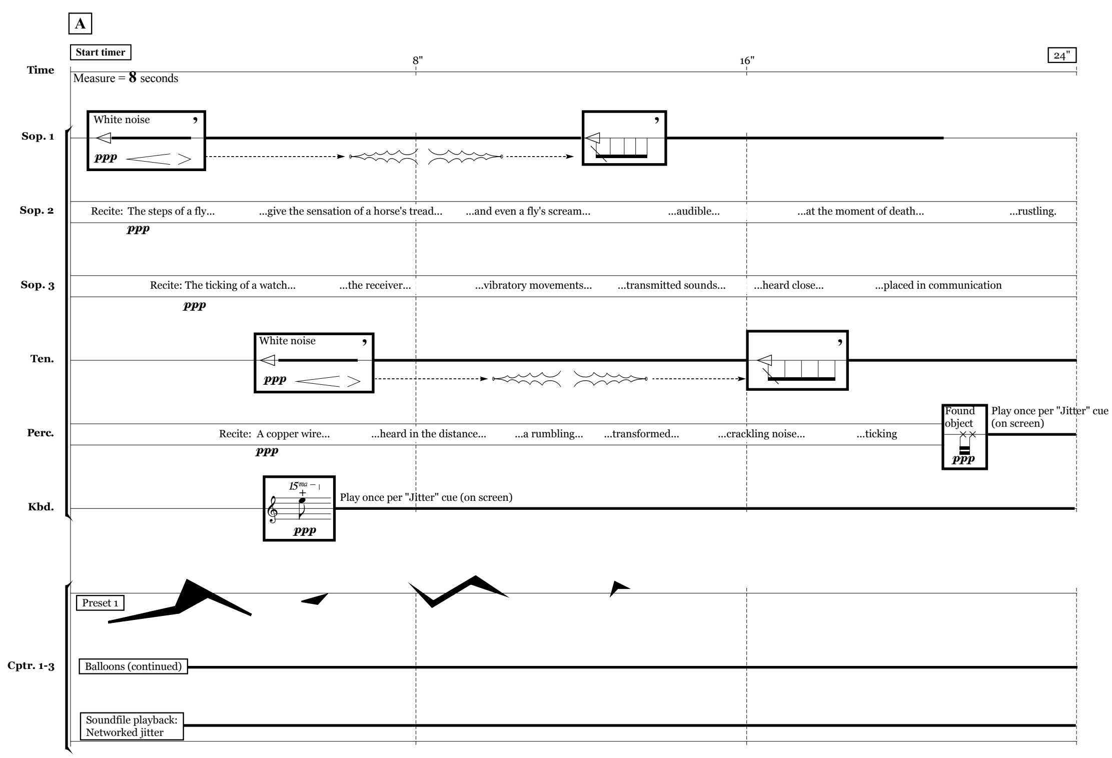

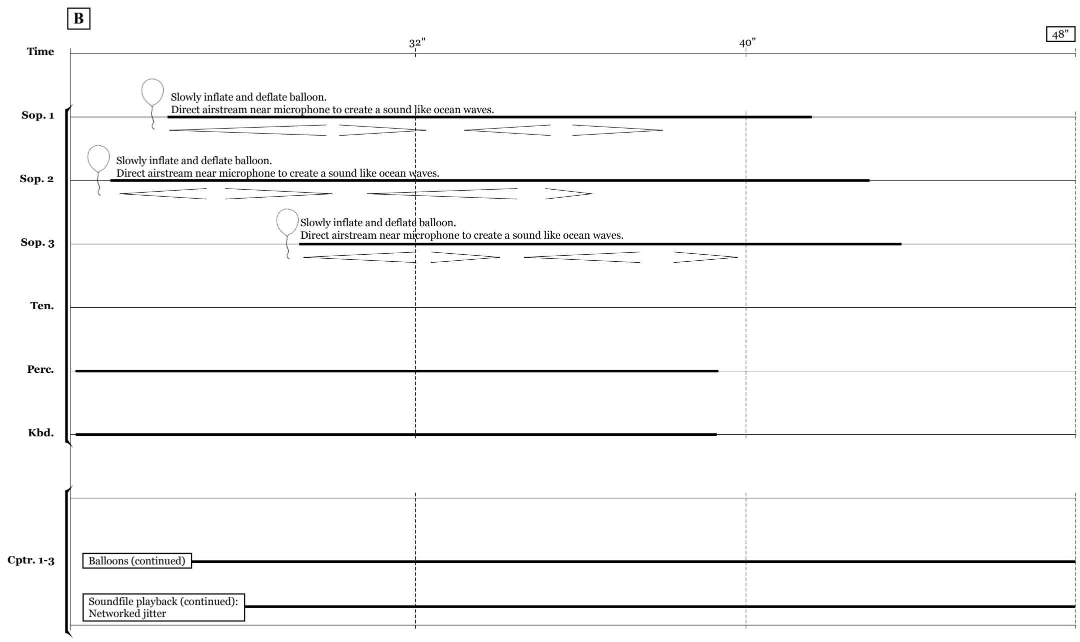
B
Time
32"
40"
48"
Sop. 1
Slowly inflate and deflate balloon.
Direct airstream near microphone to create a sound like ocean waves.
Sop. 2
Slowly inflate and deflate balloon.
Direct airstream near microphone to create a sound like ocean waves.
Sop. 3
Slowly inflate and deflate balloon.
Direct airstream near microphone to create a sound like ocean waves.
Ten.
Perc.
Kbd.
Cptr. 1-3
Balloons (continued)
Soundfile playback (continued):
Networked jitter

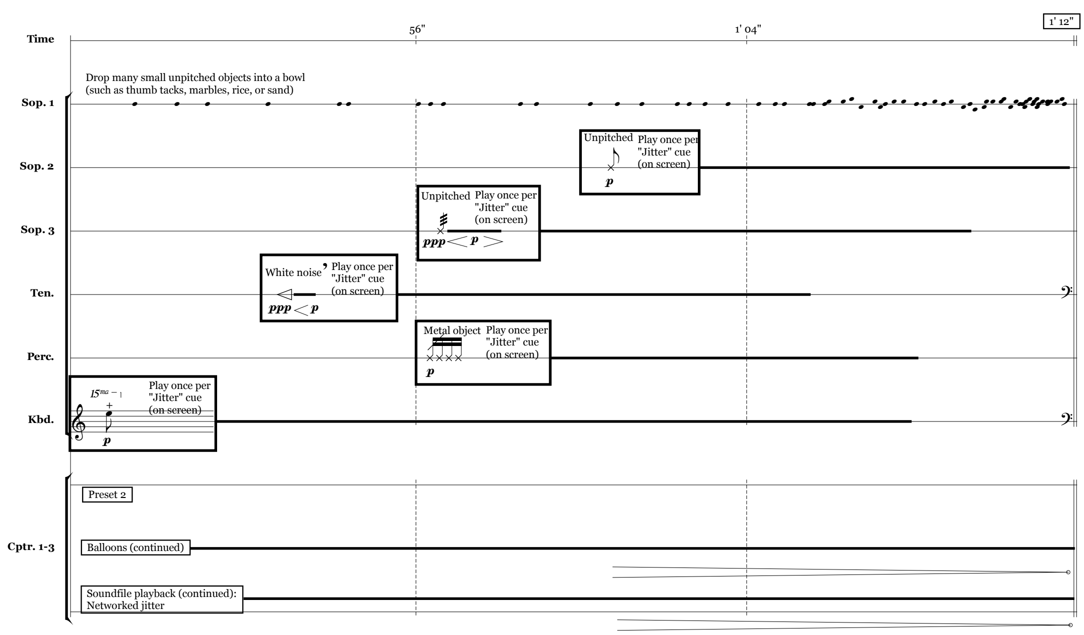
Time
56"
1' 04"
1' 12"
Sop. 1
Drop many small unpitched objects into a bowl
(such as thumb tacks, marbles, rice, or sand)
Sop. 2
Unpitched
Play once per "Jitter" cue (on screen)
p
Sop. 3
Unpitched
Play once per "Jitter" cue (on screen)
ppp p
Ten.
White noise
Play once per "Jitter" cue (on screen)
ppp p
Perc.
Metal object
Play once per "Jitter" cue (on screen)
p
Kbd.
15ma
Play once per "Jitter" cue (on screen)
p
Cptr. 1-3
Preset 2
Balloons (continued)
Soundfile playback (continued):
Networked jitter

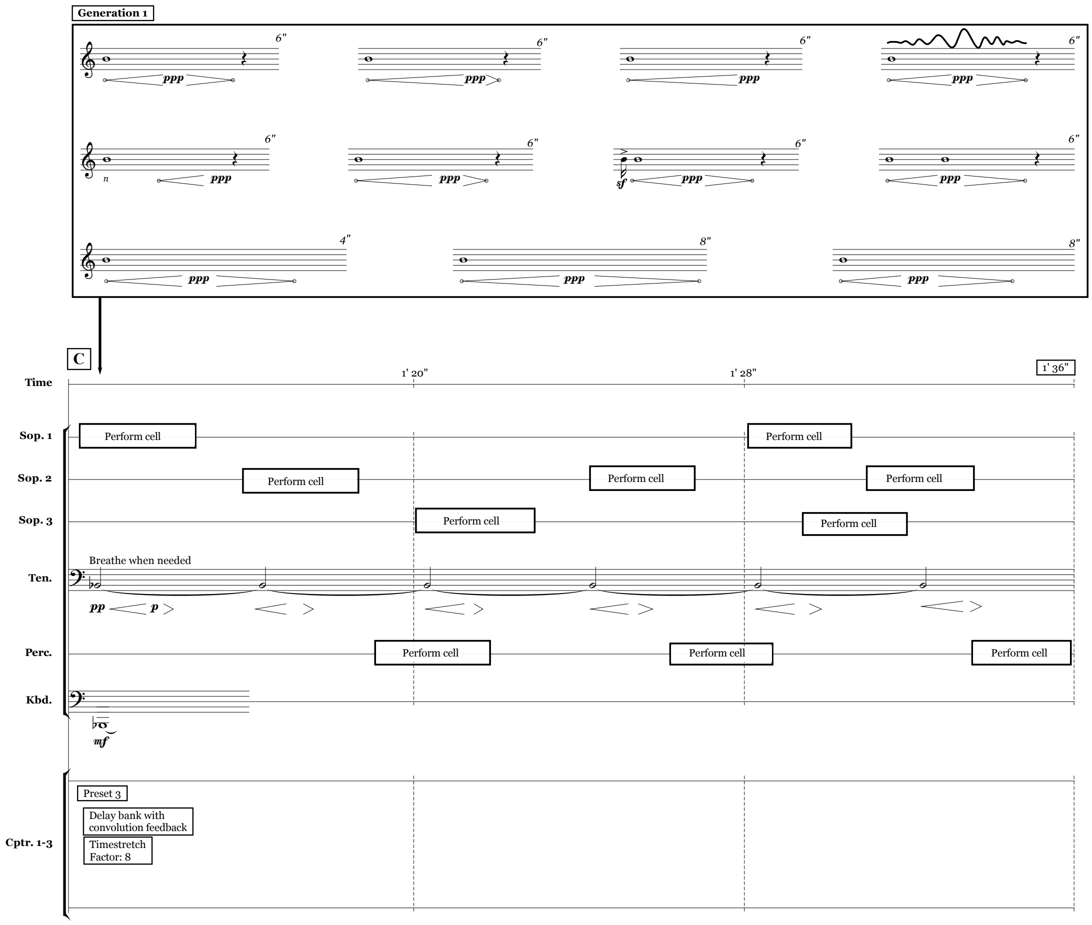
Generation 1
6"
ppp
4"
8"
sf
n
C
Time
1' 20"
1' 28"
1' 36"
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
Perform cell
Breathe when needed
pp
p
mf
Preset 3
Delay bank with convolution feedback
Timestretch Factor: 8

4

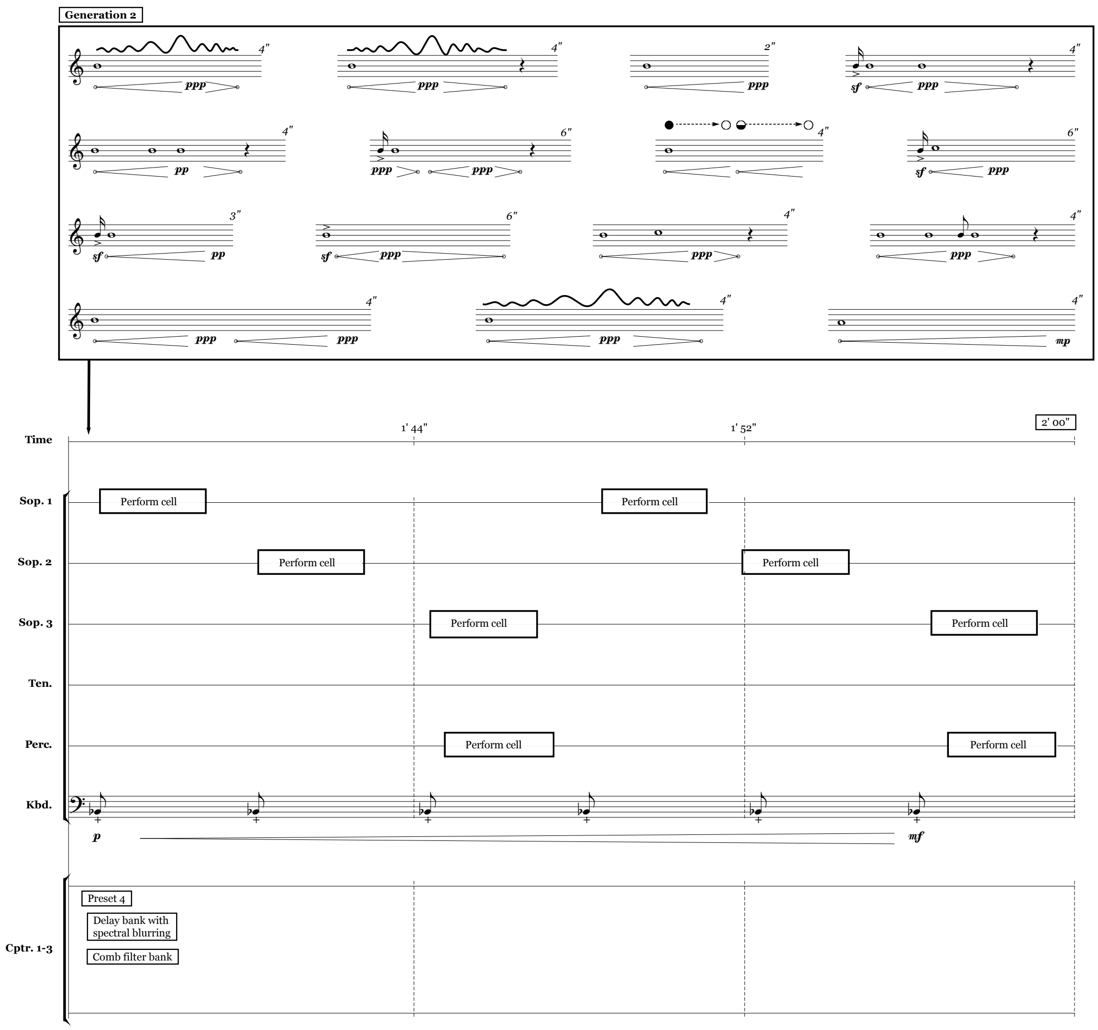
Generation 2
4"
ppp
4"
ppp
2"
ppp
4"
sf
ppp
4"
pp
6"
ppp
ppp
4"
6"
sf
ppp
3"
sf
pp
6"
sf
ppp
4"
ppp
4"
ppp
4"
ppp
ppp
4"
ppp
4"
mp
Time
1' 44"
1' 52"
2' 00"
Sop. 1
Perform cell
Perform cell
Sop. 2
Perform cell
Perform cell
Sop. 3
Perform cell
Perform cell
Ten.
Perc.
Perform cell
Perform cell
Kbd.
p
mf
Cptr. 1-3
Preset 4
Delay bank with spectral blurring
Comb filter bank

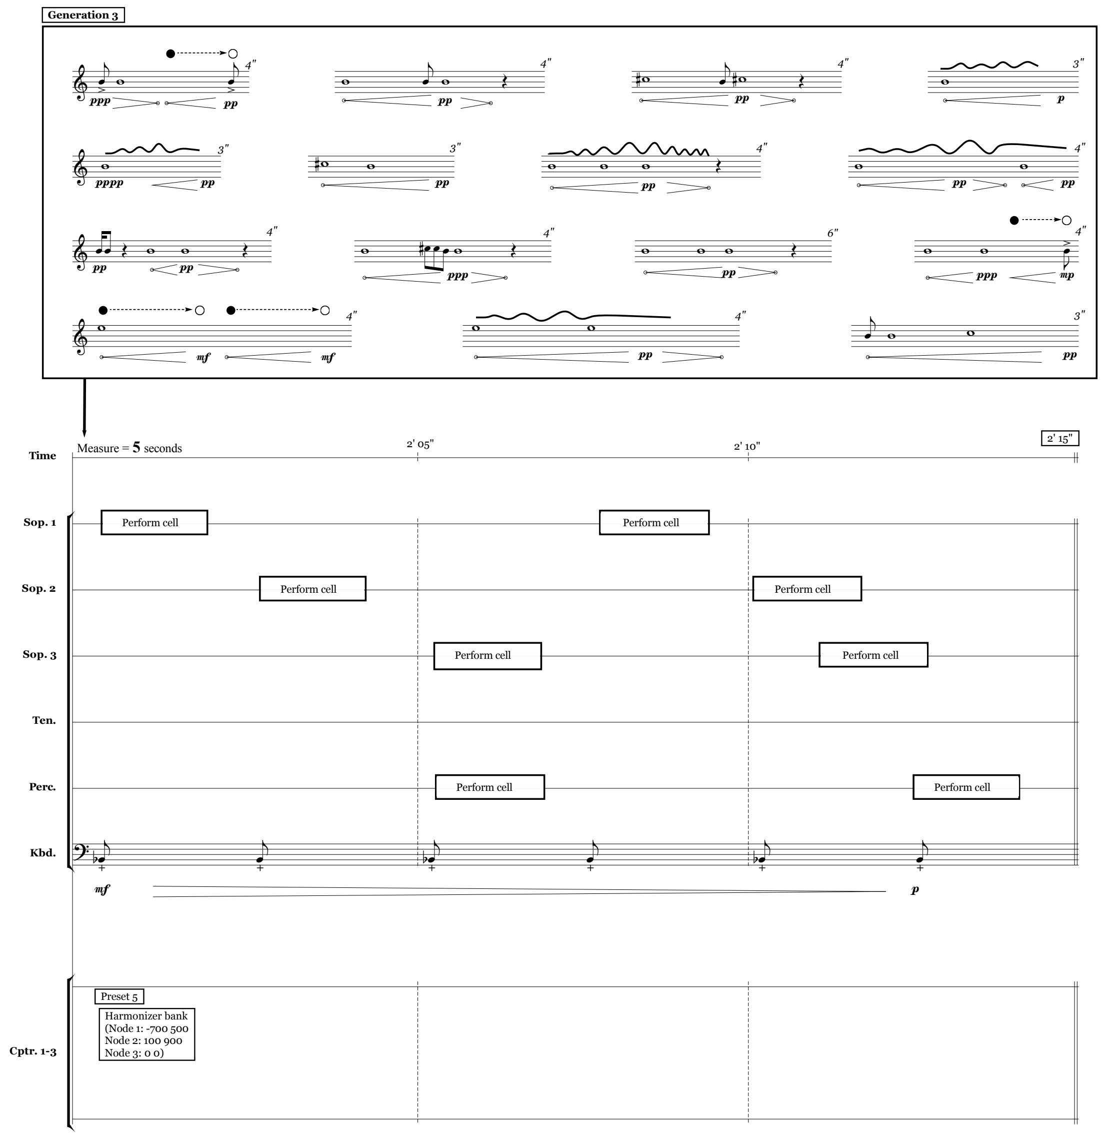
Generation 3
Measure = 5 seconds
2' 05"
2' 10"
2' 15"
Time
Sop. 1
Perform cell
Perform cell
Sop. 2
Perform cell
Perform cell
Sop. 3
Perform cell
Perform cell
Ten.
Perc.
Perform cell
Perform cell
Kbd.
mf
p
Preset 5
Harmonizer bank
(Node 1: -700 500
Node 2: 100 900
Node 3: 0 0)
Cptr. 1-3

6

Generation 4

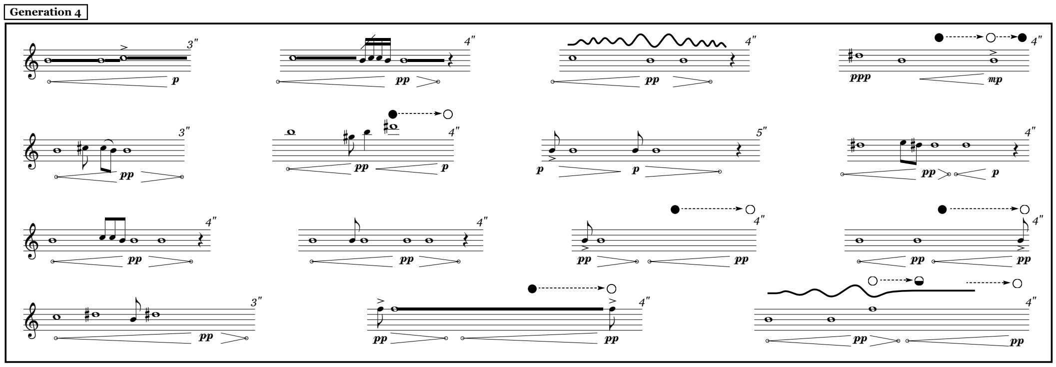

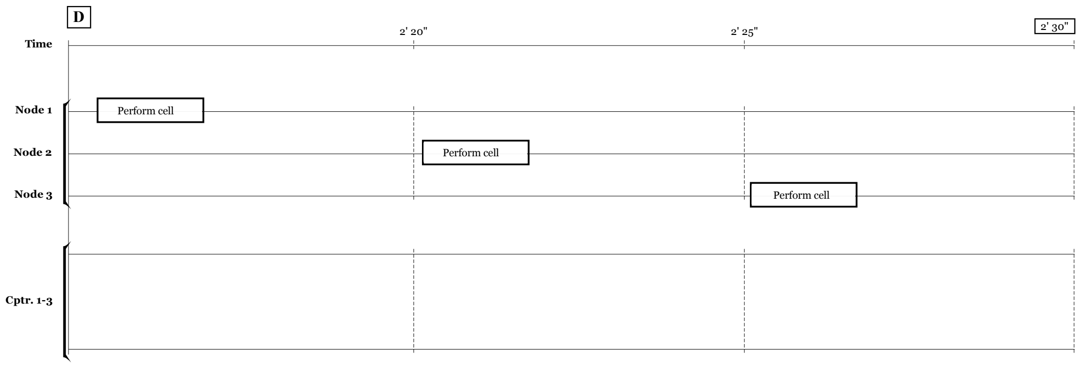

Generation 5

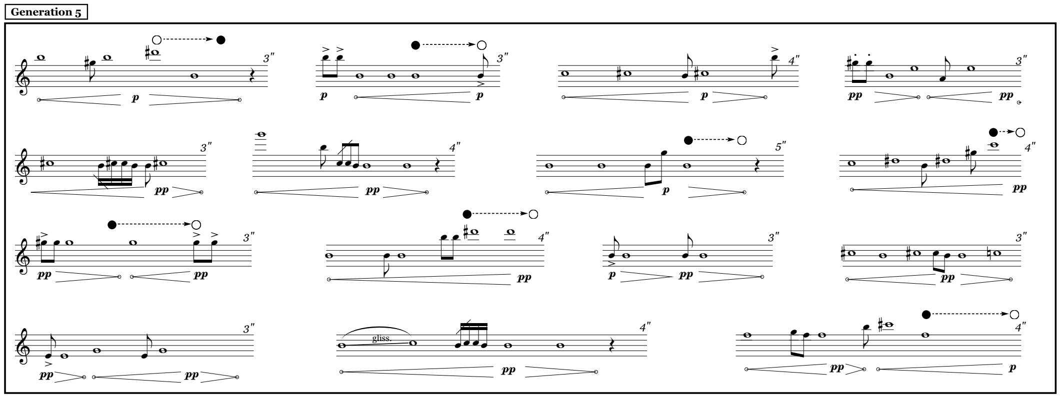

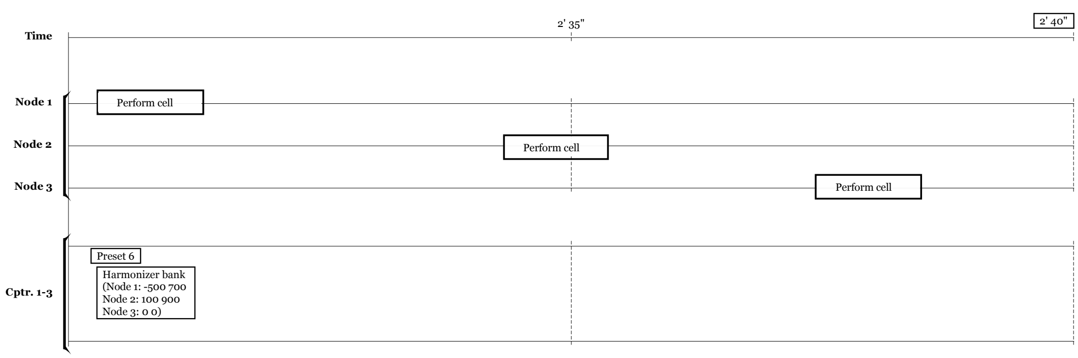

Generation 6

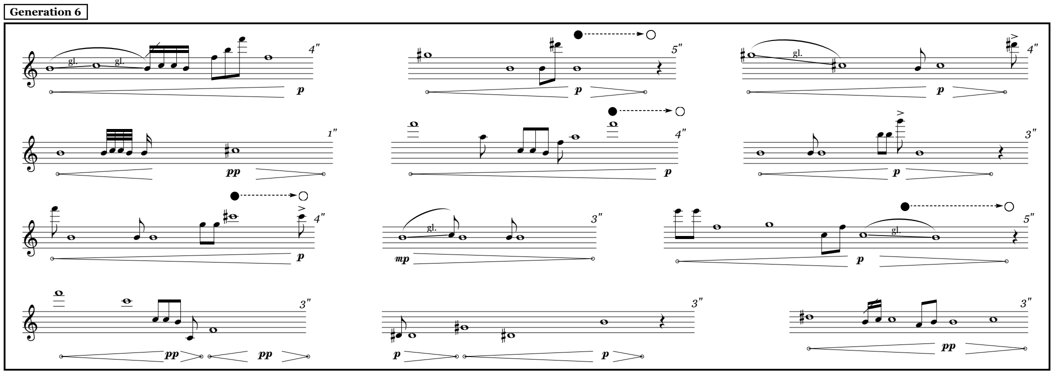

Generation 7

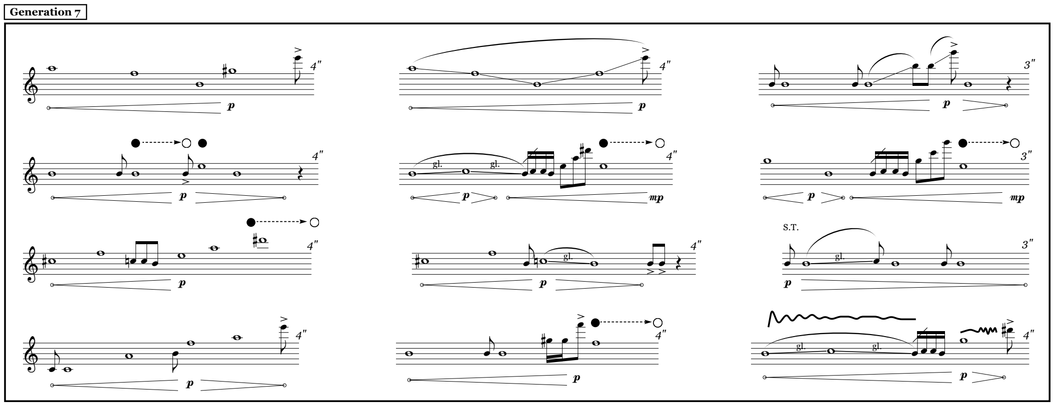

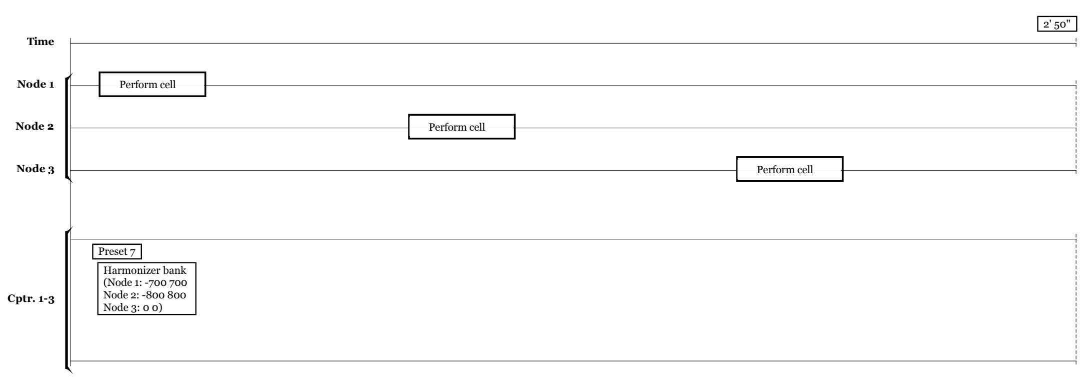

8

Generation 8

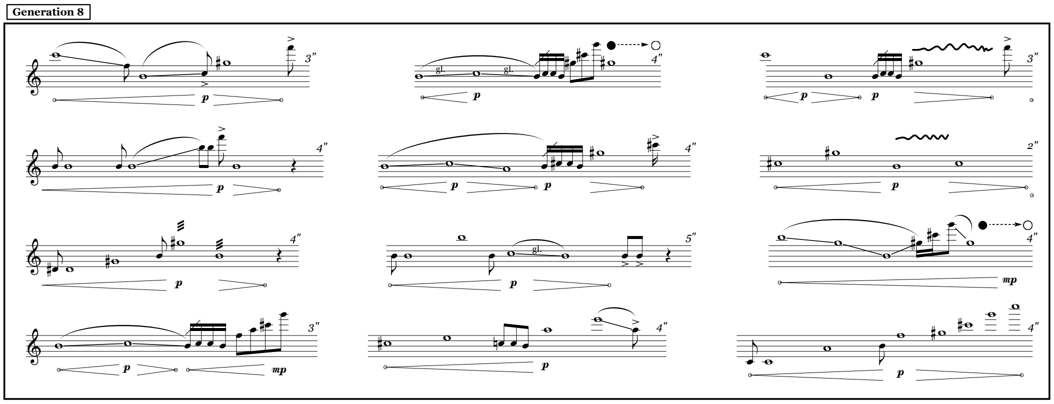

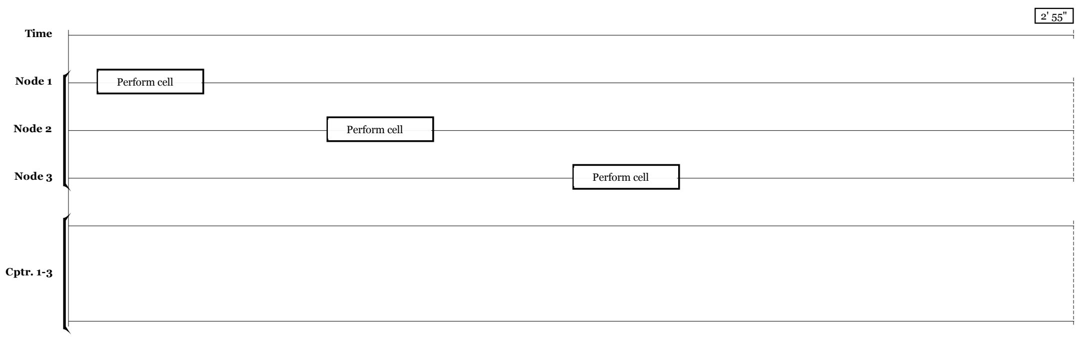

Generation 9

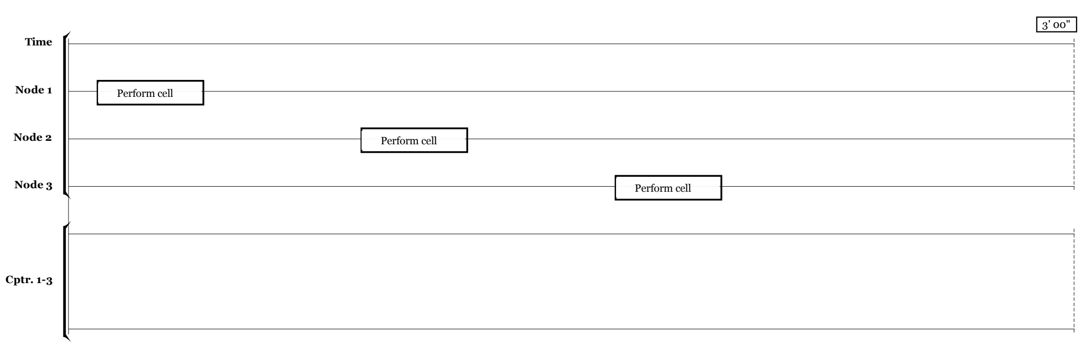

Generation 10

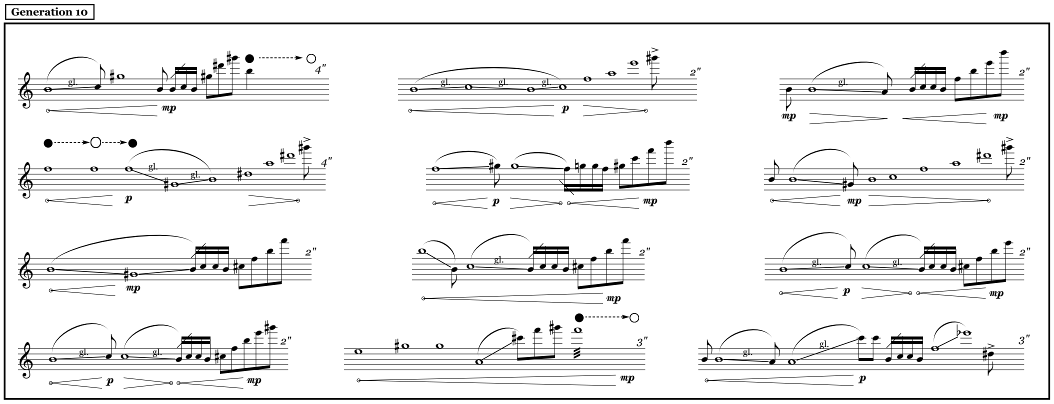

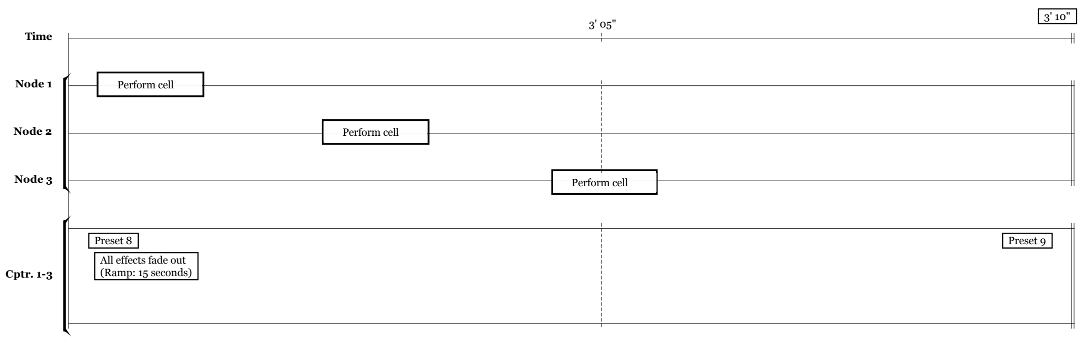

10

II. In Strange Lines and Distances

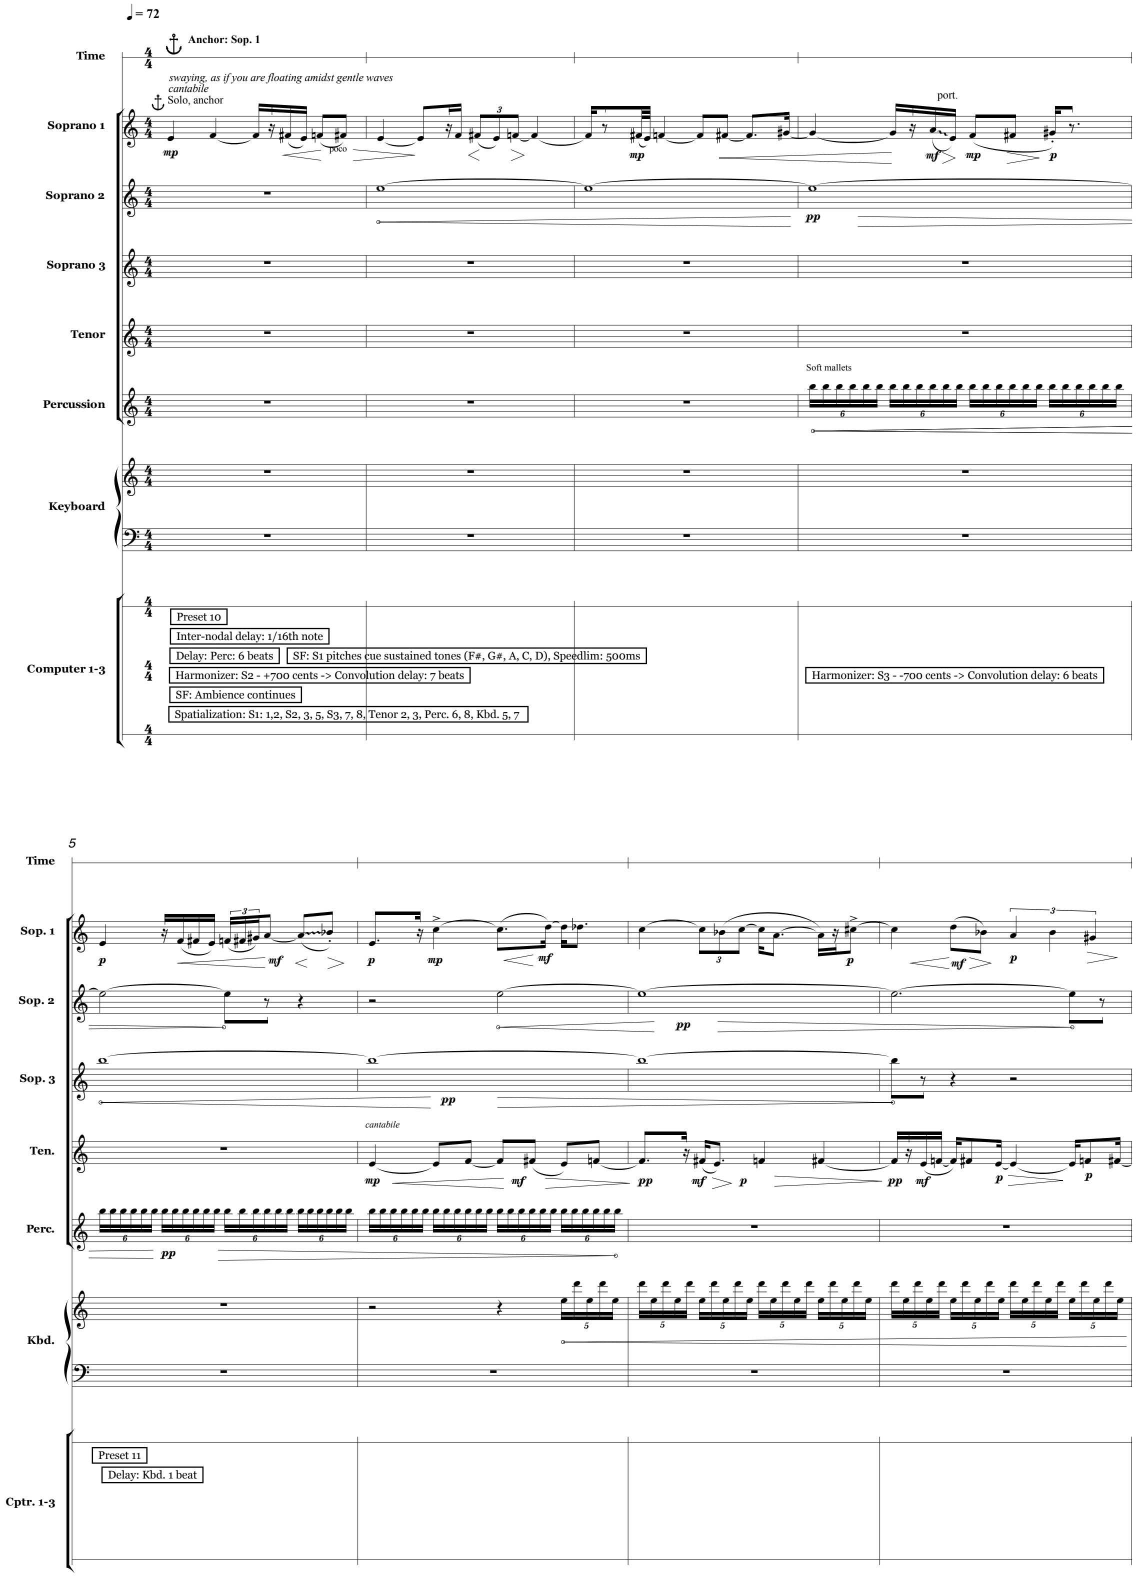


9
Time
End anchor
piu mosso
End solo, end anchor
Sop. 1
subito
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Preset 12
Harmonizer: Perc - +10 cents
Cptr. 1-3
E
♩= 84
14
Anchor: Kbd.
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Solo, anchor
like windchimes swaying in the wind
Kbd.
Use pedal liberally ad. lib.
Preset 13
Delay with spectral blurring: Kbd.
Timestretch: Kbd.
Cptr. 1-3


18
Time
End anchor
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
End solo,
end anchor
Cptr. 1-3
G#5 triggers sample -> speedlim 2000ms
F
= 96
22
Anchor: Ten.
Soli
Soli, anchor
gently swaying, as if floating amidst waves
Medium hard mallets
Preset 14
Fade out processing
(Ramp: 20 seconds)


27
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
port.
32
End anchor
Anchor: Sop. 2
Soli, anchor
End solo,
end anchor
Preset 15
Spatialization: S1 - 1,3,5, S2 - 2,4,6


37
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
15
poco a poco cresc.

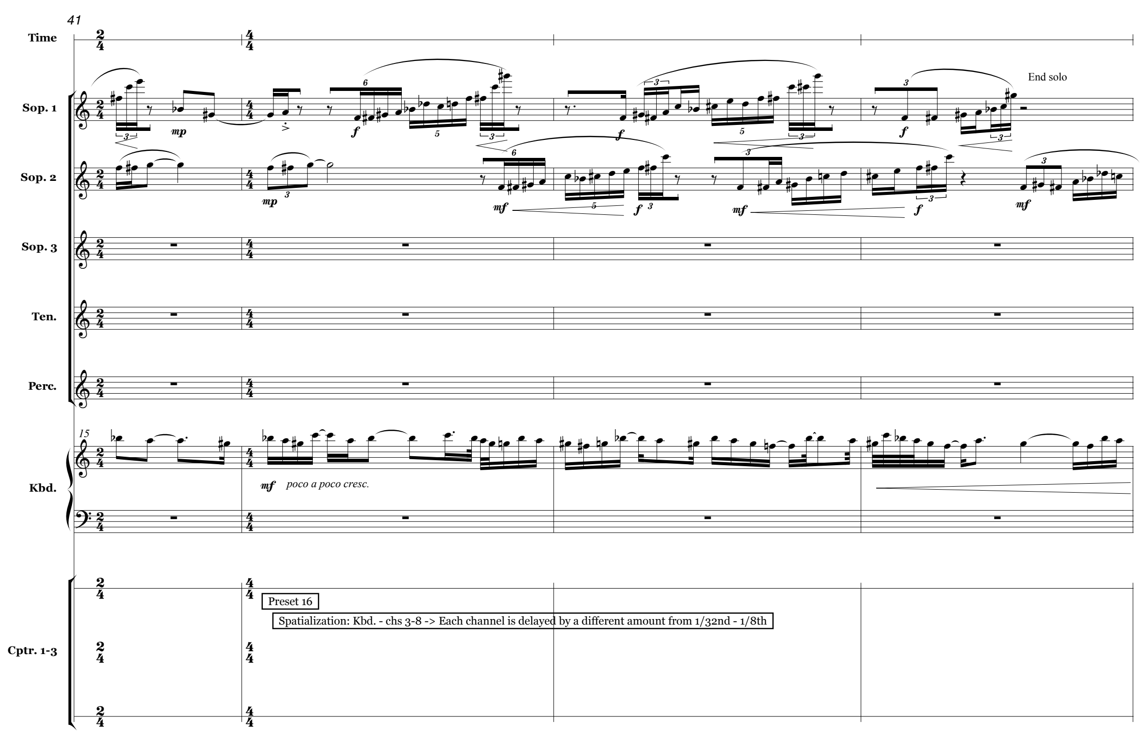
41
Time
Sop. 1
End solo
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
15
poco a poco cresc.
Preset 16
Spatialization: Kbd. - chs 3-8 -> Each channel is delayed by a different amount from 1/32nd - 1/8th

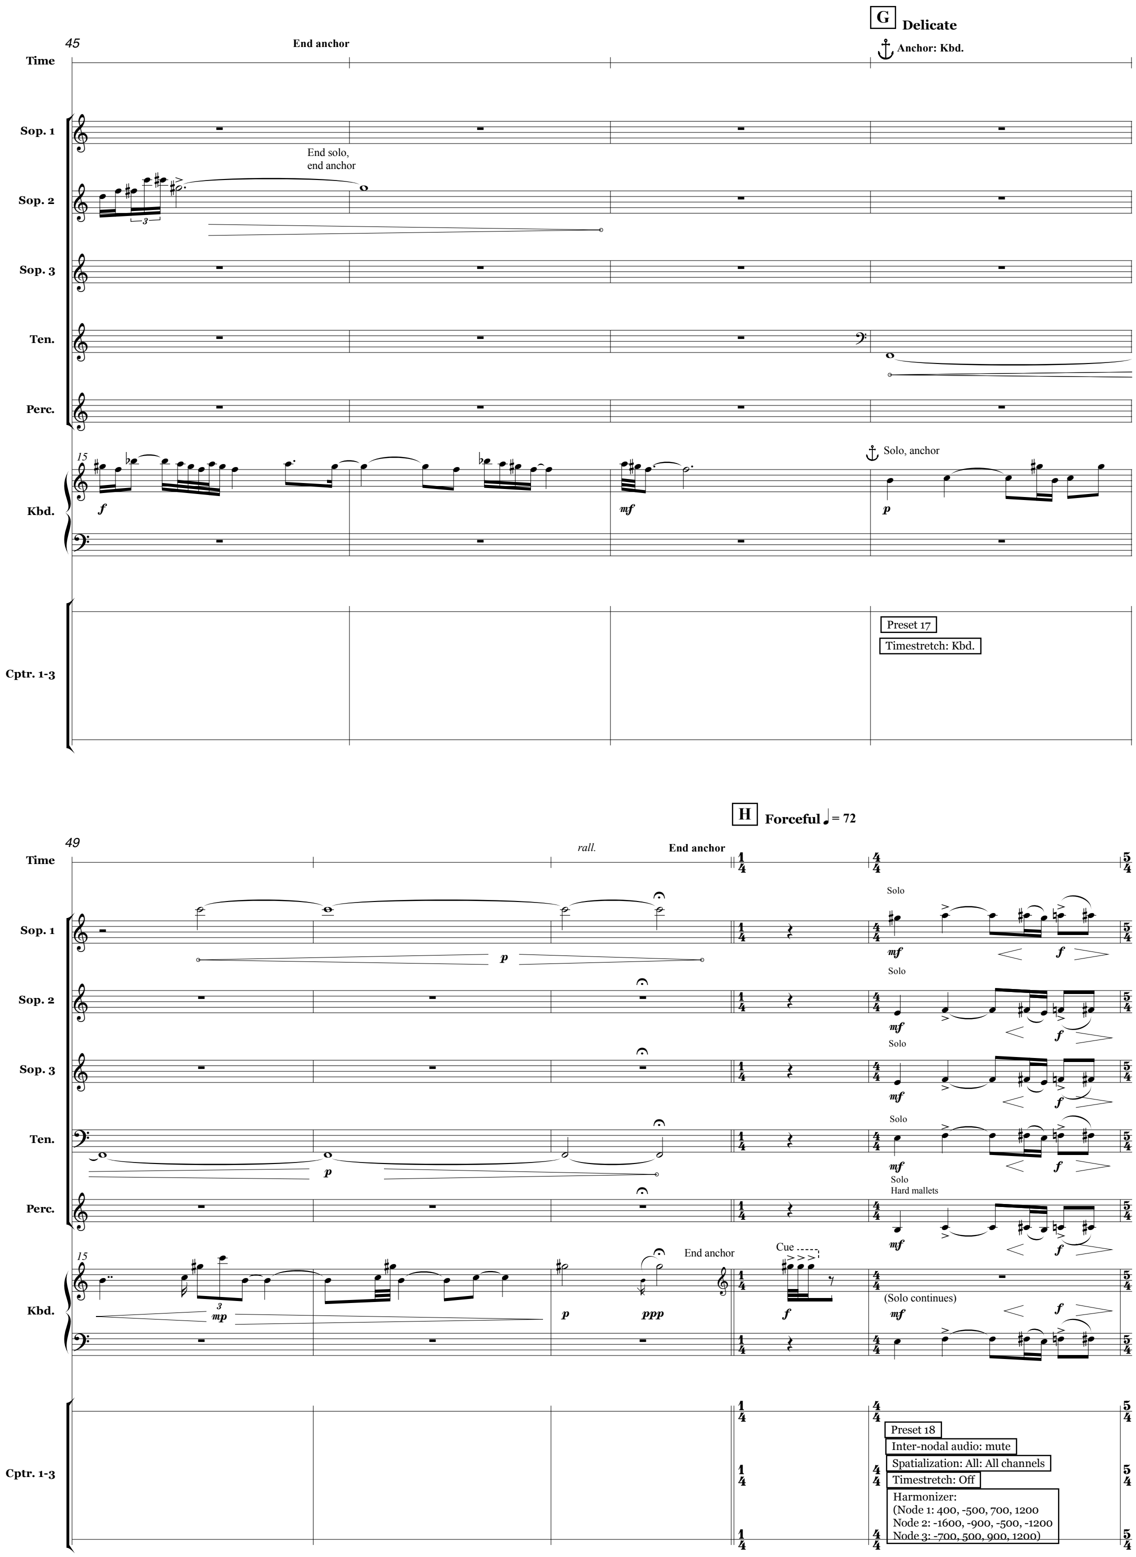
G Delicate
Anchor: Kbd.
45
End anchor
Time
Sop. 1
End solo,
end anchor
Sop. 2
Sop. 3
Ten.
Perc.
Solo, anchor
Kbd.
Preset 17
Timestretch: Kbd.
Cptr. 1-3
H Forceful ♩= 72
49
rall.
End anchor
Time
Solo
Sop. 1
Sop. 2
Sop. 3
Ten.
Solo
Hard mallets
Perc.
End anchor
Cue
(Solo continues)
Kbd.
Preset 18
Inter-nodal audio: mute
Spatialization: All: All channels
Timestretch: Off
Harmonizer:
(Node 1: 400, -500, 700, 1200
Node 2: -1600, -900, -500, -1200
Node 3: -700, 500, 900, 1200)
Cptr. 1-3


54
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
fp
mf
f

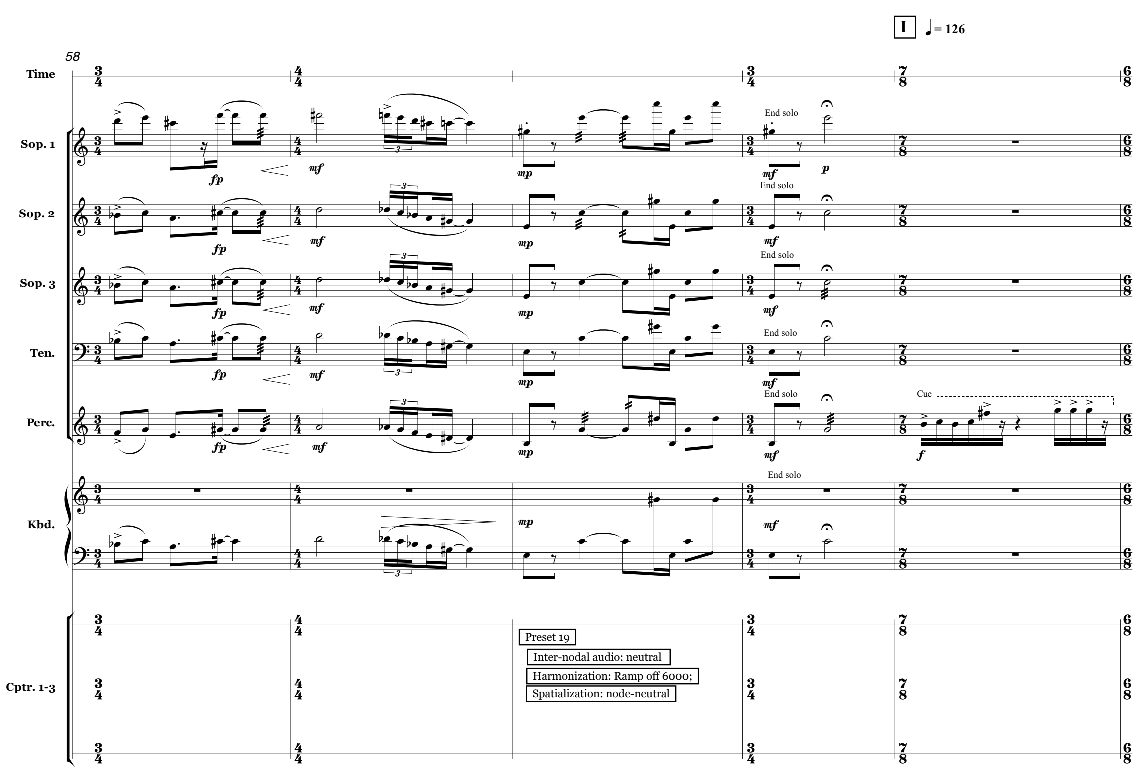
I
♩= 126
58
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
End solo
Cue
fp
mf
mp
p
f
Preset 19
Inter-nodal audio: neutral
Harmonization: Ramp off 6000;
Spatialization: node-neutral

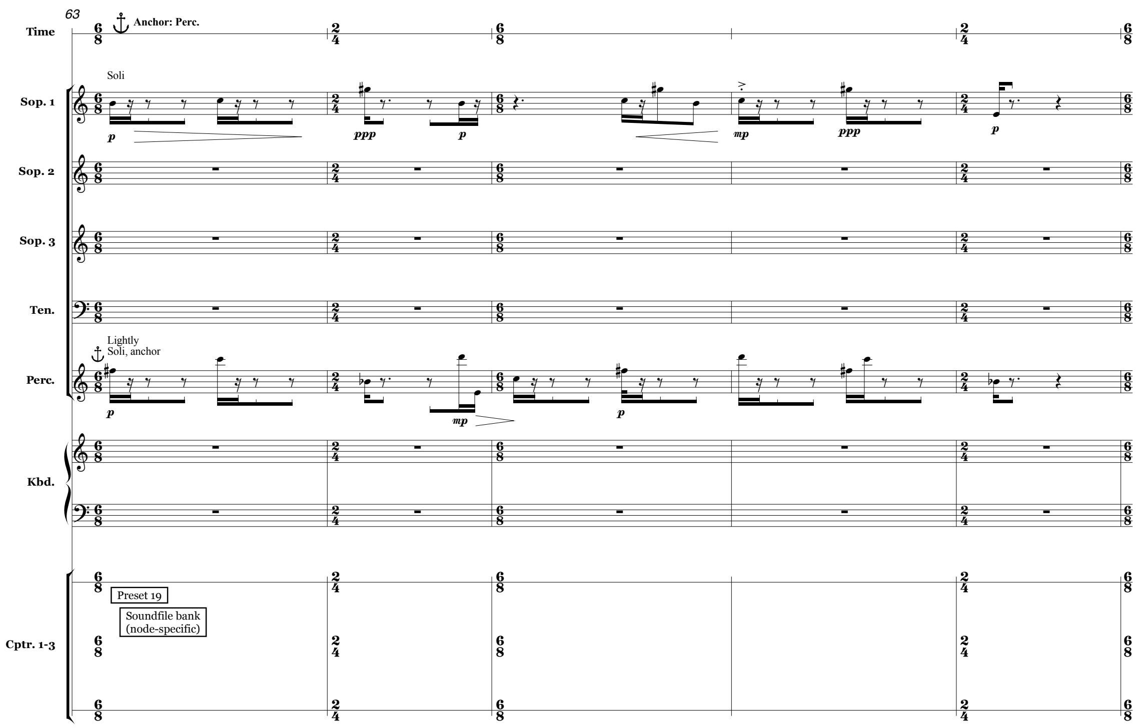
63
Anchor: Perc.
Time
Soli
Sop. 1
Sop. 2
Sop. 3
Ten.
Lightly
Soli, anchor
Perc.
Kbd.
Preset 19
Soundfile bank
(node-specific)
Cptr. 1-3


68
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3


73
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
79
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3


85
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
90


J
95
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
100
piu mosso


105
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
l.v.
110
piu mosso

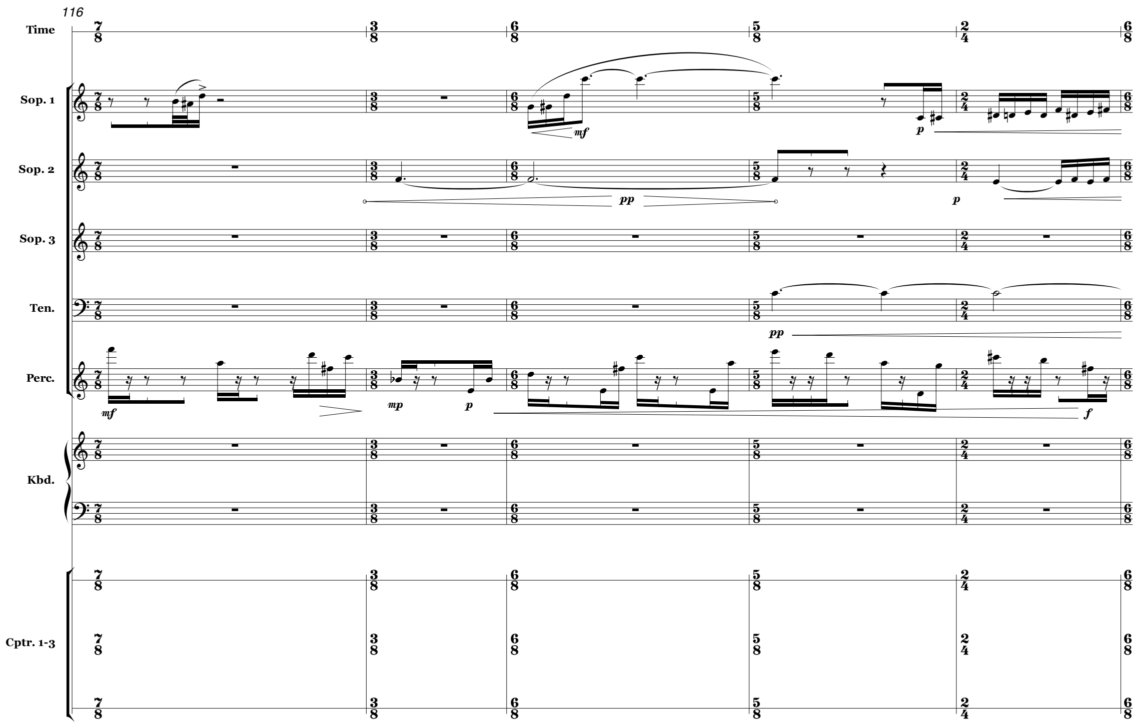
116
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3

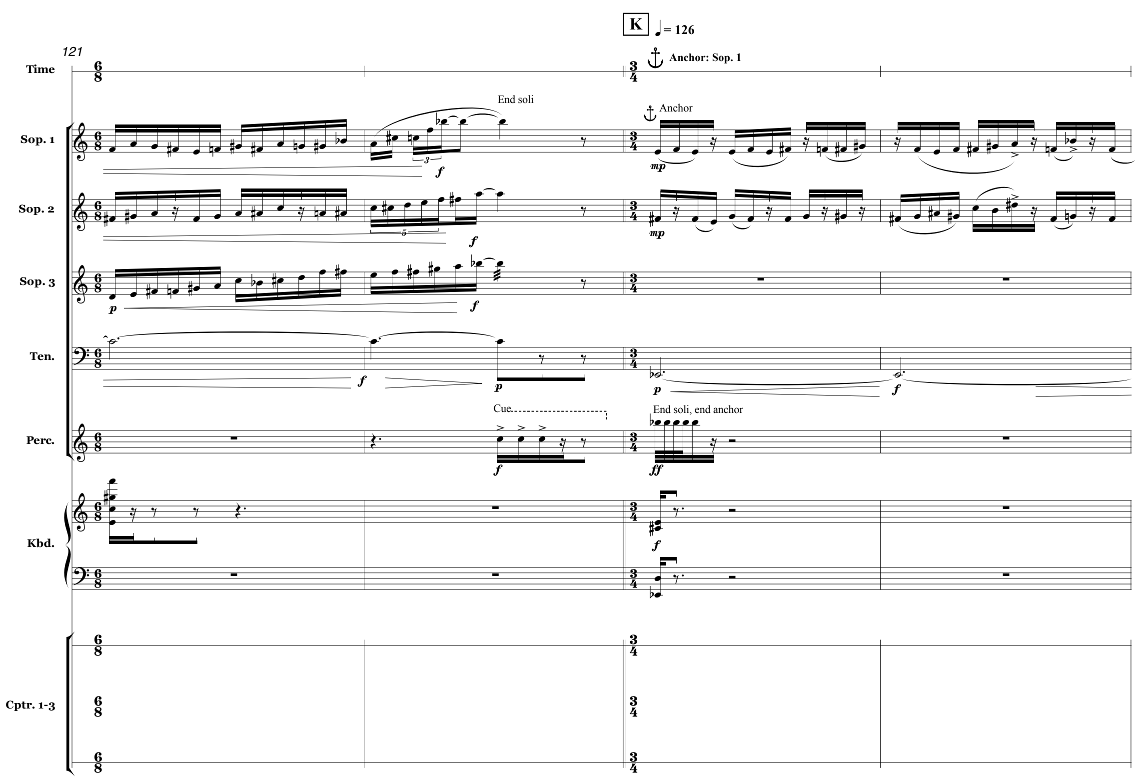
K
♩= 126
Anchor: Sop. 1
121
Time
Sop. 1
End soli
Anchor
Sop. 2
Sop. 3
Ten.
Perc.
Cue
End soli, end anchor
Kbd.
Cptr. 1-3


125
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
130
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3


135
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
mf
mp
mf
mf
3
140
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
mp
poco a poco cresc.
mp
mf
mp
mf
mp
p
f
p
f
fp
mp
p
f
p
f
p
f


145
Time
Sop. 1
mf
f
mf
Sop. 2
mf
mp
f
mf
Sop. 3
Ten.
Perc.
f
Kbd.
Cptr. 1-3


150
Time
Sop. 1
mp
f
Sop. 2
mp
f
mf
f
Sop. 3
Ten.
Perc.
Kbd.
3
Cptr. 1-3

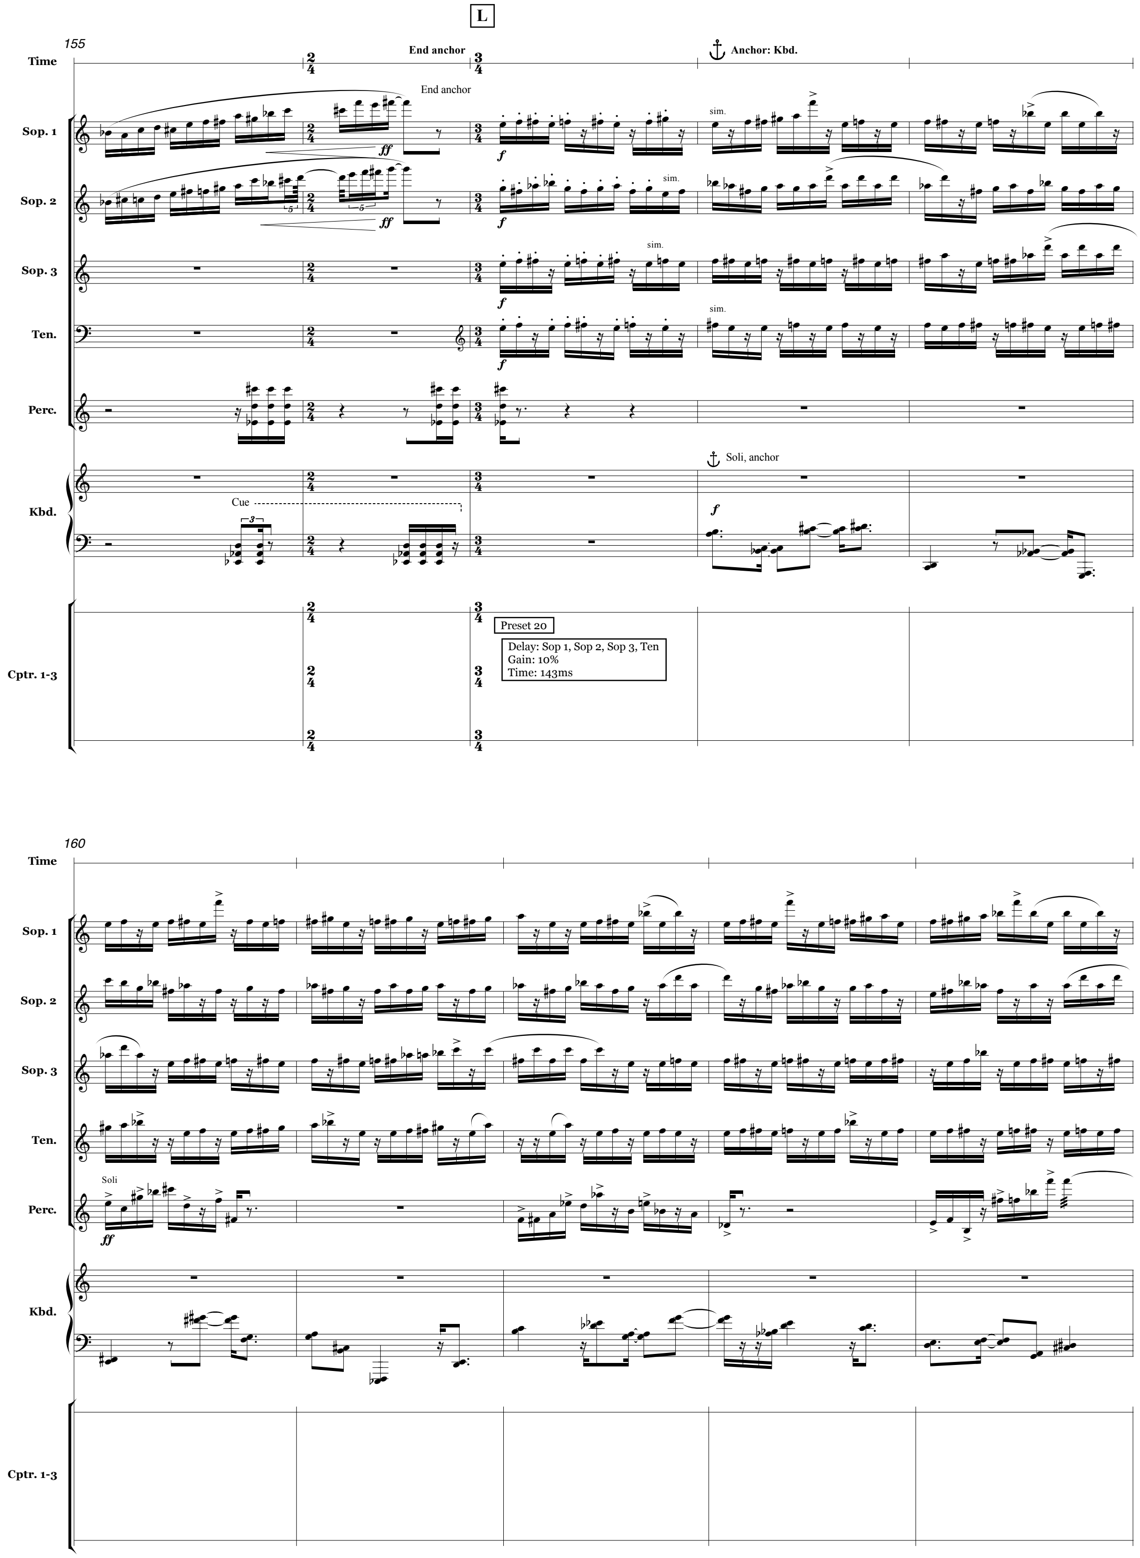
L
155
Time
End anchor
Anchor: Kbd.
End anchor
Sop. 1
sim.
Sop. 2
sim.
Sop. 3
sim.
Ten.
sim.
Perc.
Soli, anchor
Kbd.
Cue
Preset 20
Delay: Sop 1, Sop 2, Sop 3, Ten
Gain: 10%
Time: 143ms
Cptr. 1-3
160
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Soli
Perc.
Kbd.
Cptr. 1-3


165
Time
End anchor
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
End anchor
Kbd.
Cptr. 1-3
Reverb: Fade to large hall?
M
♩= 104
170
Anchor: Sop. 1
Time
Anchor
Sop. 1
p
mp
p
mp
Sop. 2
Sop. 3
Ten.
Medium mallets
Perc.
pp
p
Kbd.
Preset 21
Cptr. 1-3
Delay->Harmonizer: Perc., Kbd.
Delay time: 3396
Delay feedback: 0.1
Cents: 1200


176
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
mp p mf p mp pp p pp
181
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
p mf mp p mp pp p mp p mp pp p pp mp mp


186
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
p
mp
mp p
pp
8va
190
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3


194
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
Preset 22
Delays:
Gain: 0
Ramp: 10'
199
End anchor
End anchor

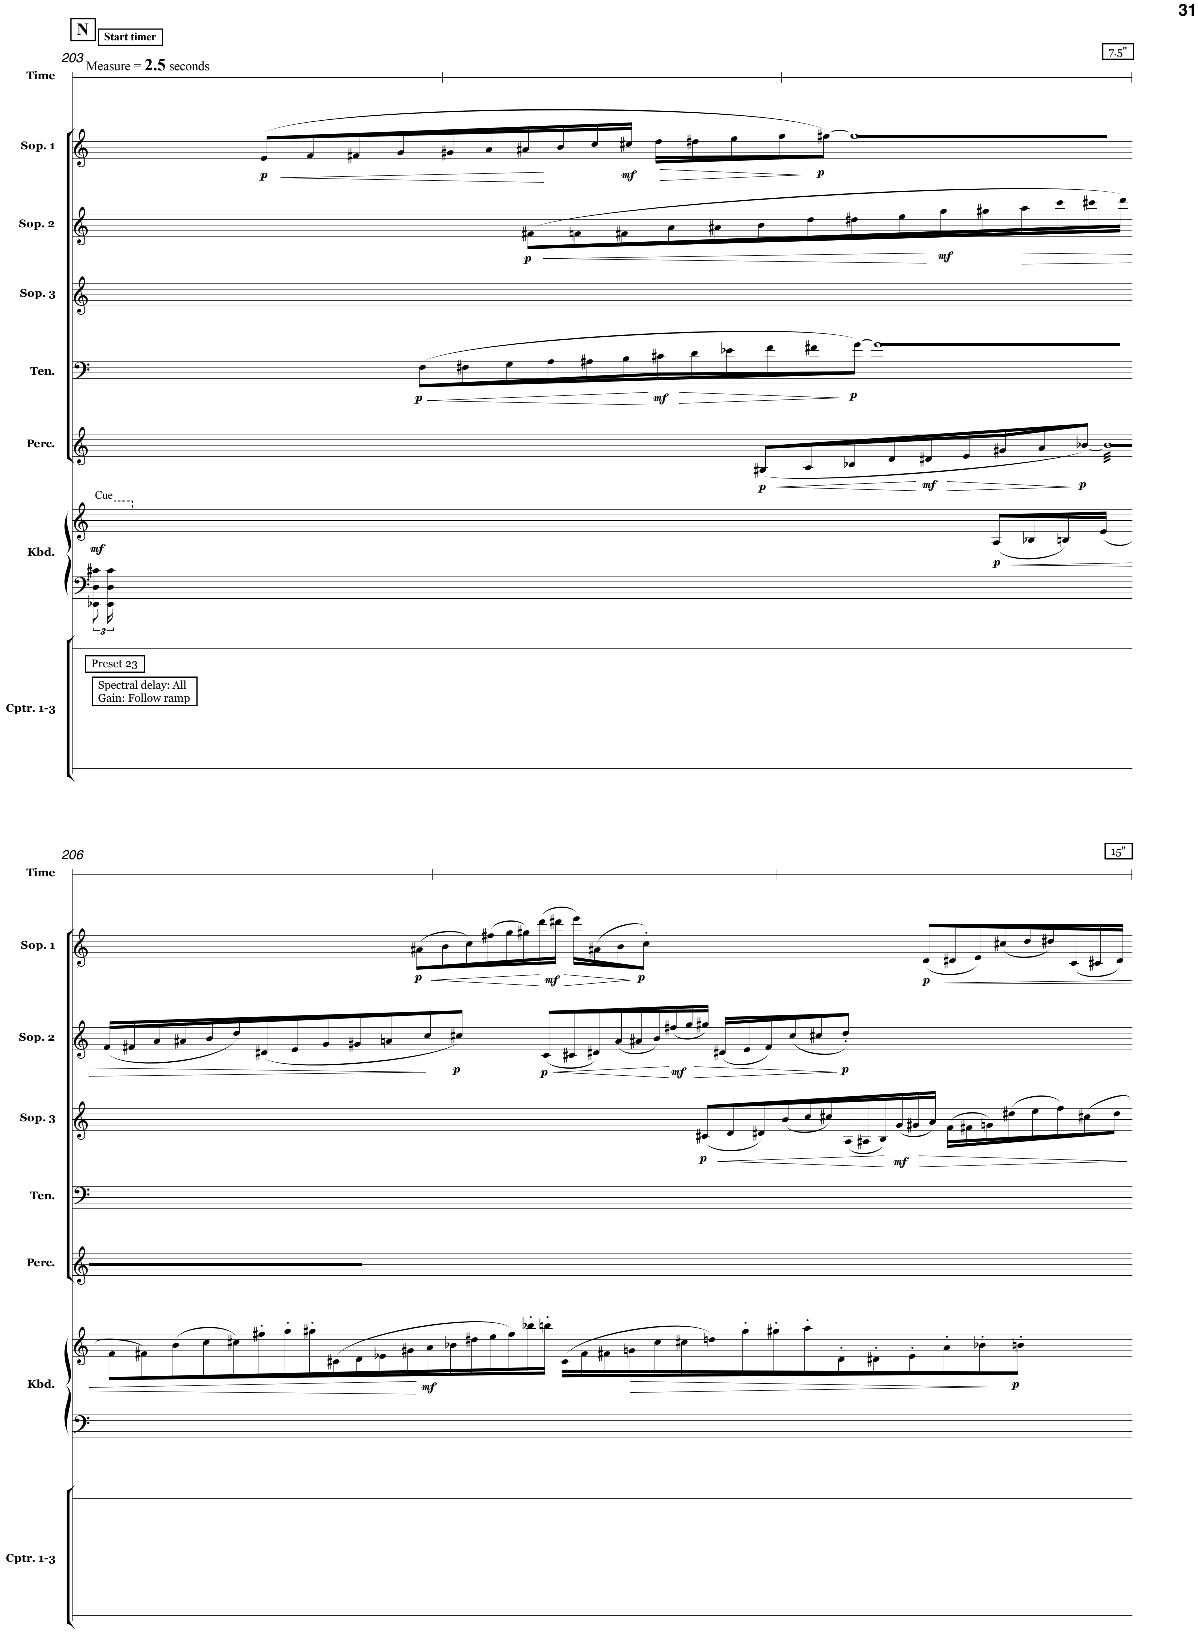
N
Start timer
203
Measure = 2.5 seconds
7.5"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Cue
Kbd.
Preset 23
Spectral delay: All
Gain: Follow ramp
Cptr. 1-3
206
15"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3

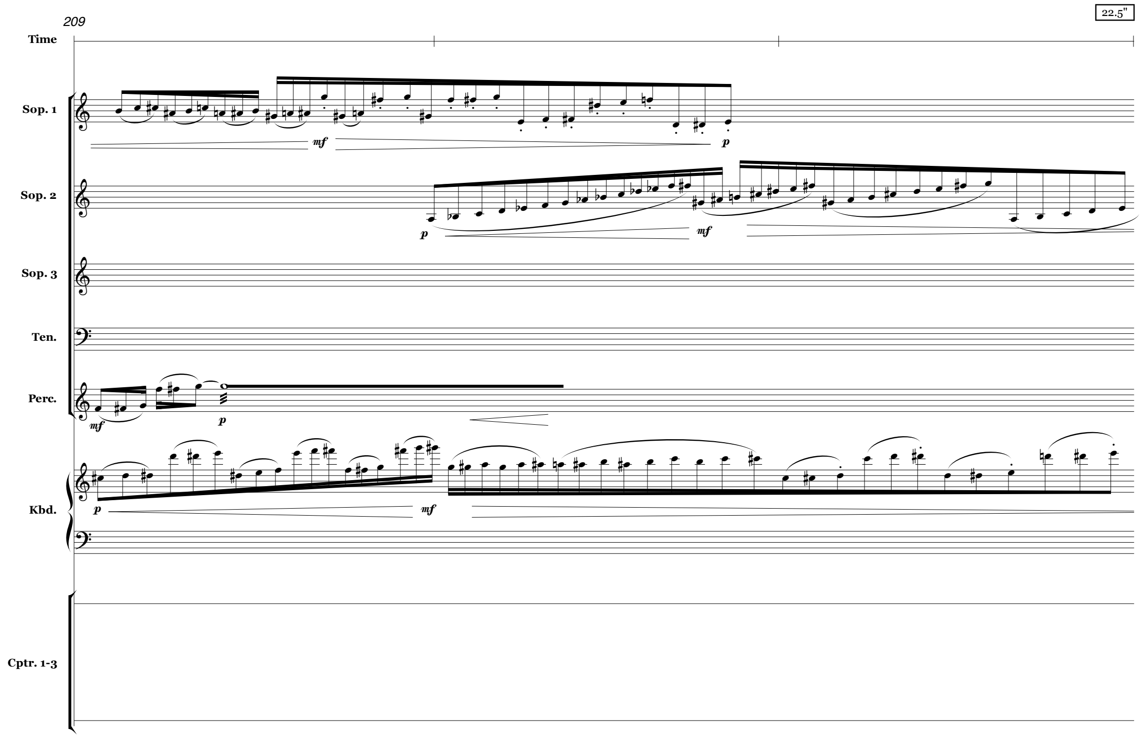
209
22.5"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
mf
p

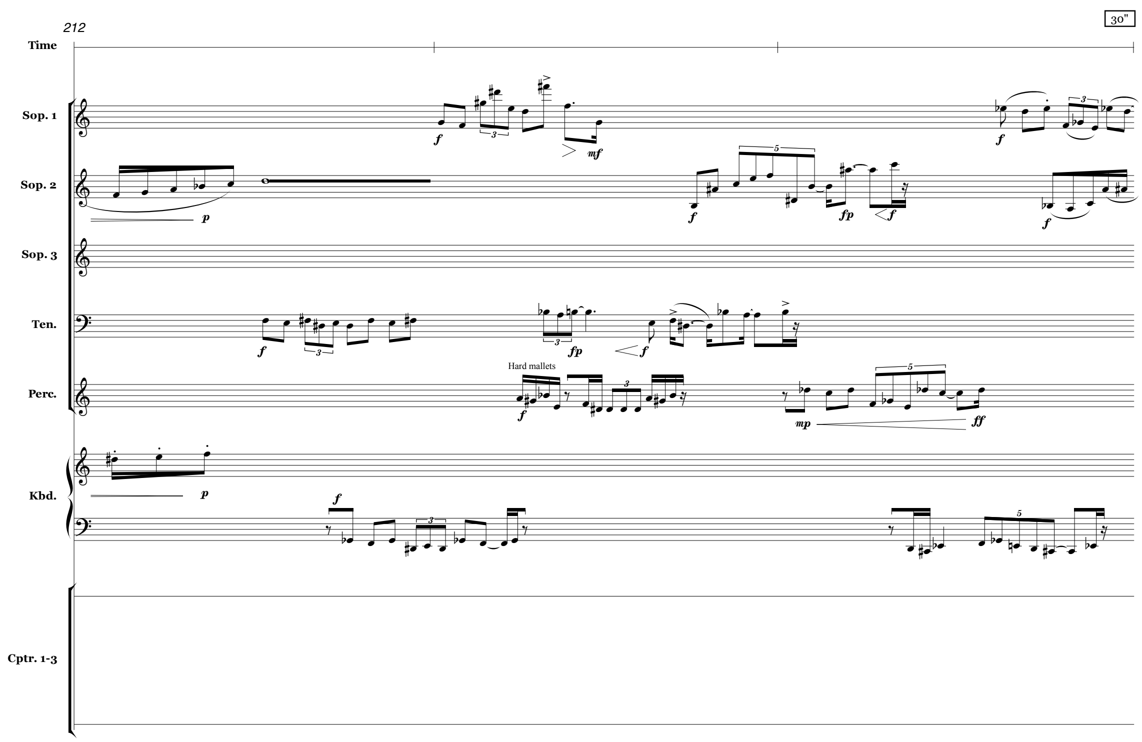
212
30"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Hard mallets
Kbd.
Cptr. 1-3
f
mf
p
fp
mp
ff

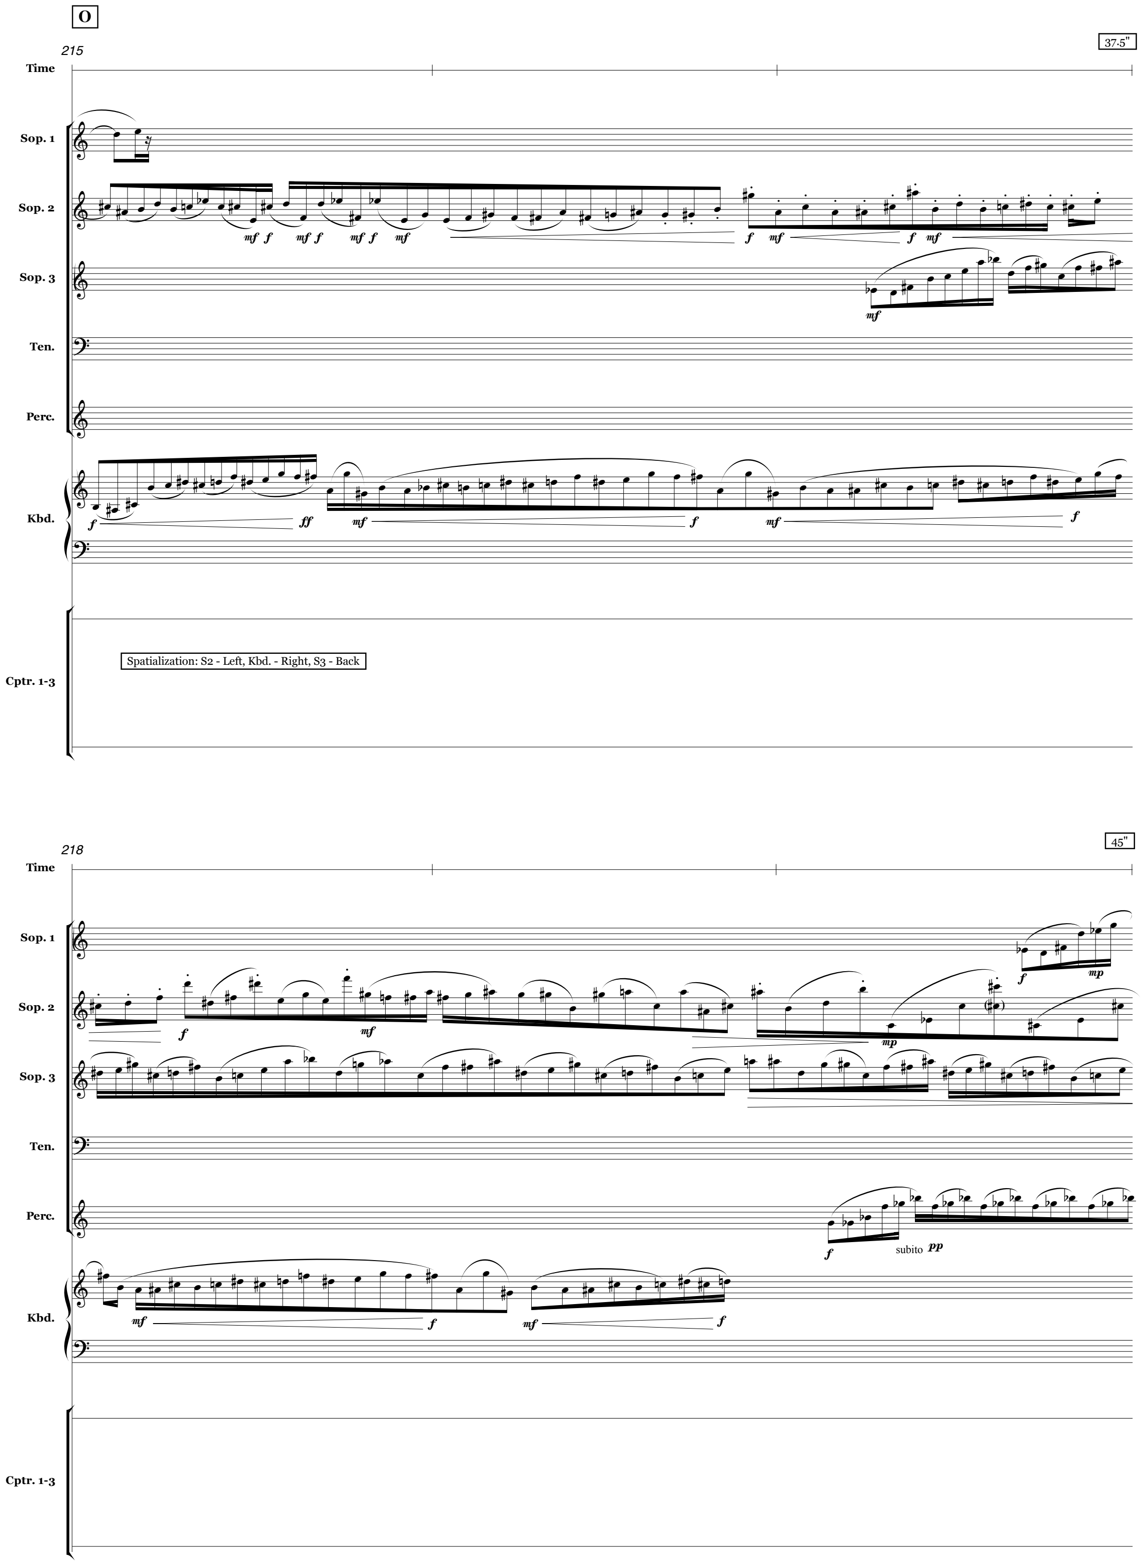
O
215
Time
37.5"
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
Spatialization: S2 - Left, Kbd. - Right, S3 - Back
218
45"
subito pp

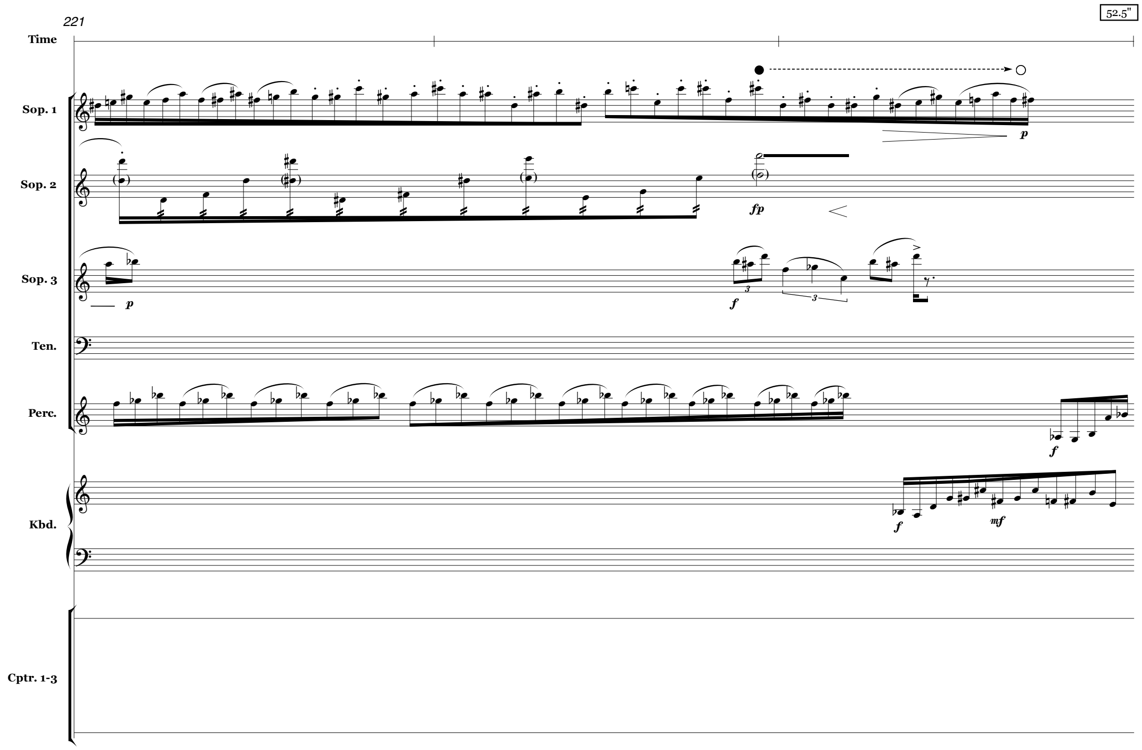
221
52.5"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3

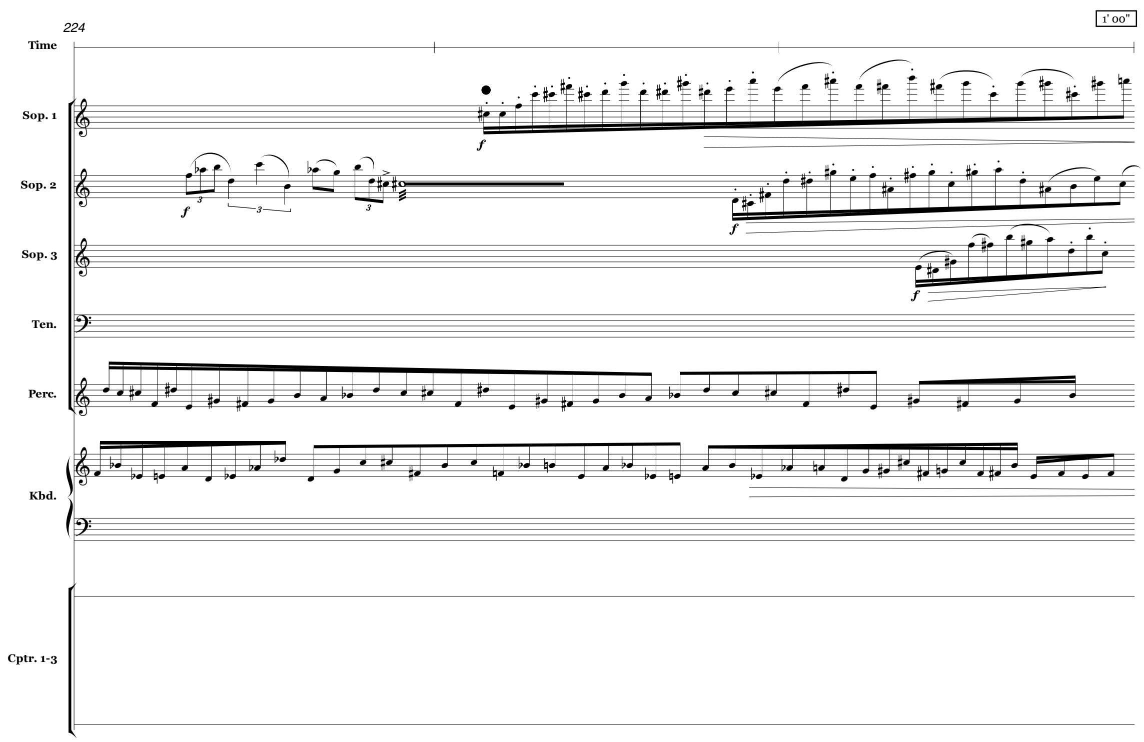
224
1' 00"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3

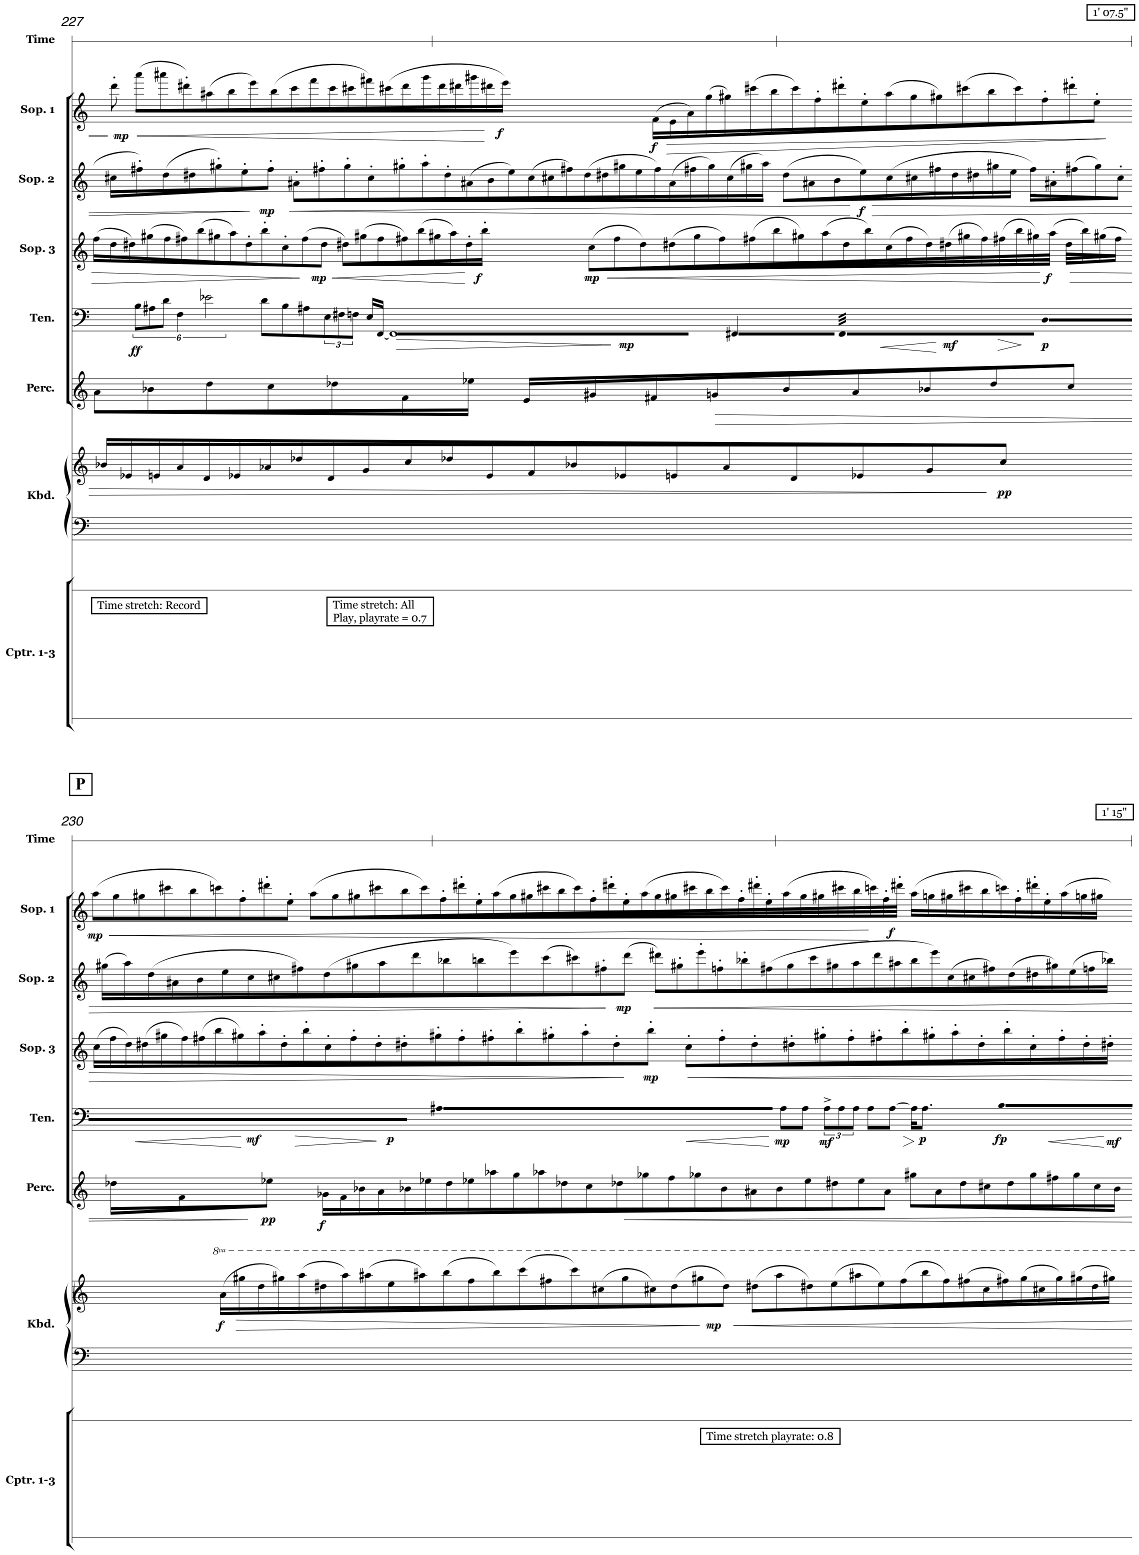
227
1' 07.5"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Time stretch: Record
Time stretch: All
Play, playrate = 0.7
Cptr. 1-3
P
230
1' 15"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Time stretch playrate: 0.8
Cptr. 1-3

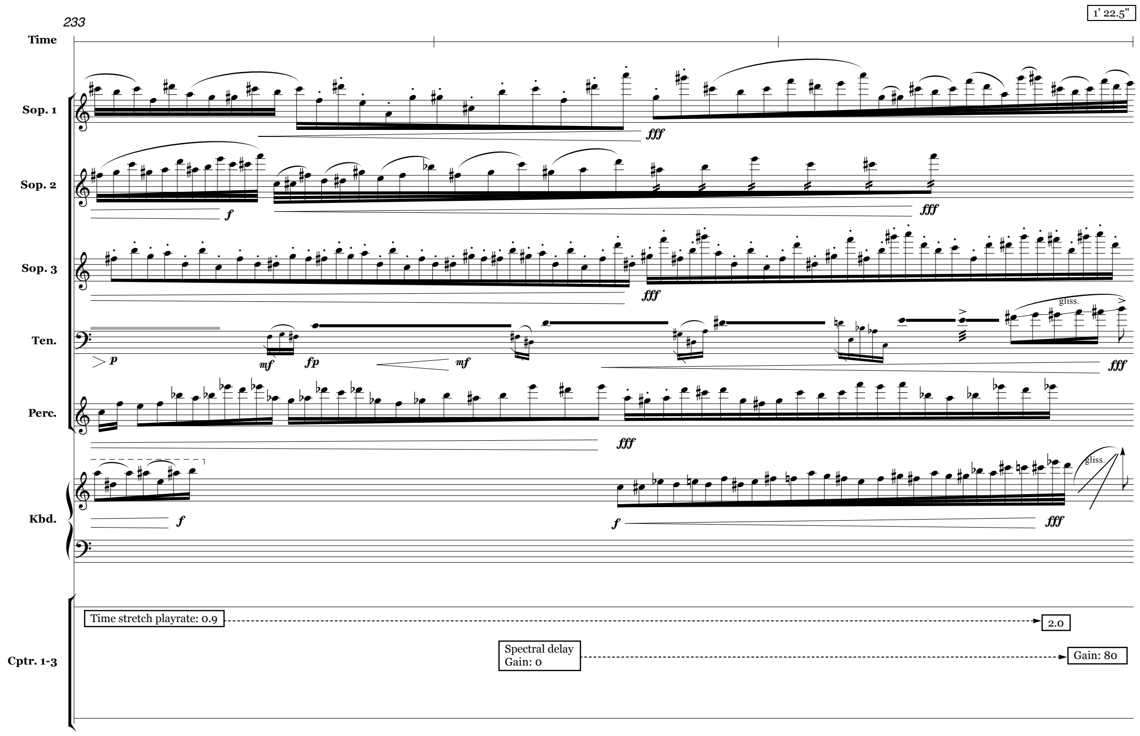
233
1' 22.5"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
fff
f
fff
fff
p
mf
fp
mf
fff
fff
f
f
fff
gliss.
gliss.
Time stretch playrate: 0.9
2.0
Spectral delay
Gain: 0
Gain: 80

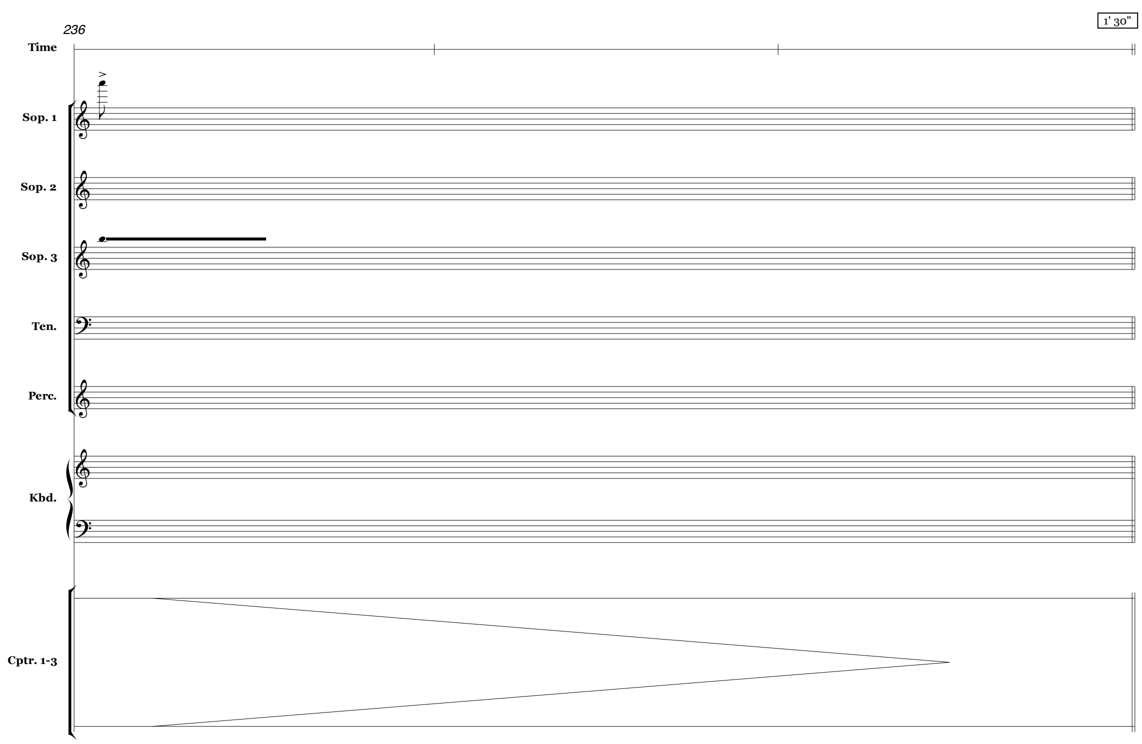
236
1' 30"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3

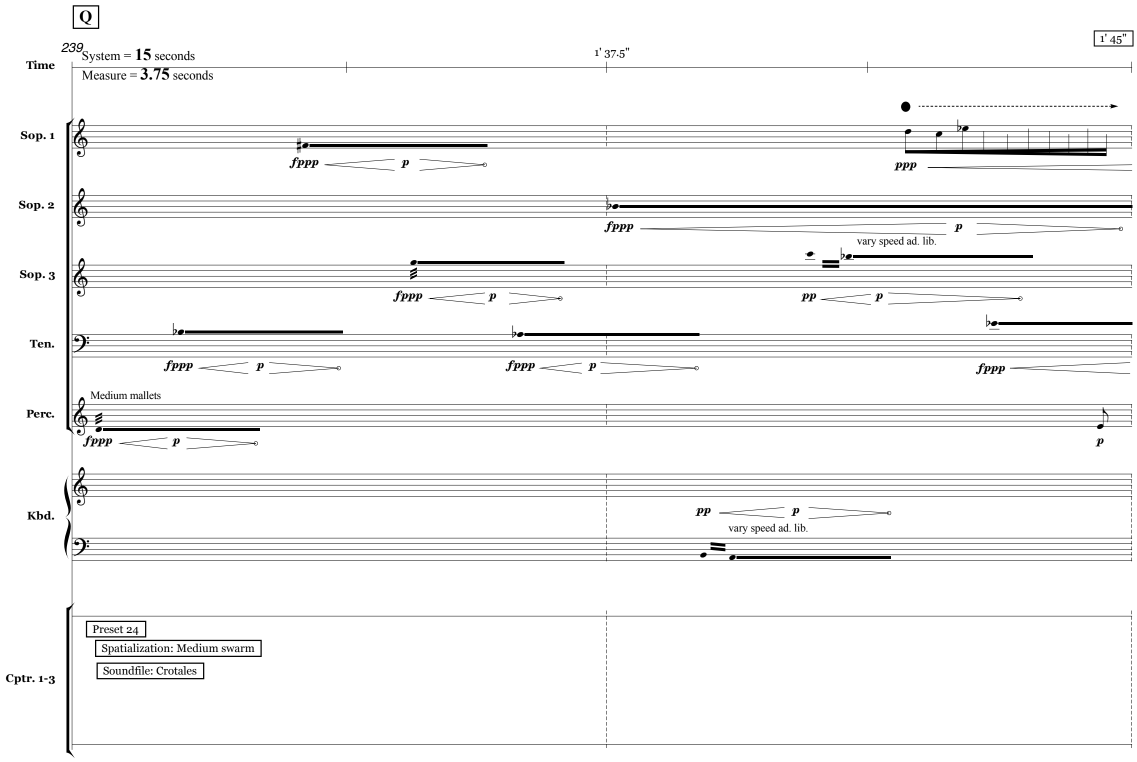
Q
239
System = 15 seconds
Measure = 3.75 seconds
1' 37.5"
1' 45"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
fppp
p
ppp
vary speed ad. lib.
pp
Medium mallets
Preset 24
Spatialization: Medium swarm
Soundfile: Crotales

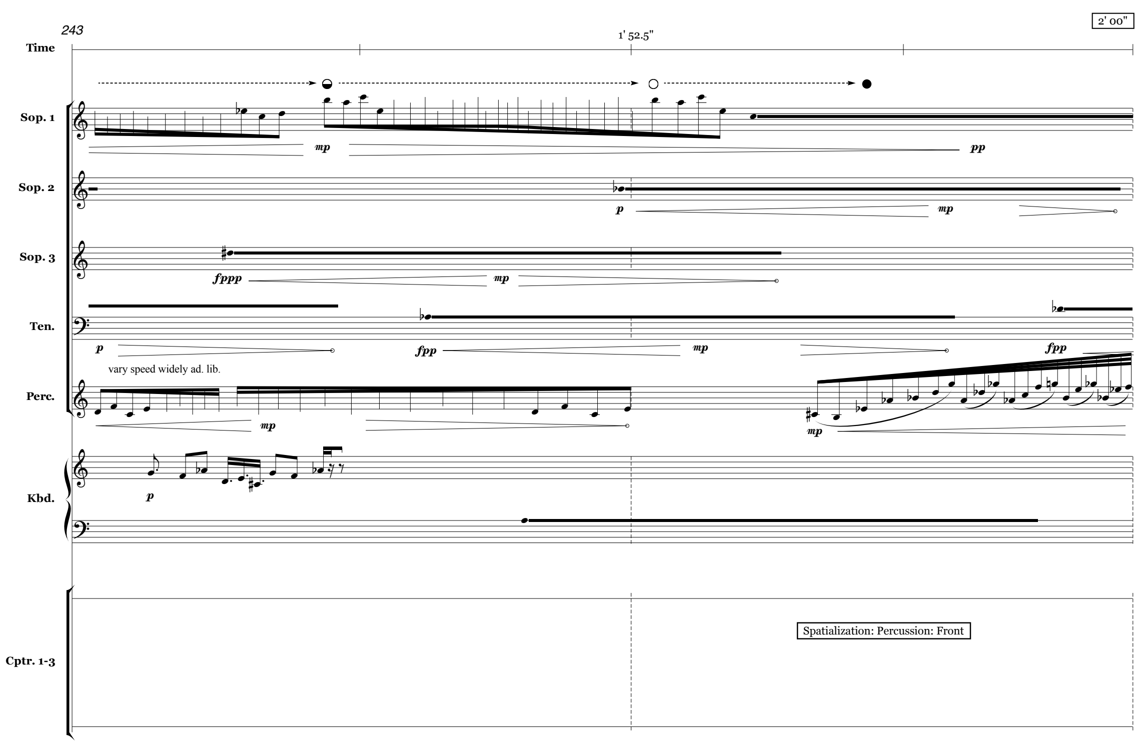
243
1' 52.5"
2' 00"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
mp
pp
p
fppp
fpp
vary speed widely ad. lib.
Spatialization: Percussion: Front

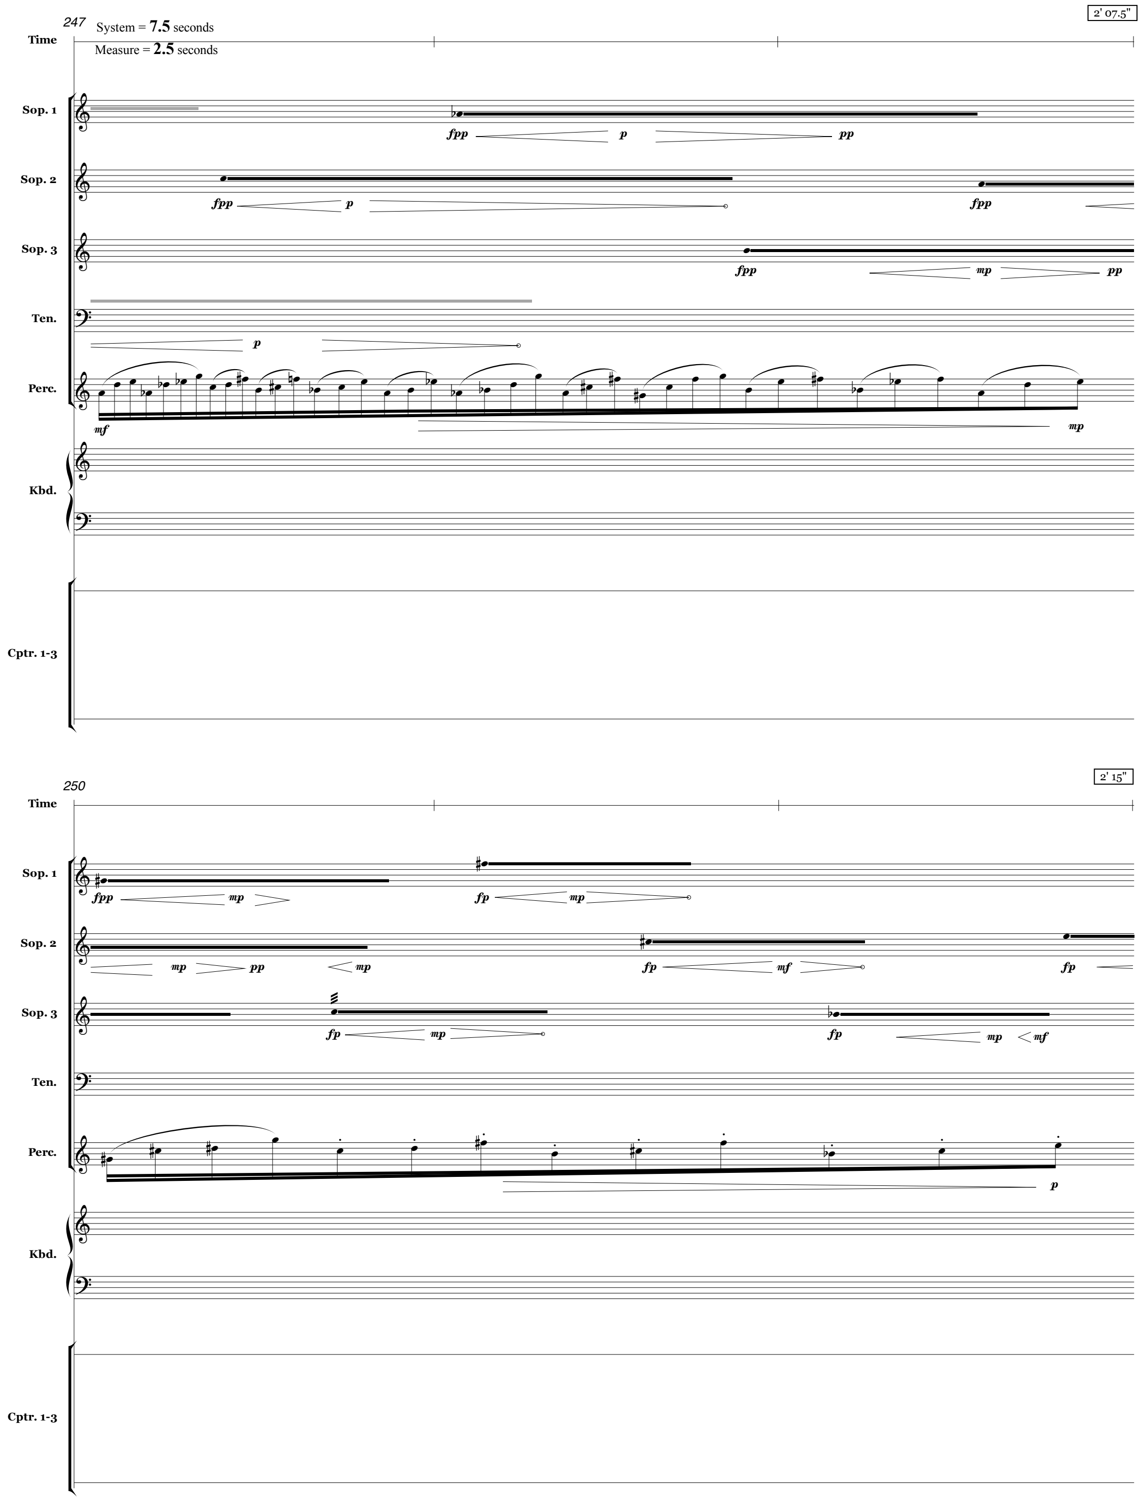
247
System = 7.5 seconds
Measure = 2.5 seconds
2' 07.5"
Time
Sop. 1
fpp
p
pp
Sop. 2
fpp
p
fpp
Sop. 3
fpp
mp
pp
Ten.
p
Perc.
mf
mp
Kbd.
Cptr. 1-3
250
2' 15"
Time
Sop. 1
fpp
mp
fp
mp
Sop. 2
mp
pp
mp
fp
mf
fp
Sop. 3
fp
mp
fp
mp
mf
Ten.
Perc.
p
Kbd.
Cptr. 1-3

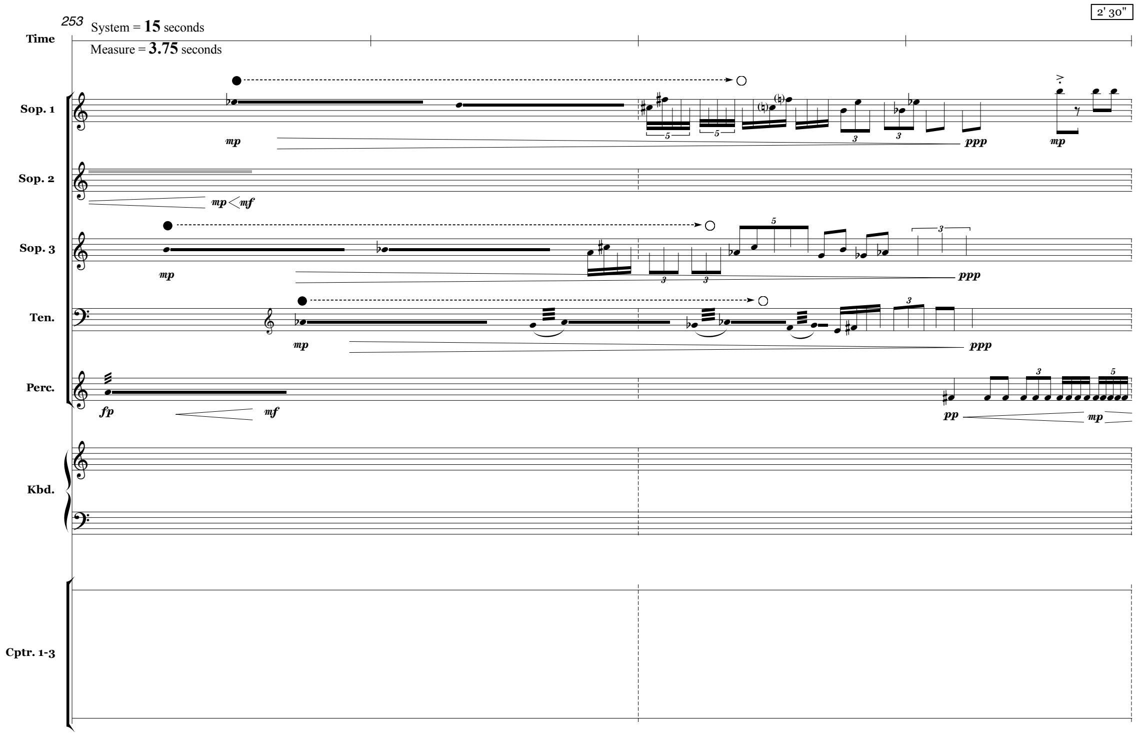
253
2' 30"
System = 15 seconds
Measure = 3.75 seconds
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3

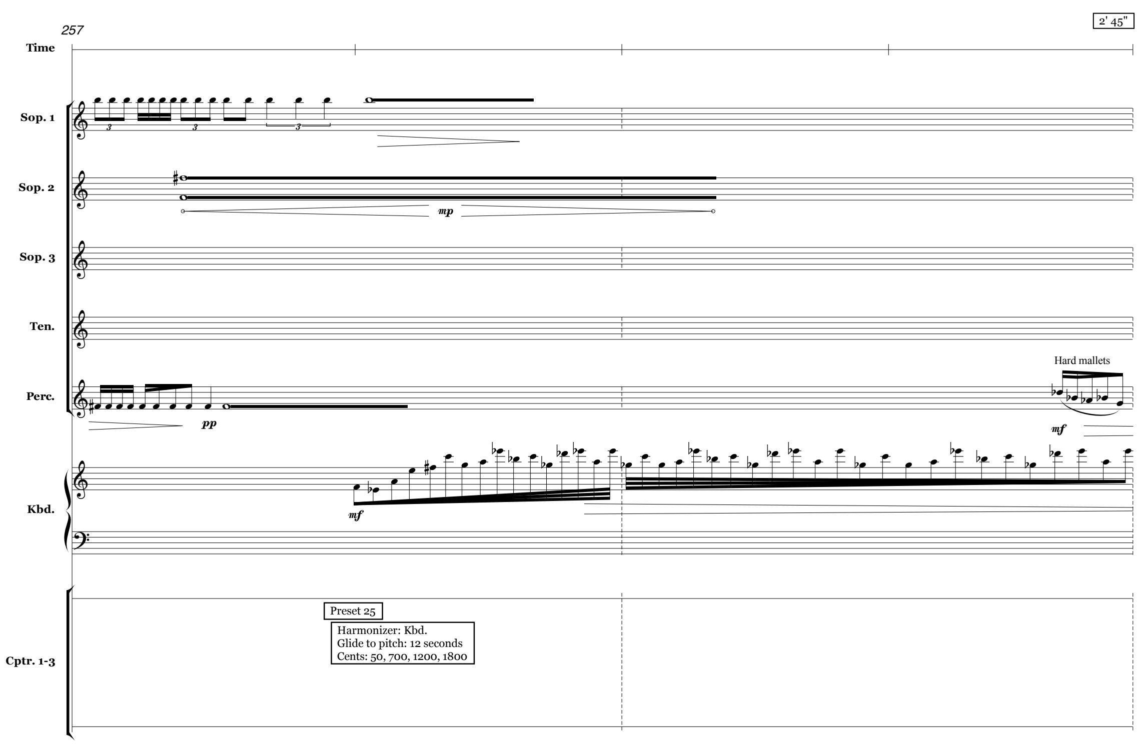
257
2' 45"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Hard mallets
Kbd.
Cptr. 1-3
Preset 25
Harmonizer: Kbd.
Glide to pitch: 12 seconds
Cents: 50, 700, 1200, 1800

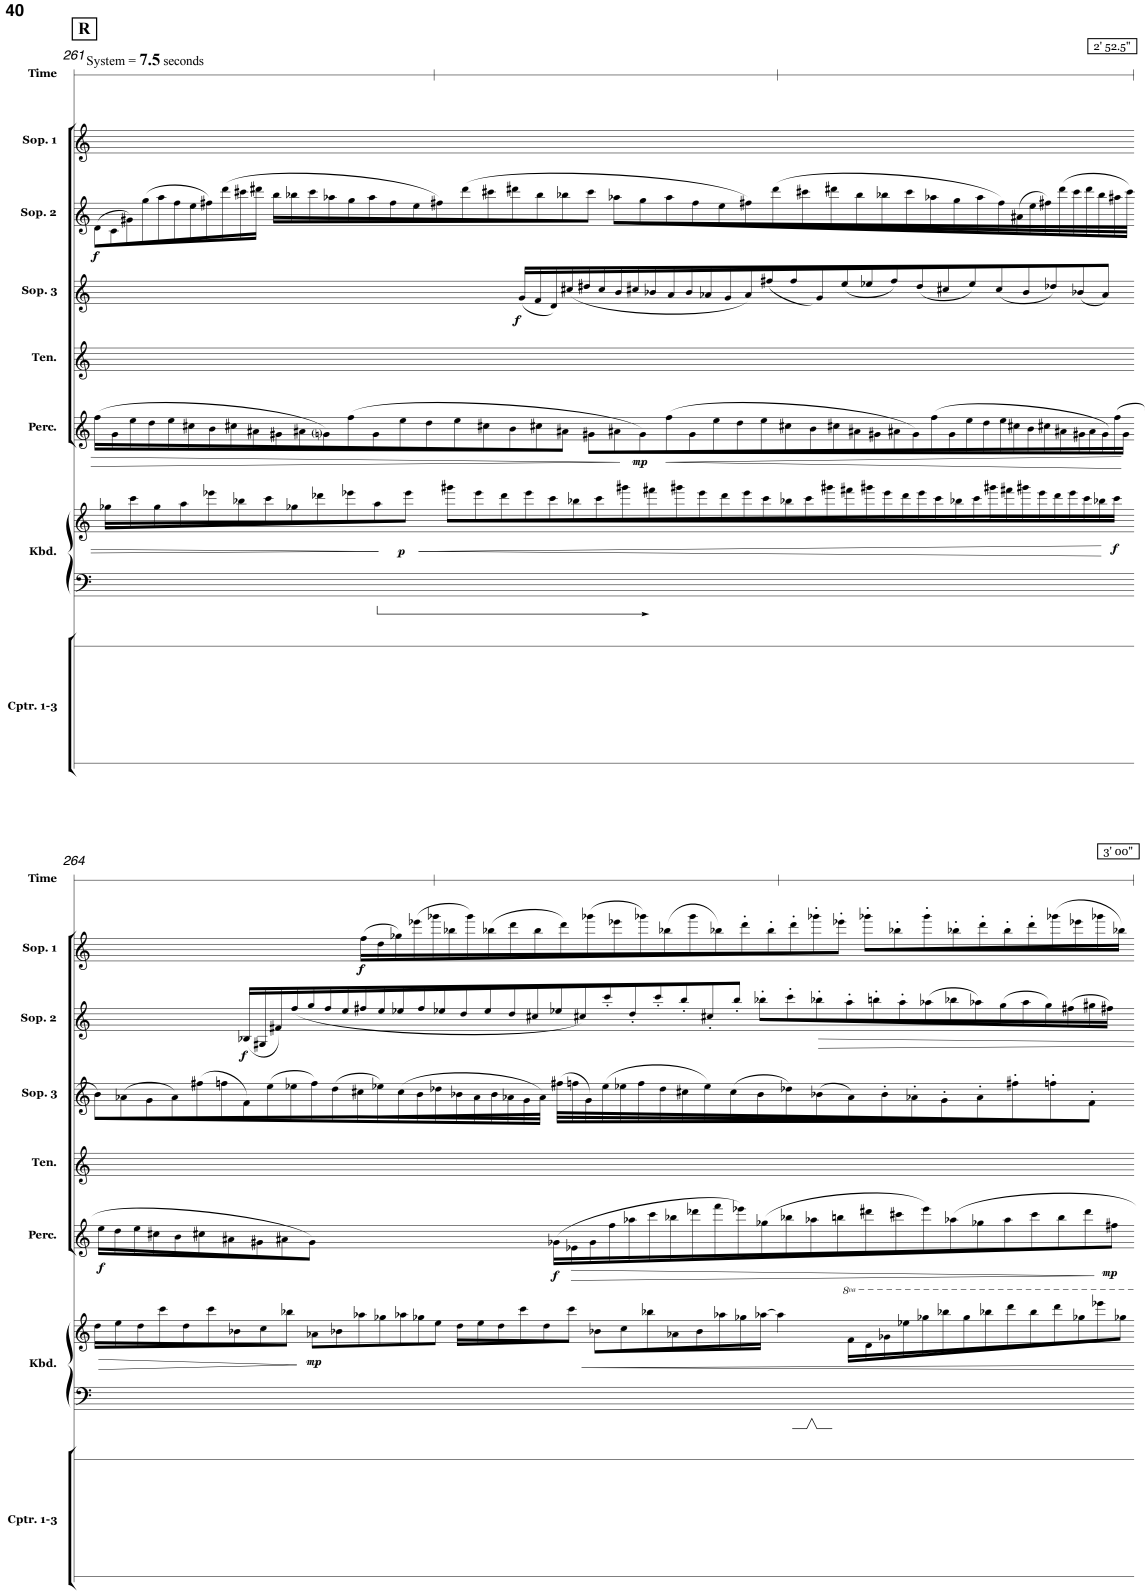
R
261 System = 7.5 seconds
2' 52.5"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
f
f
mp
p
f
264
3' 00"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
f
f
f
f
mp
mp
8va

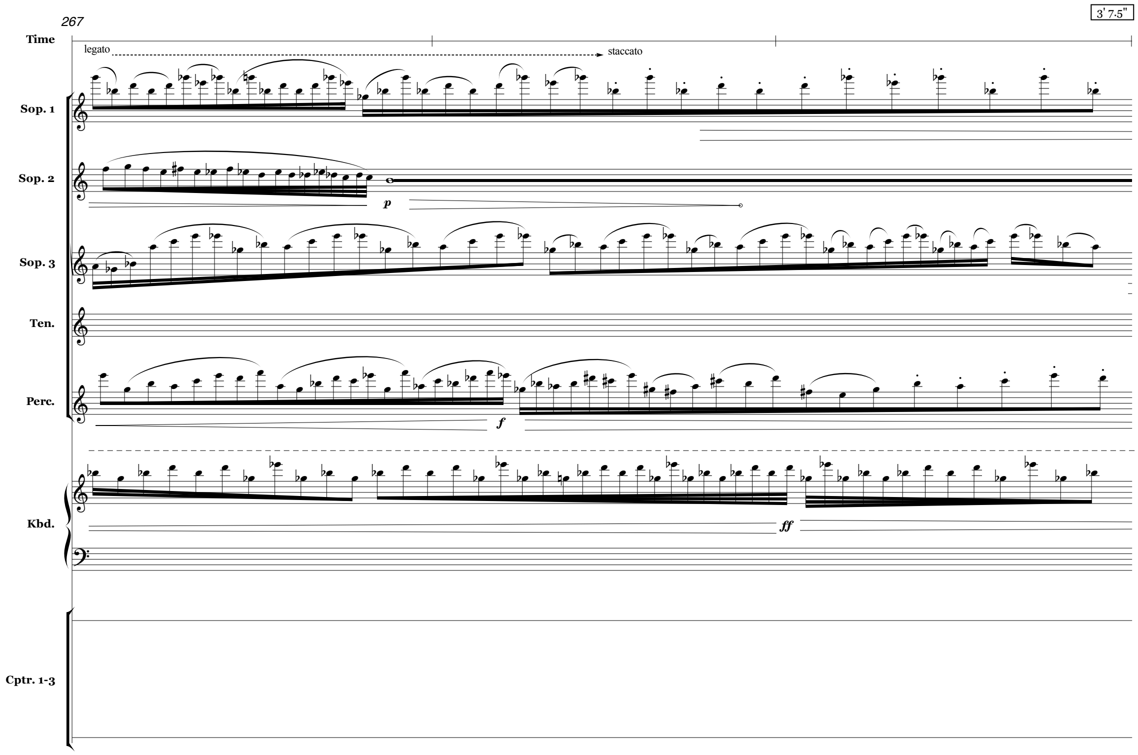
267
3' 7.5"
Time
legato
staccato
Sop. 1
Sop. 2
p
Sop. 3
Ten.
Perc.
f
Kbd.
ff
Cptr. 1-3

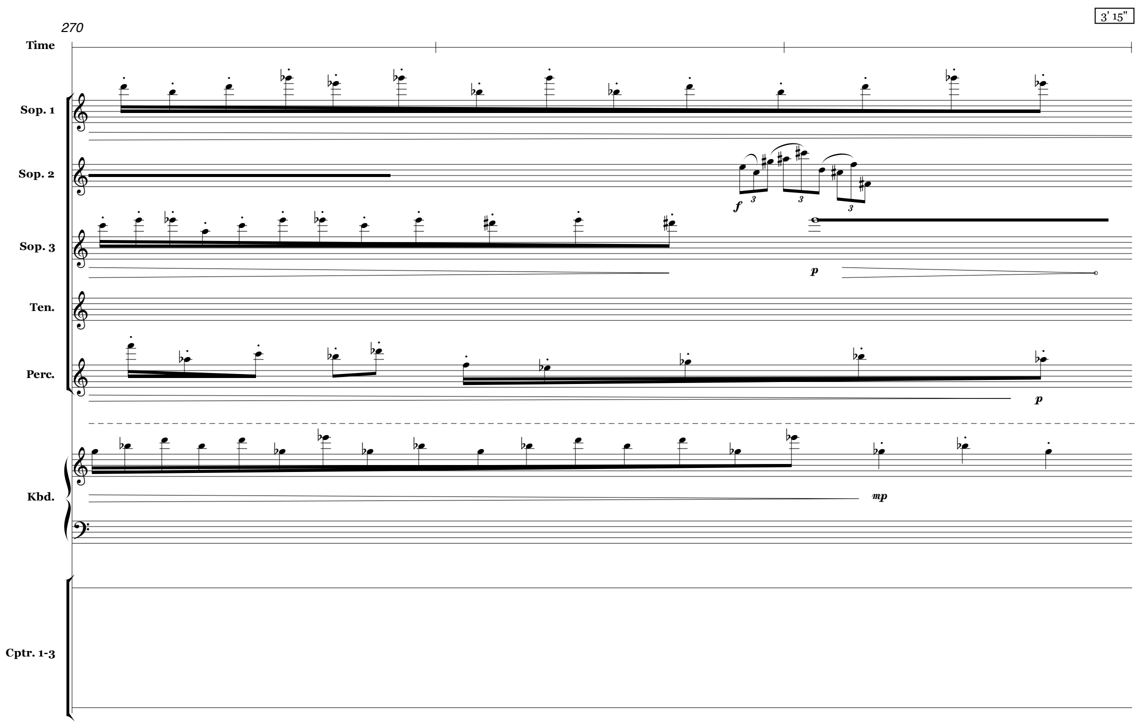
270
3' 15"
Time
Sop. 1
Sop. 2
f
3
3
3
Sop. 3
p
Ten.
Perc.
p
Kbd.
mp
Cptr. 1-3

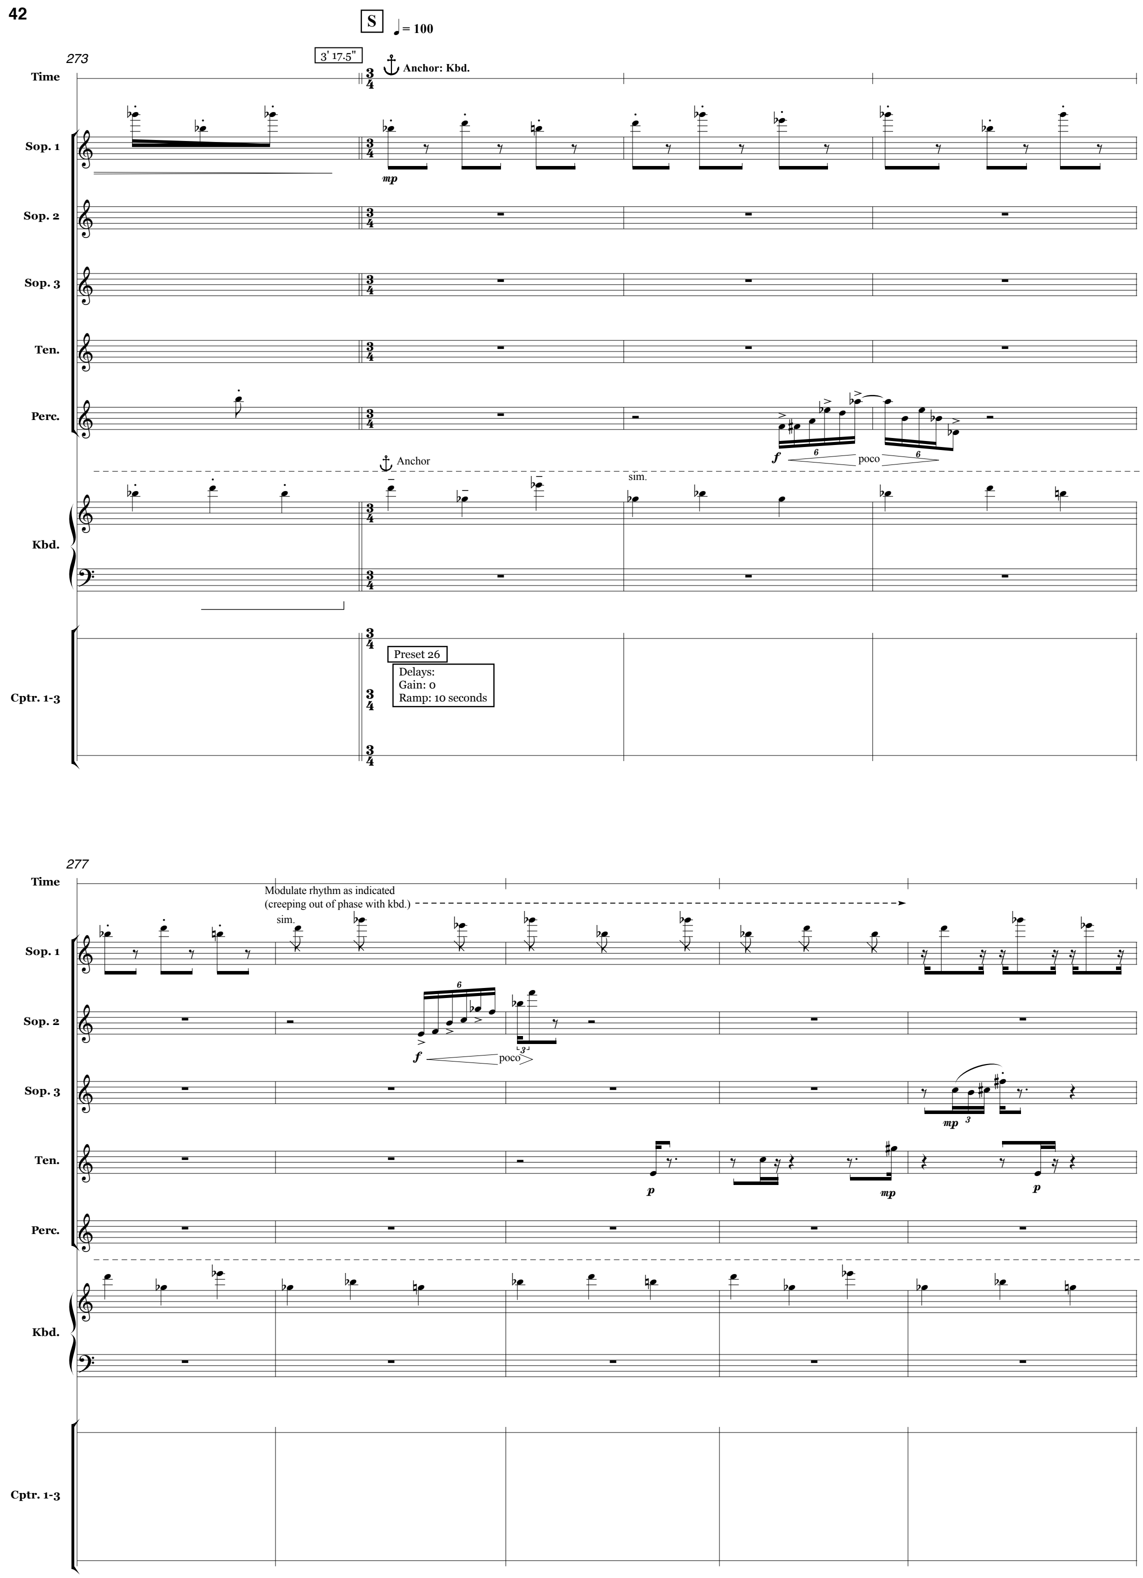
42
S
♩= 100
273
3' 17.5"
Anchor: Kbd.
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Anchor
sim.
poco
Kbd.
Preset 26
Delays:
Gain: 0
Ramp: 10 seconds
Cptr. 1-3
277
Modulate rhythm as indicated
(creeping out of phase with kbd.)


282
Time
modulate beat position
sim.
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
mp
mf
fp
ppp
p
287
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
p
ppp
mf

44

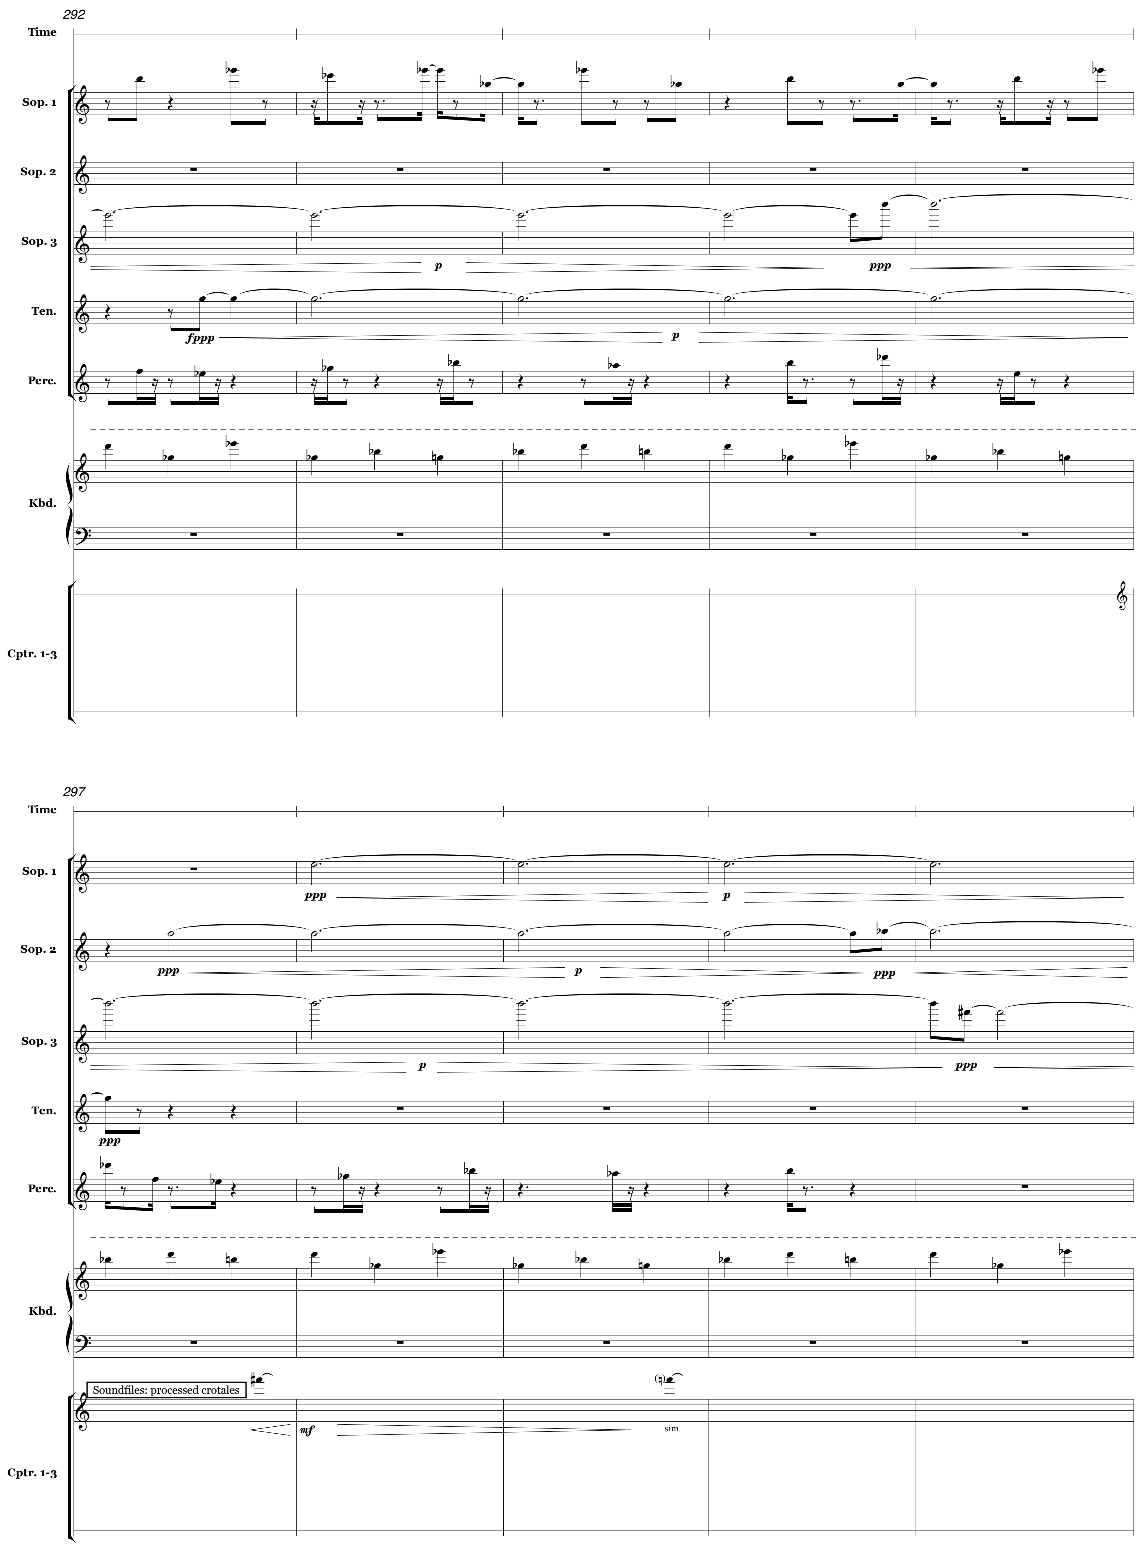
292
Time
Sop. 1
Sop. 2
Sop. 3
p
ppp
Ten.
fppp
p
Perc.
Kbd.
Cptr. 1-3
297
Time
Sop. 1
ppp
p
Sop. 2
ppp
p
ppp
Sop. 3
p
ppp
Ten.
ppp
Perc.
Kbd.
Soundfiles: processed crotales
mf
sim.
Cptr. 1-3


T
302
Time
Sop. 1
ppp
p
Sop. 2
p
ppp
Sop. 3
p
ppp
Ten.
Perc.
mf
5
Kbd.
Cptr. 1-3
307
Time
Sop. 1
ppp
mp
3
Sop. 2
Sop. 3
ppp
Ten.
Cue
mf
p
3
Perc.
5
p
mf
Kbd.
Cptr. 1-3

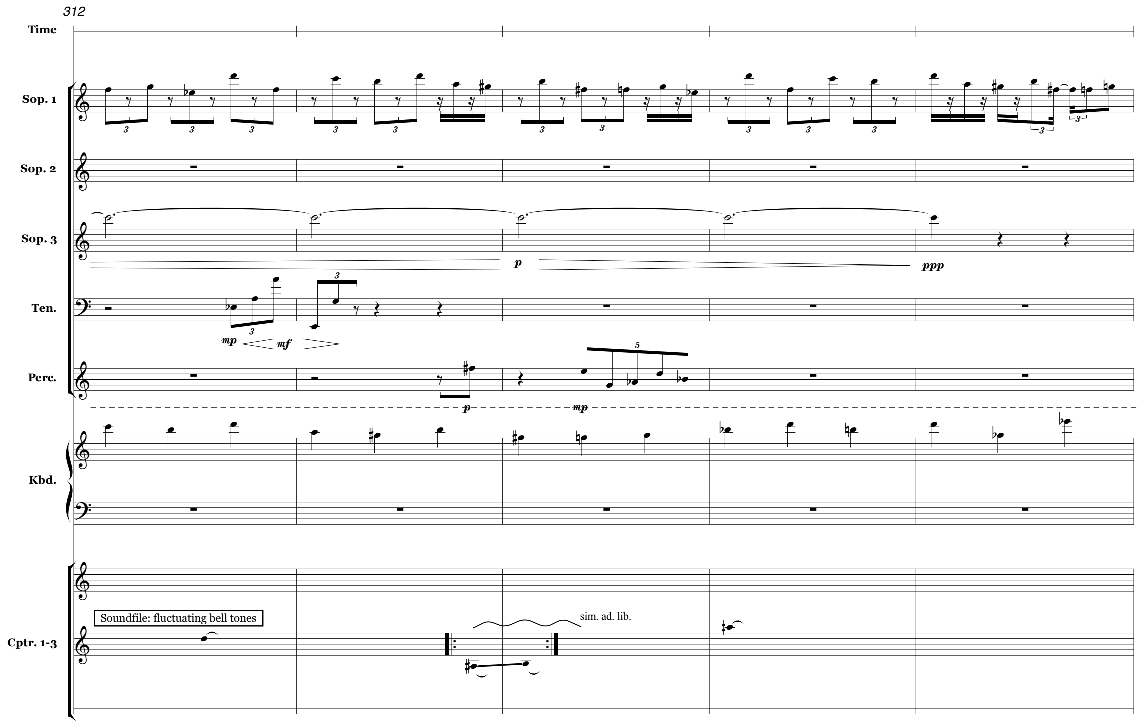
312
Time
Sop. 1
Sop. 2
Sop. 3
p
ppp
Ten.
mp
mf
Perc.
p
mp
Kbd.
Cptr. 1-3
Soundfile: fluctuating bell tones
sim. ad. lib.


317
Time
Sop. 1
mf
f
Sop. 2
Sop. 3
Ten.
mf
p
mf
Perc.
Kbd.
Cptr. 1-3

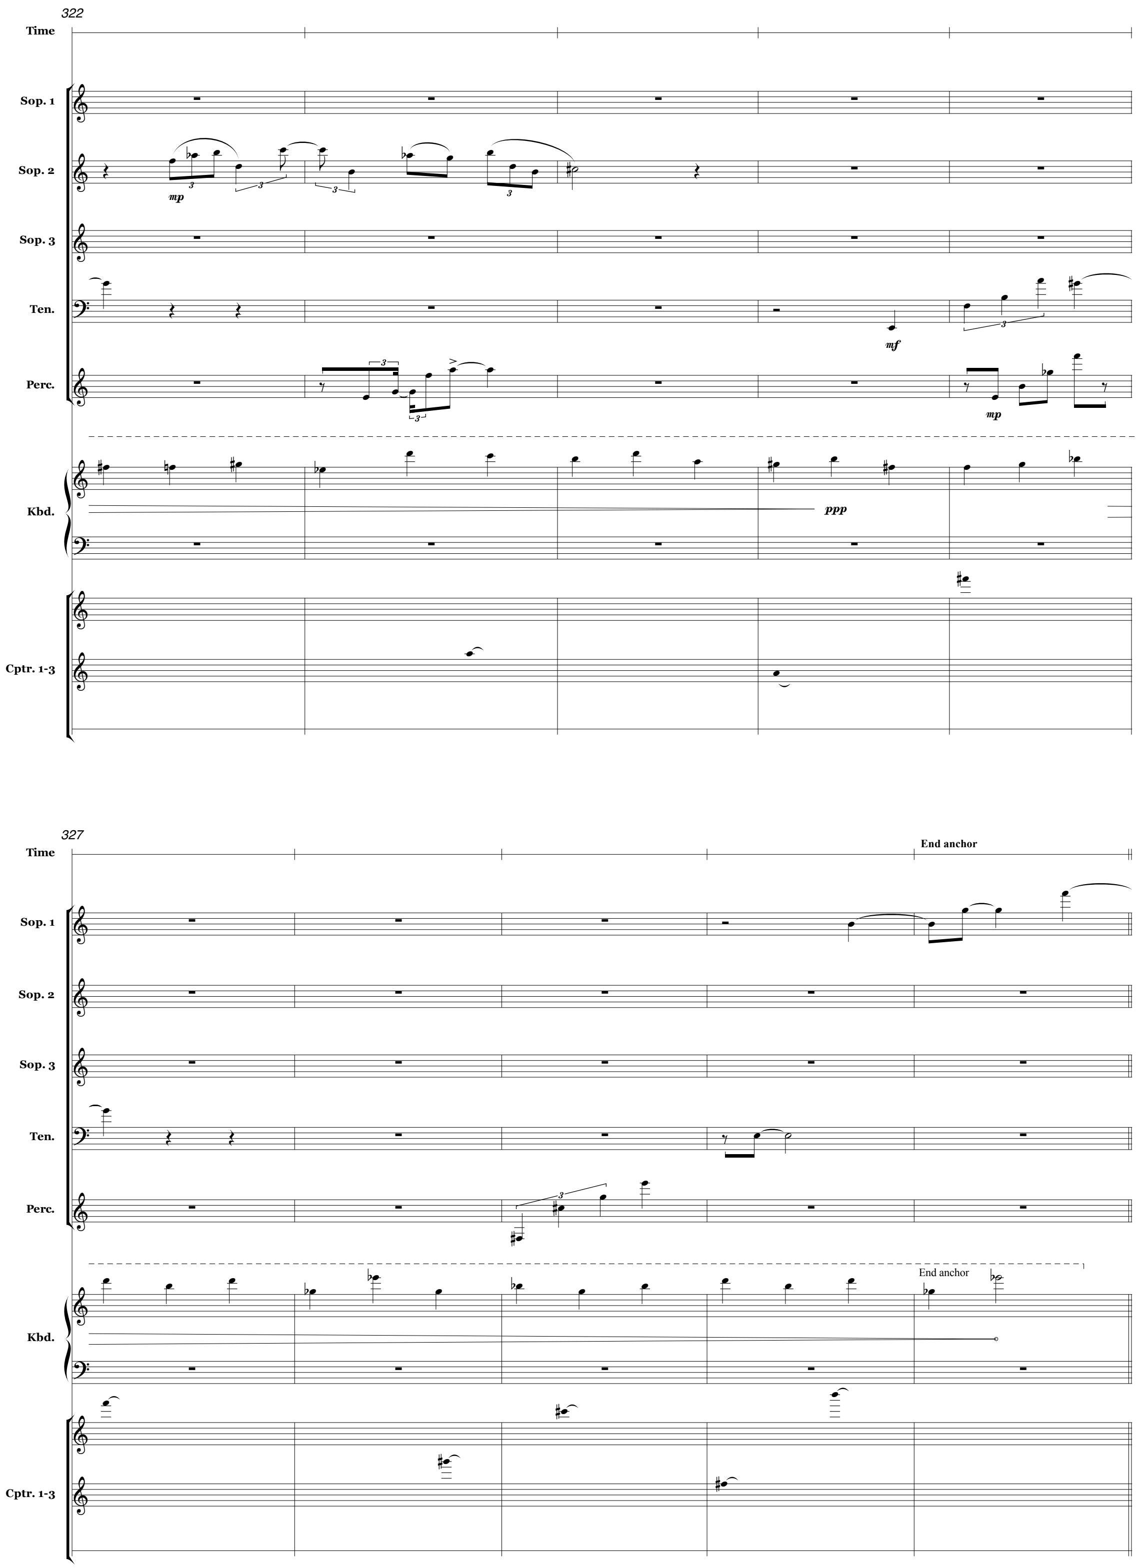
322
Time
Sop. 1
Sop. 2
mp
Sop. 3
Ten.
mf
Perc.
mp
Kbd.
ppp
Cptr. 1-3
327
Time
End anchor
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
End anchor
Kbd.
Cptr. 1-3

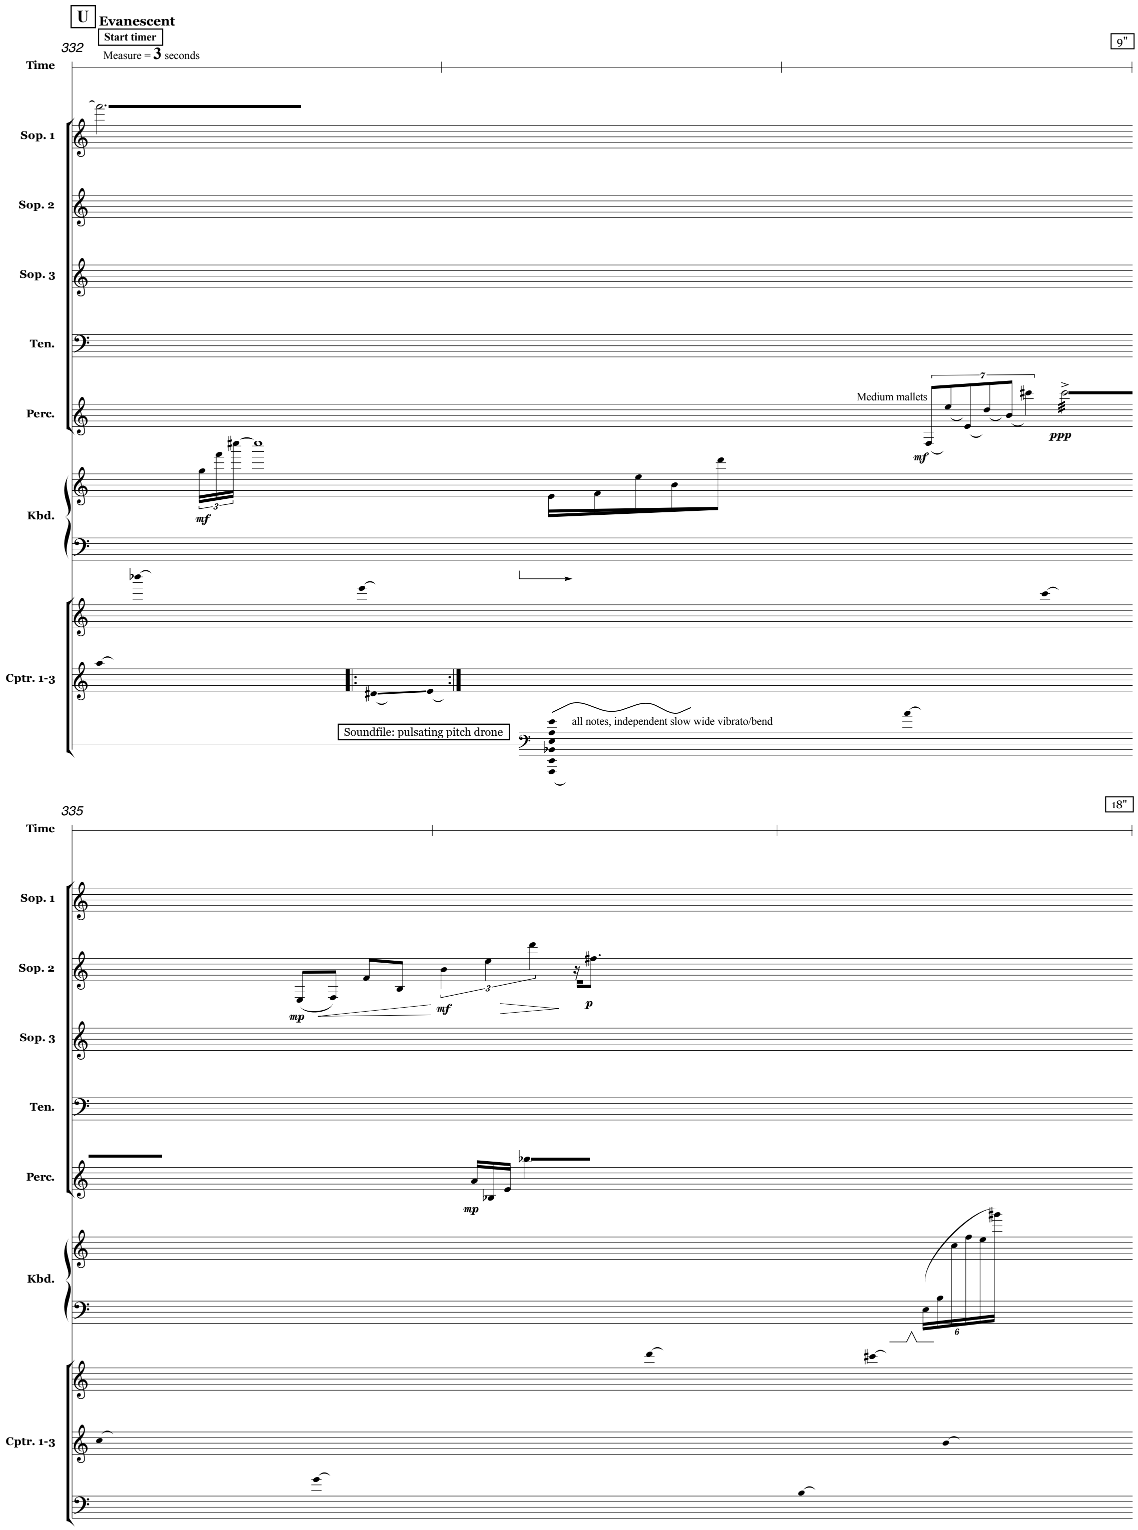
U
Evanescent
Start timer
332
Measure = 3 seconds
9"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Medium mallets
7
mf
ppp
Kbd.
3
mf
Cptr. 1-3
Soundfile: pulsating pitch drone
all notes, independent slow wide vibrato/bend
335
18"
Time
Sop. 1
Sop. 2
mp
mf
3
p
Sop. 3
Ten.
Perc.
mp
Kbd.
6
Cptr. 1-3

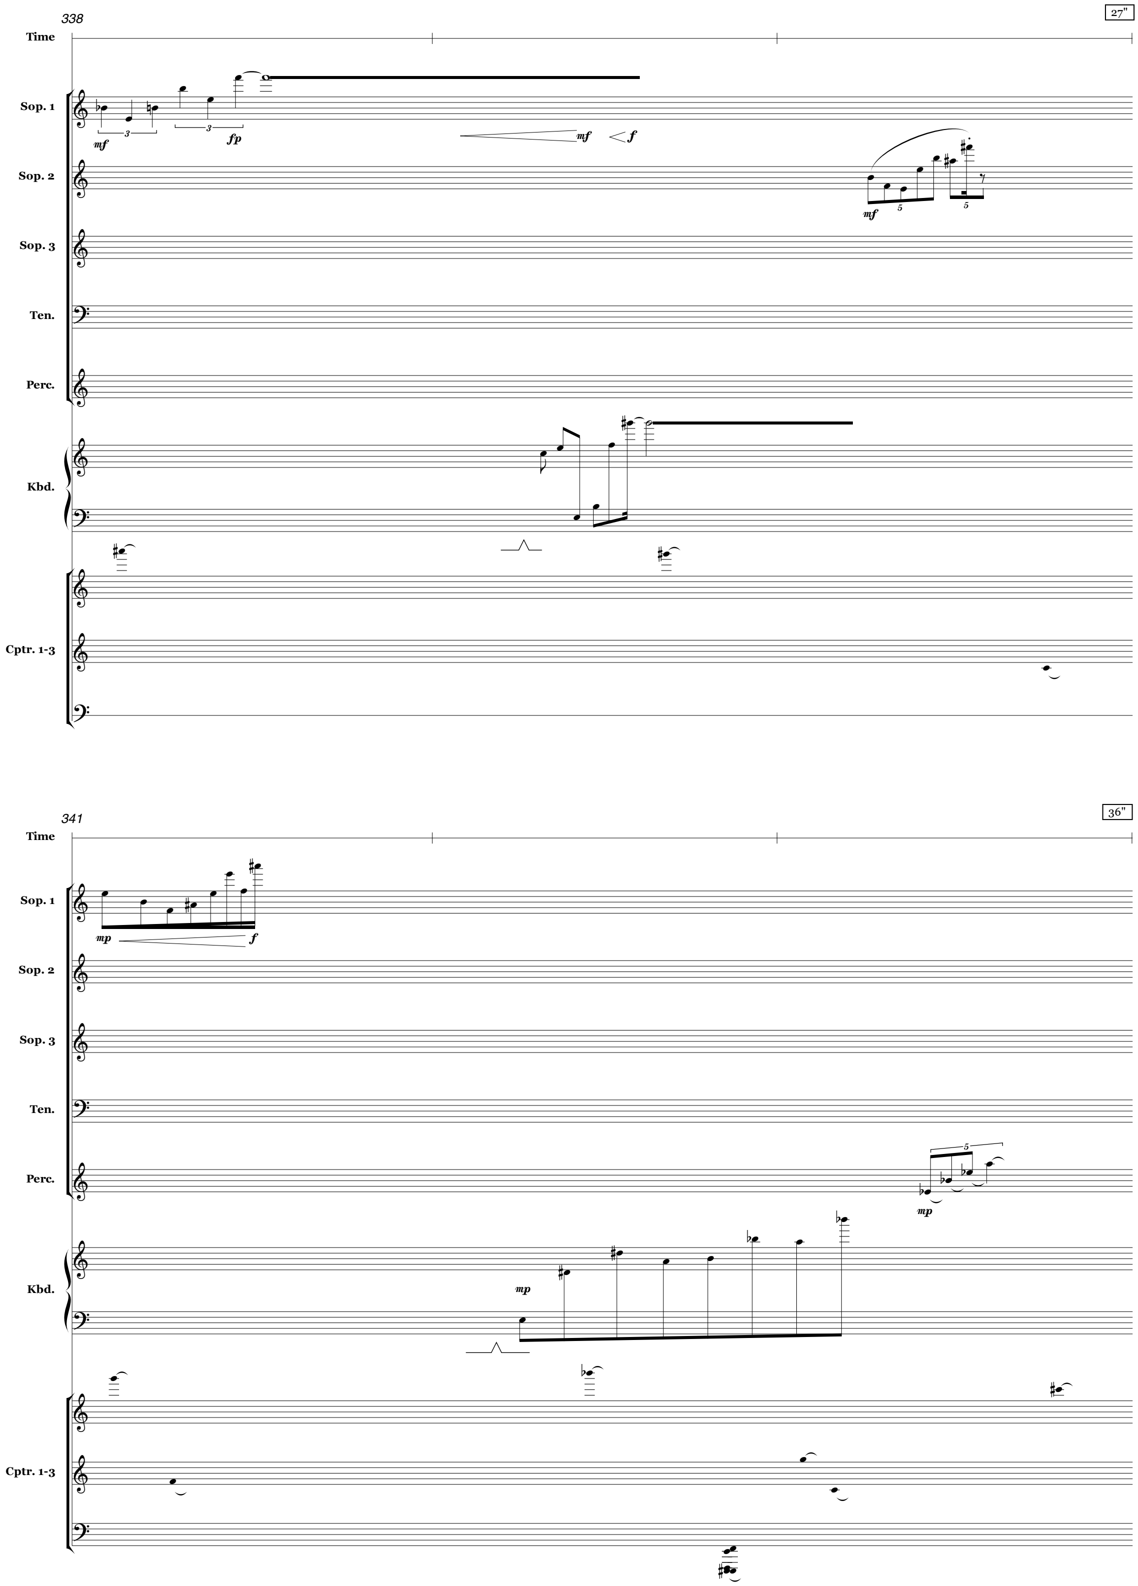
338
27"
Time
Sop. 1
mf
3
3
fp
mf
f
Sop. 2
mf
5
5
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
341
36"
Time
Sop. 1
mp
f
Sop. 2
Sop. 3
Ten.
Perc.
5
mp
Kbd.
mp
Cptr. 1-3

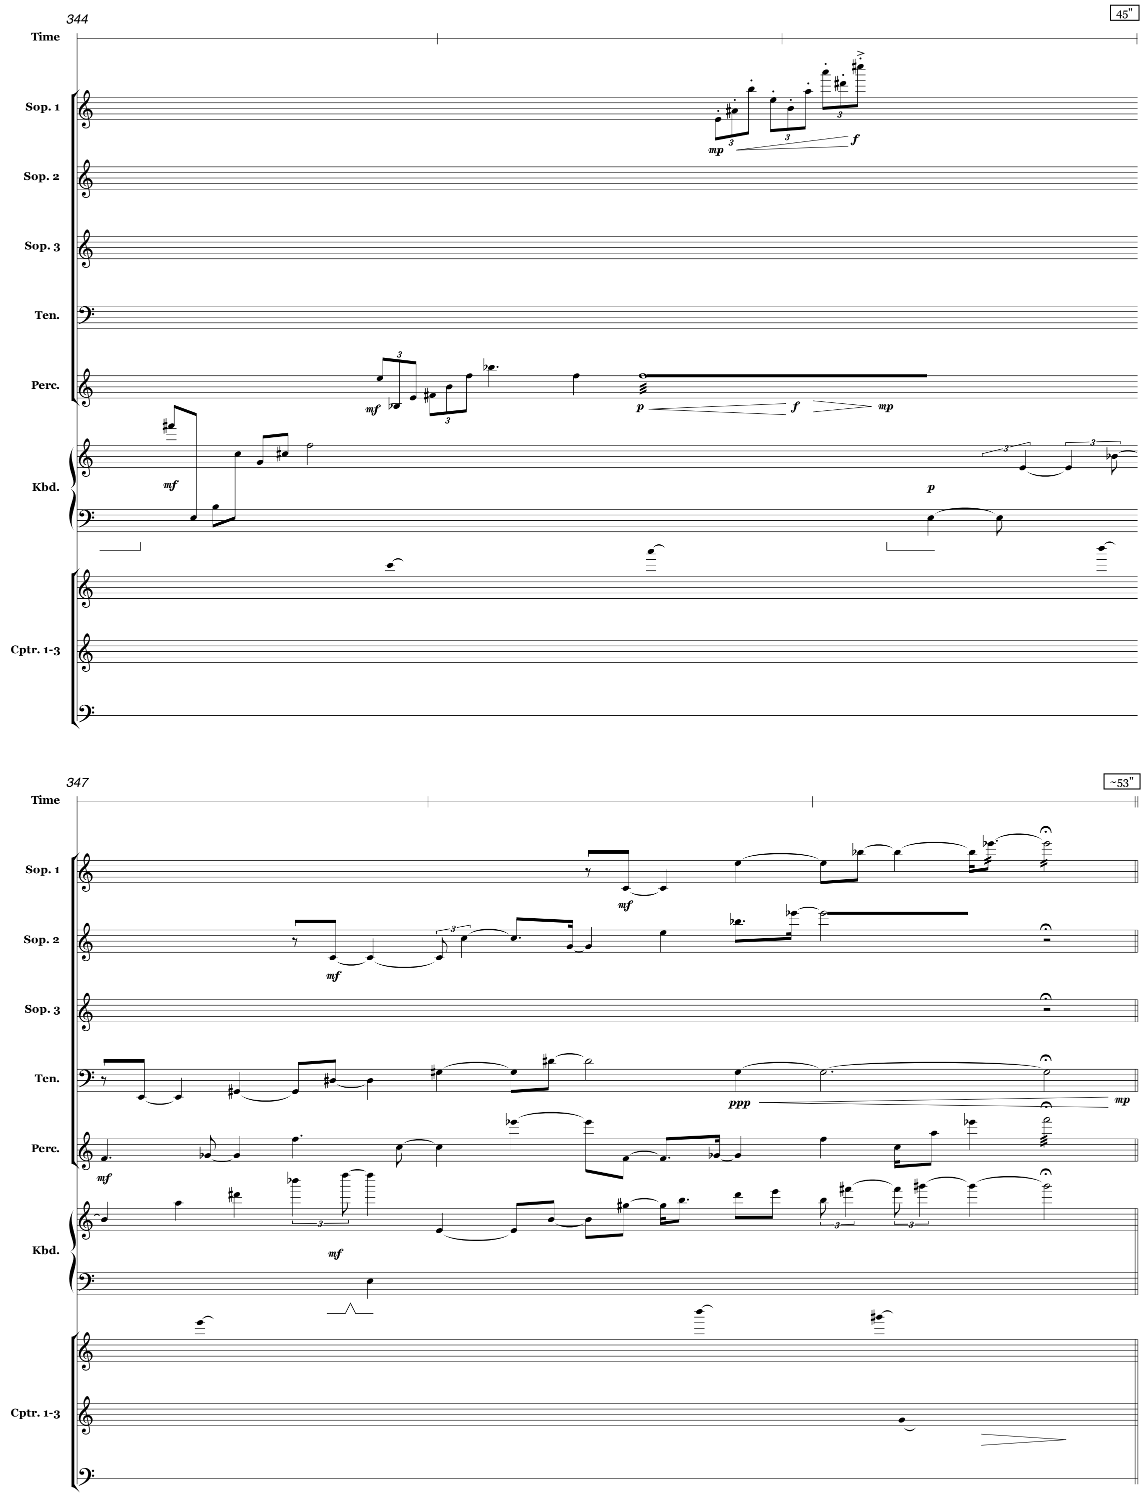
344
45"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
mp
f
mf
p
mp
347
~53"
mf
ppp
mp

III. A Twisted Pair

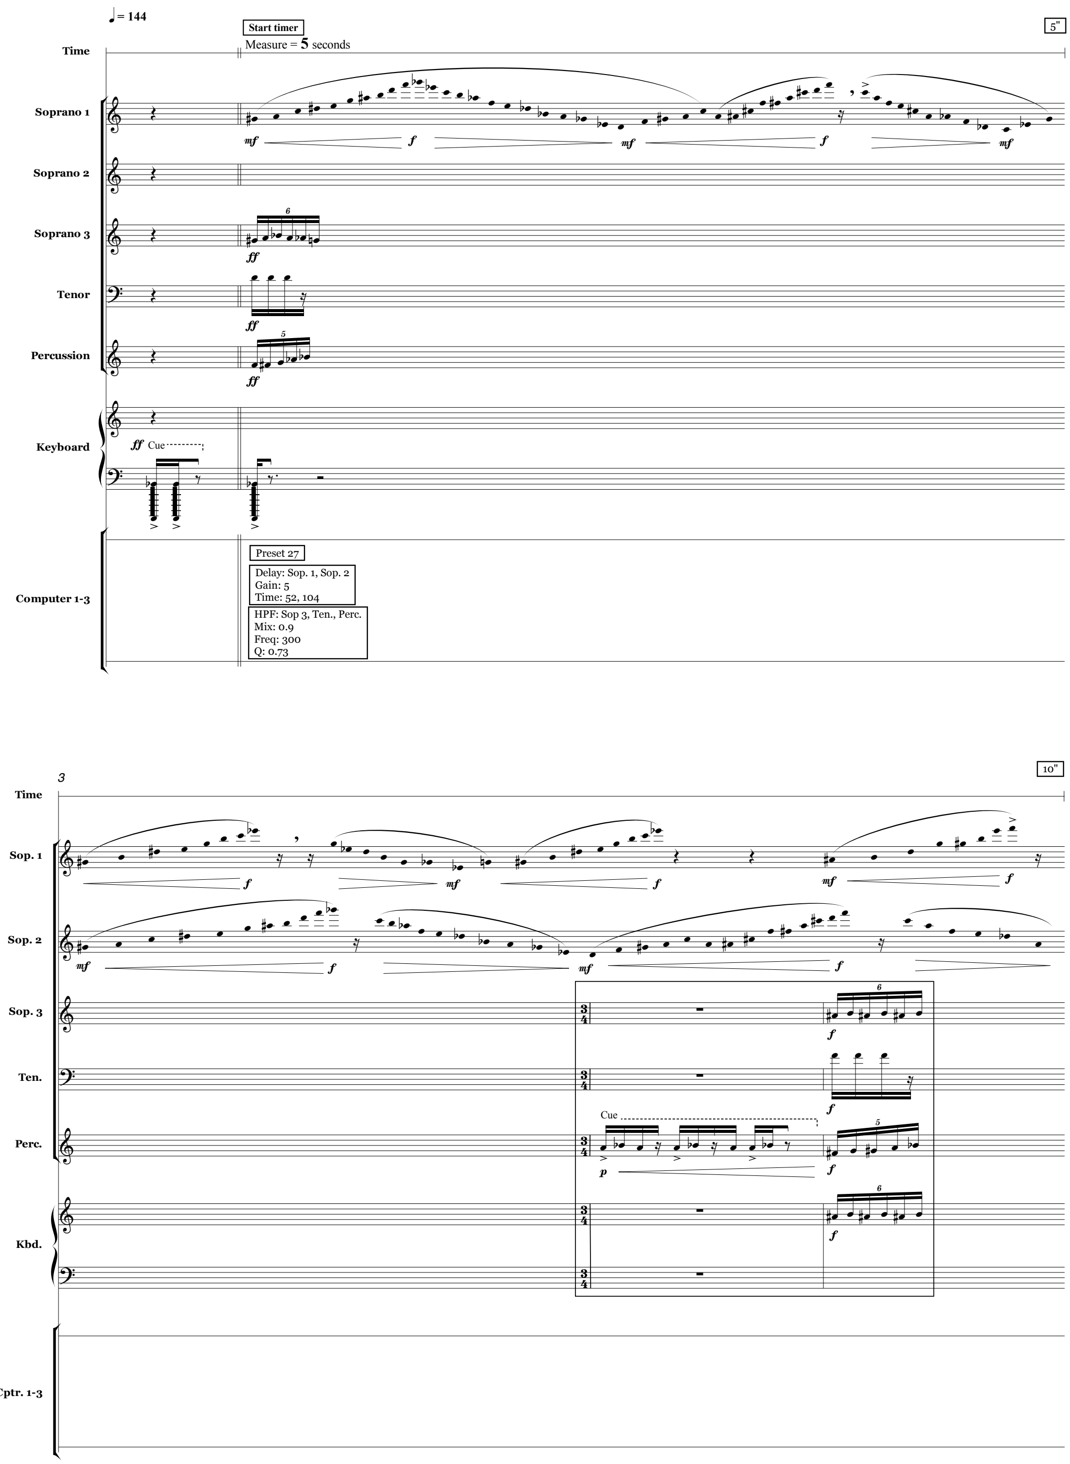

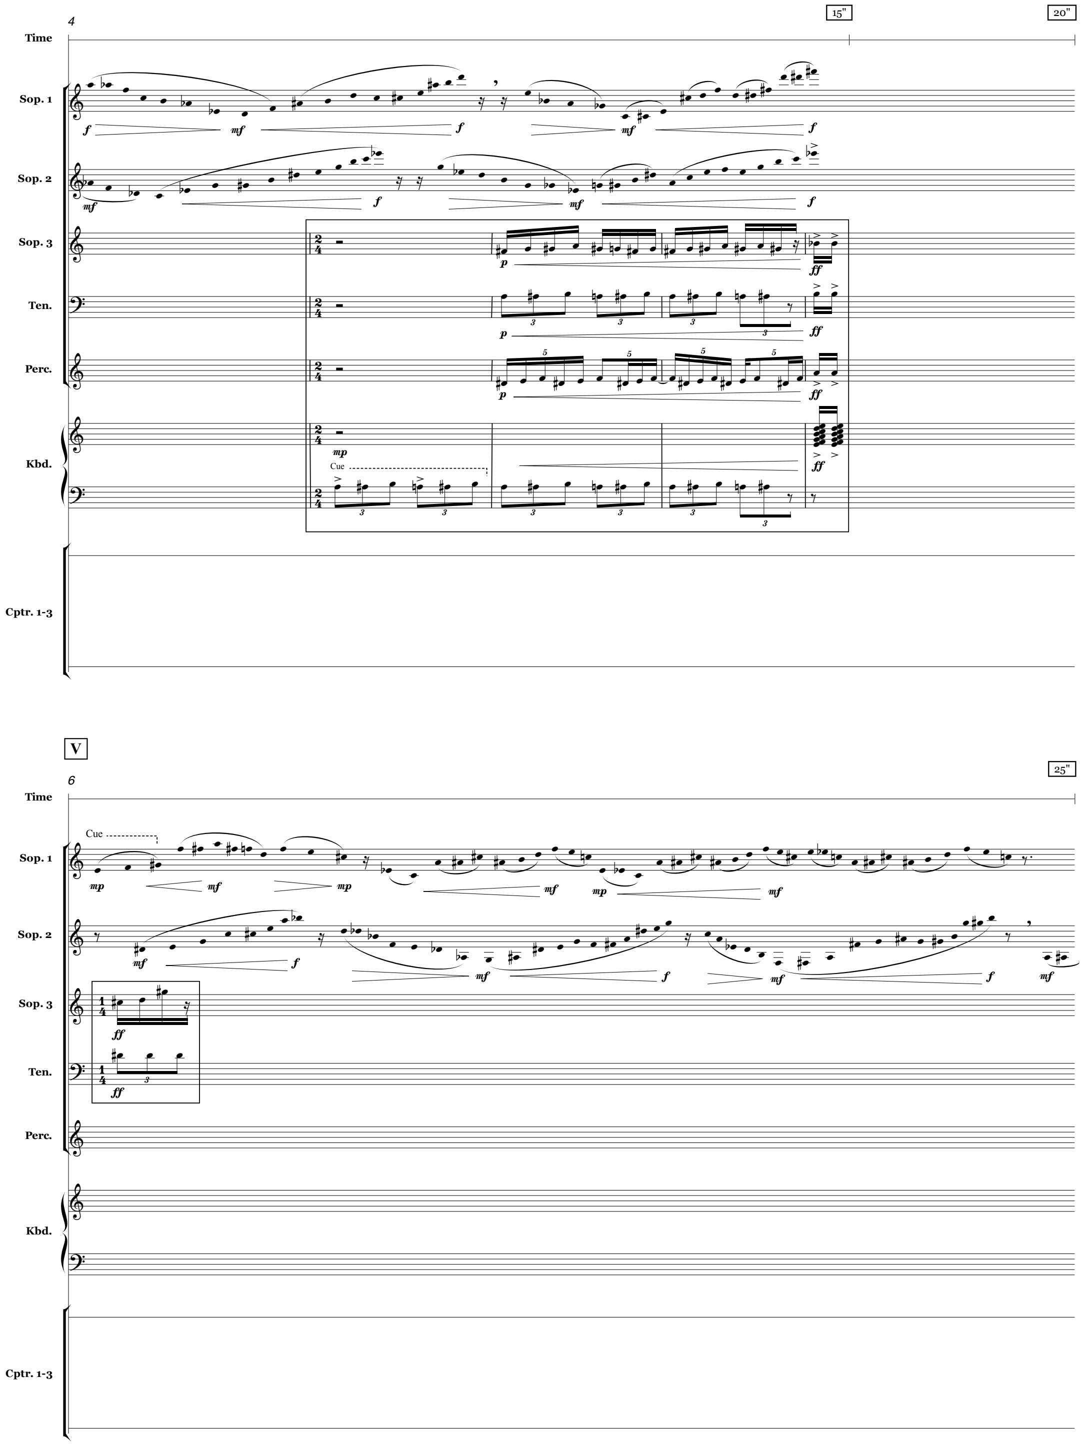

4
15"
20"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cue
Cptr. 1-3
V
6
25"
Time
Sop. 1
Cue
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3

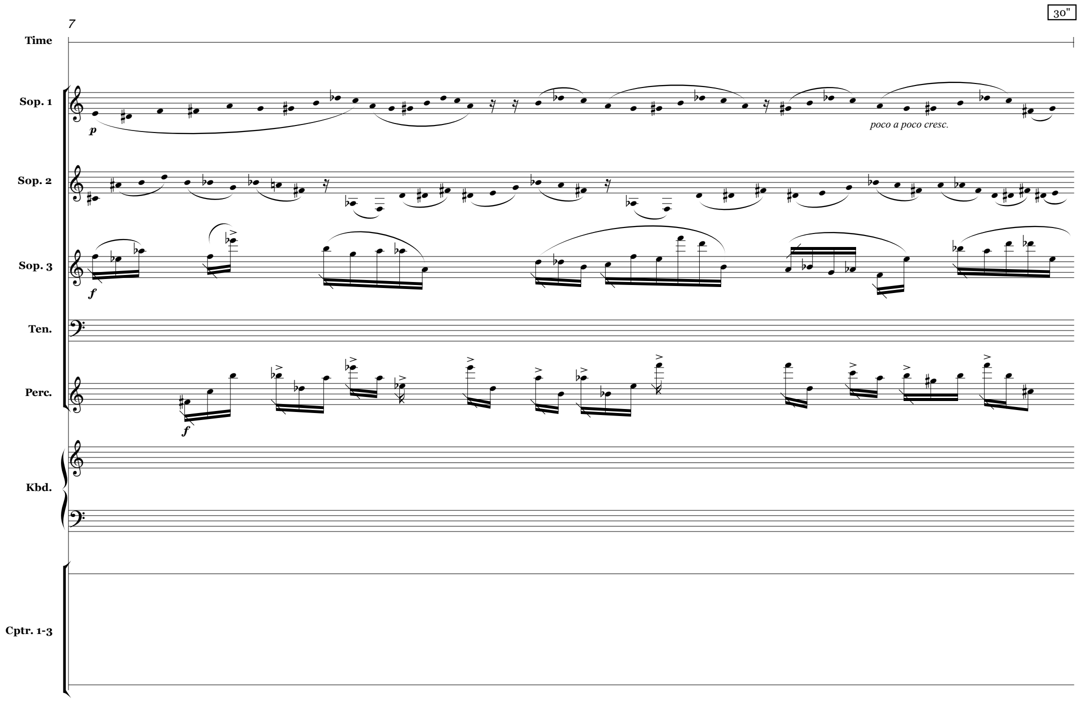
30"
7
Time
Sop. 1
p
poco a poco cresc.
Sop. 2
Sop. 3
f
Ten.
Perc.
f
Kbd.
Cptr. 1-3

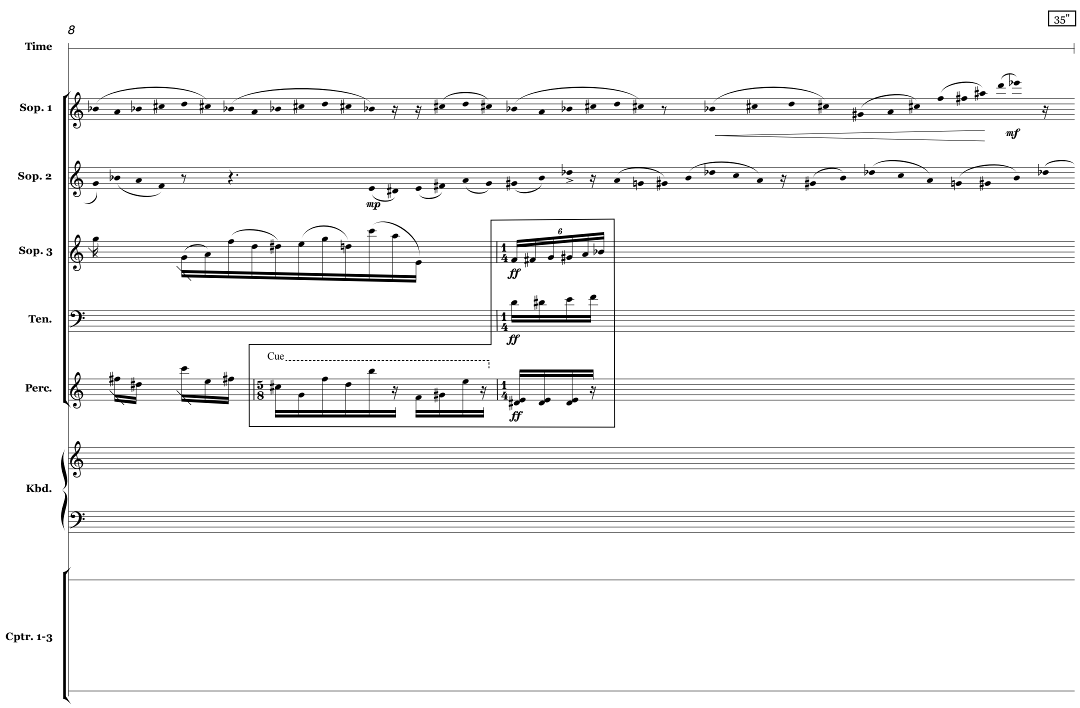
35"
8
Time
Sop. 1
mf
Sop. 2
mp
Sop. 3
6
ff
Ten.
ff
Cue
Perc.
ff
Kbd.
Cptr. 1-3

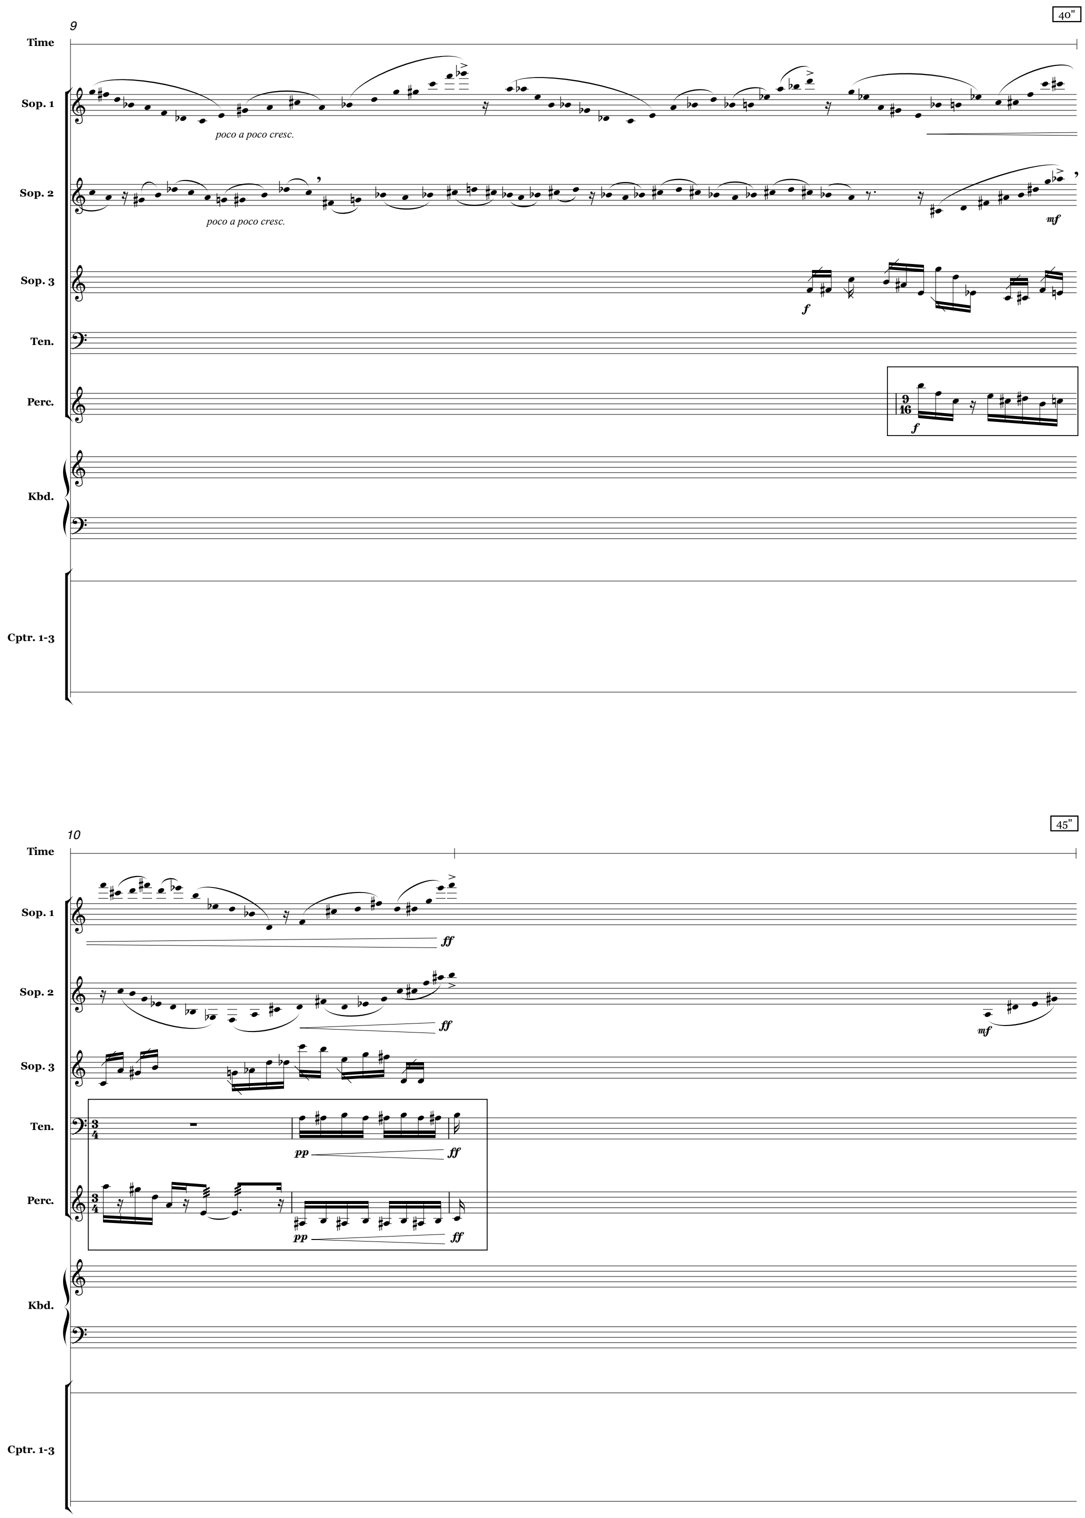
9
40"
Time
Sop. 1
poco a poco cresc.
Sop. 2
poco a poco cresc.
mf
Sop. 3
f
Ten.
Perc.
9
16
f
Kbd.
Cptr. 1-3
10
45"
Time
Sop. 1
ff
Sop. 2
ff
mf
Sop. 3
Ten.
3
4
pp
ff
Perc.
3
4
pp
ff
Kbd.
Cptr. 1-3

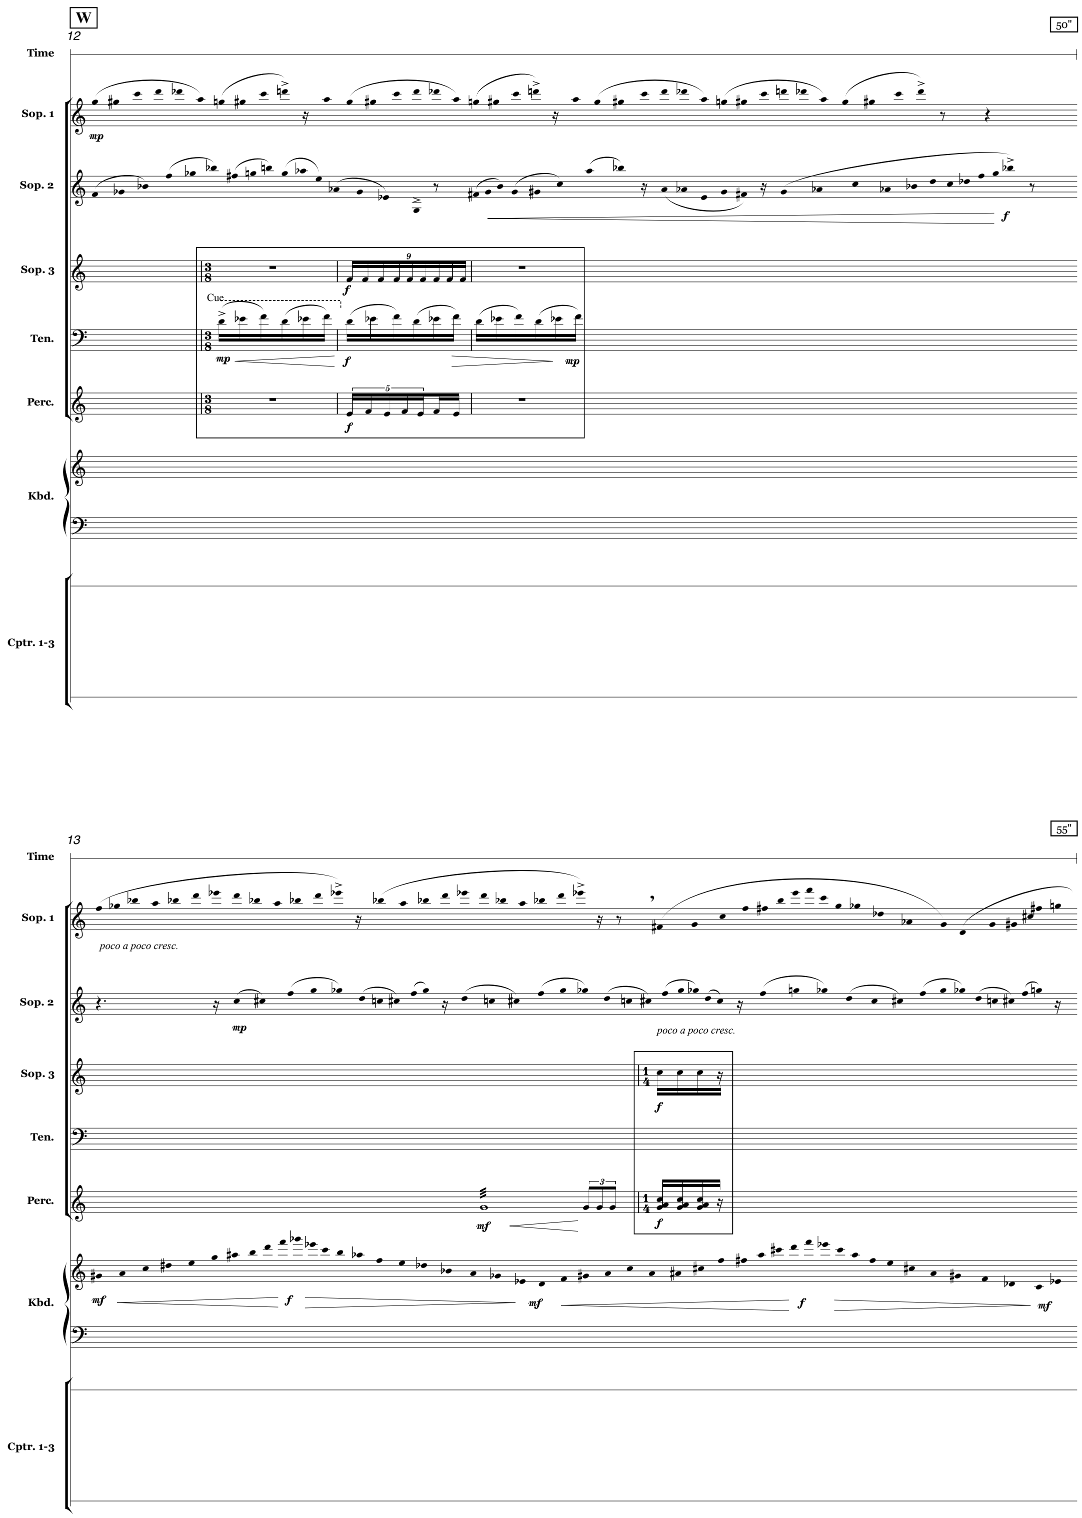
W
12
50"
Time
Sop. 1
mp
Sop. 2
f
Sop. 3
9
Cue
Ten.
mp
f
mp
Perc.
5
Kbd.
Cptr. 1-3
13
55"
Time
Sop. 1
poco a poco cresc.
Sop. 2
mp
poco a poco cresc.
Sop. 3
Ten.
Perc.
mf
3
Kbd.
mf
f
mf
f
mf
Cptr. 1-3

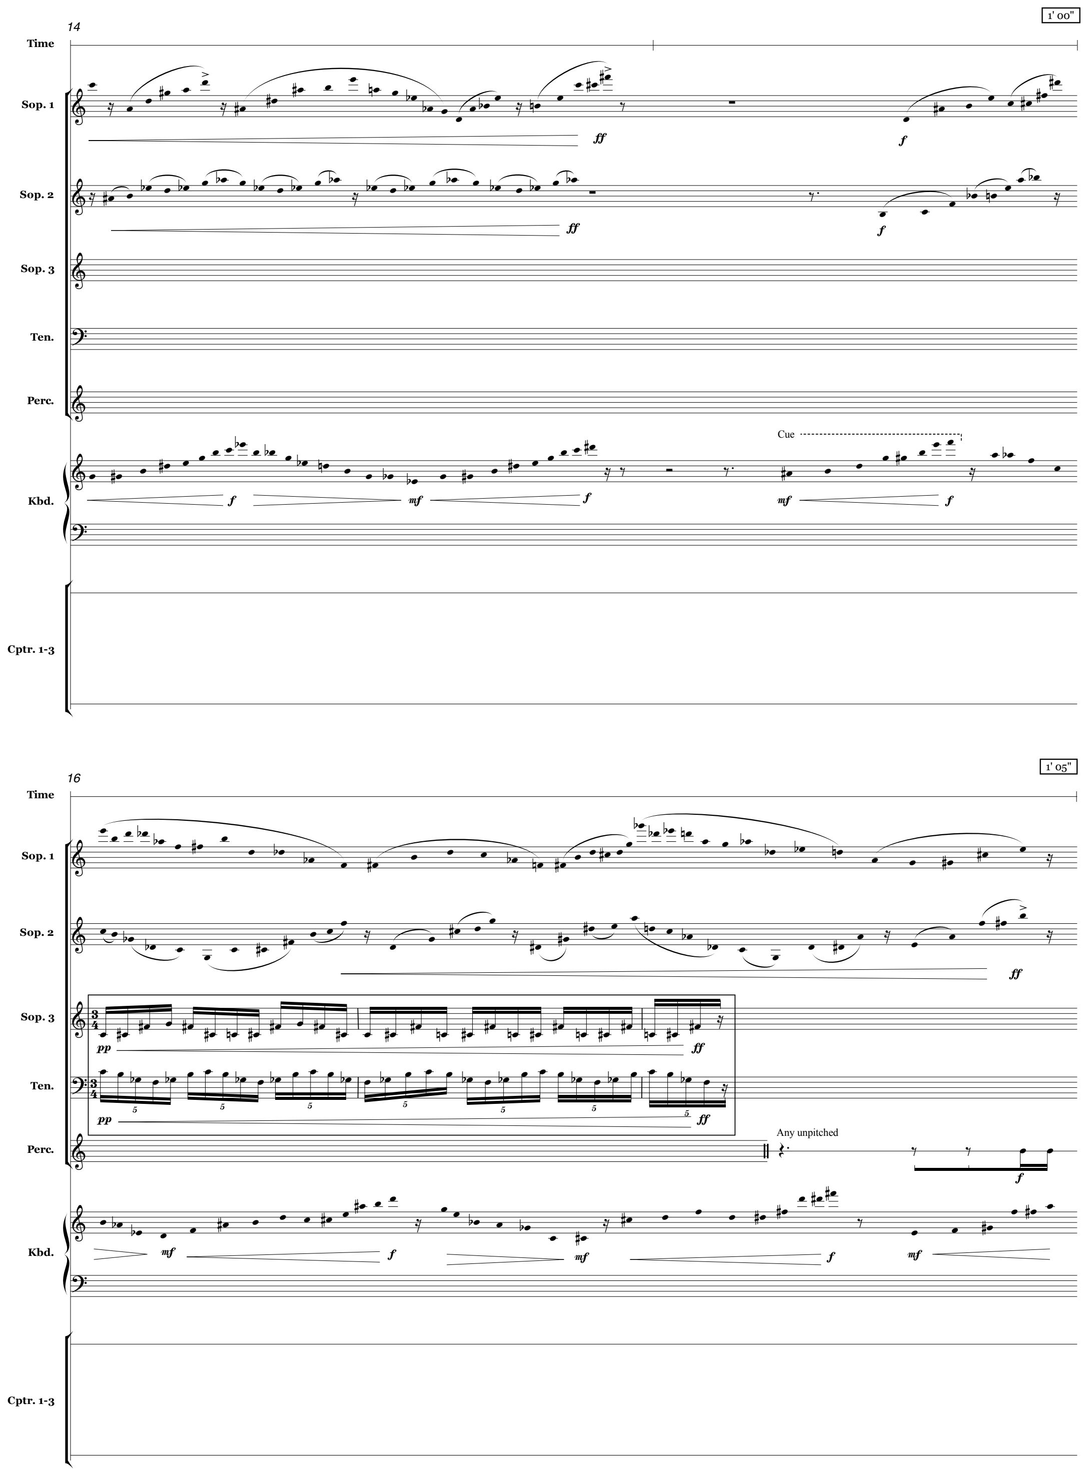
14
1' 00"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Cue
Kbd.
Cptr. 1-3
16
1' 05"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Any unpitched
Perc.
Kbd.
Cptr. 1-3

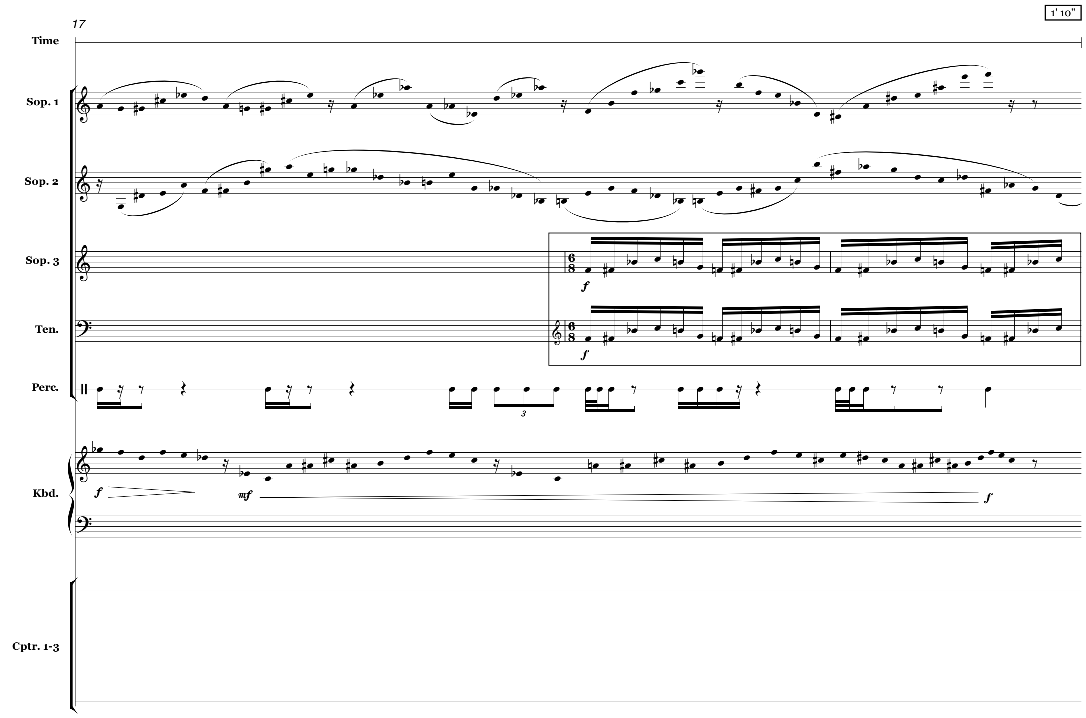
17
1' 10"
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
f
mf
3

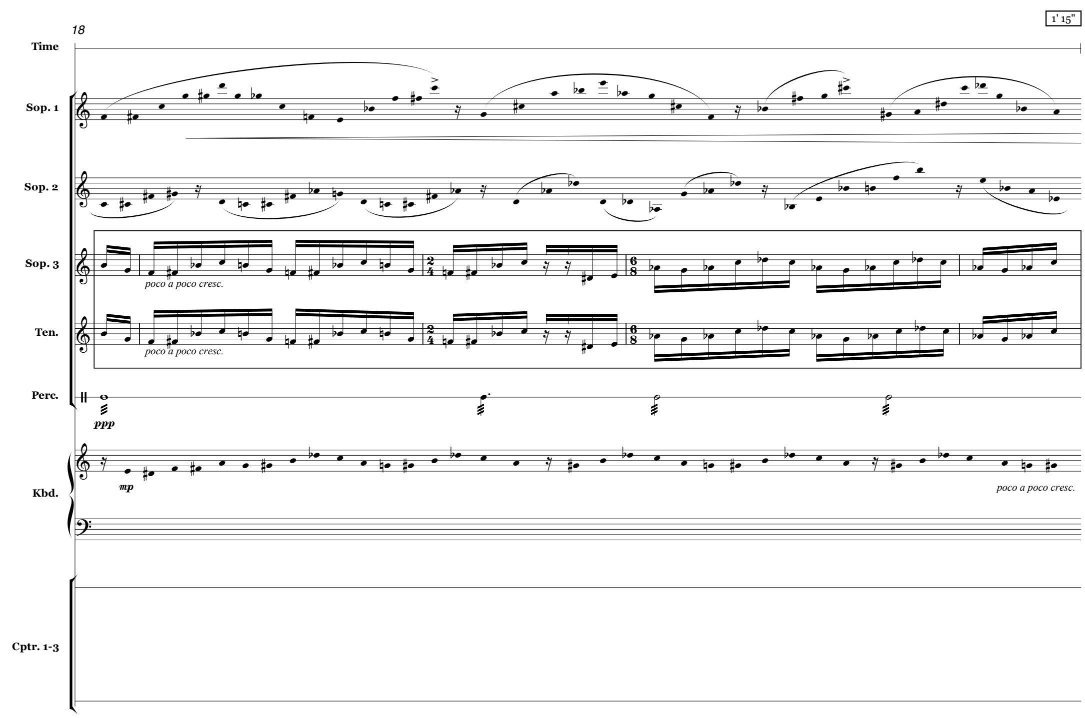
18
1' 15"
Time
Sop. 1
Sop. 2
Sop. 3
poco a poco cresc.
Ten.
poco a poco cresc.
Perc.
ppp
Kbd.
mp
poco a poco cresc.
Cptr. 1-3

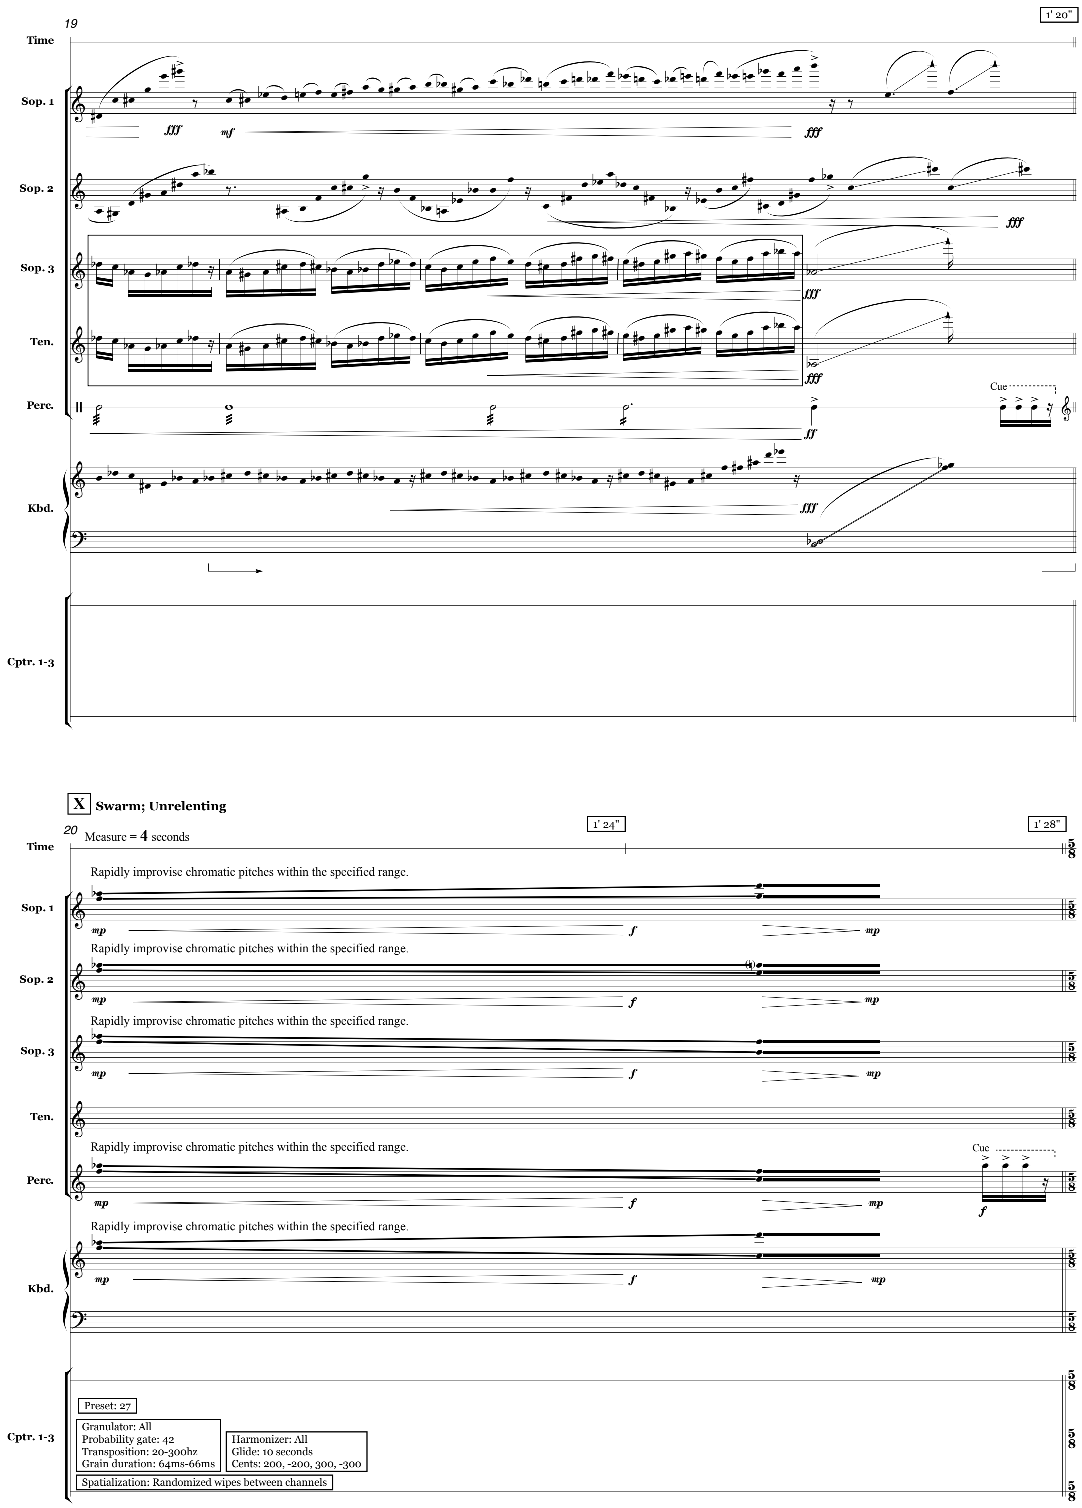

19
1' 20"
Time
Sop. 1
fff
mf
fff
Sop. 2
fff
Sop. 3
fff
Ten.
fff
Perc.
ff
Cue
Kbd.
fff
Cptr. 1-3
X Swarm; Unrelenting
20 Measure = 4 seconds
1' 24"
1' 28"
Time
Rapidly improvise chromatic pitches within the specified range.
Sop. 1
mp
f
mp
Rapidly improvise chromatic pitches within the specified range.
Sop. 2
mp
f
mp
Rapidly improvise chromatic pitches within the specified range.
Sop. 3
mp
f
mp
Ten.
Rapidly improvise chromatic pitches within the specified range.
Perc.
mp
f
mp
Cue
f
Rapidly improvise chromatic pitches within the specified range.
Kbd.
mp
f
mp
Cptr. 1-3
Preset: 27
Granulator: All
Probability gate: 42
Transposition: 20-300hz
Grain duration: 64ms-66ms
Harmonizer: All
Glide: 10 seconds
Cents: 200, -200, 300, -300
Spatialization: Randomized wipes between channels
5/8

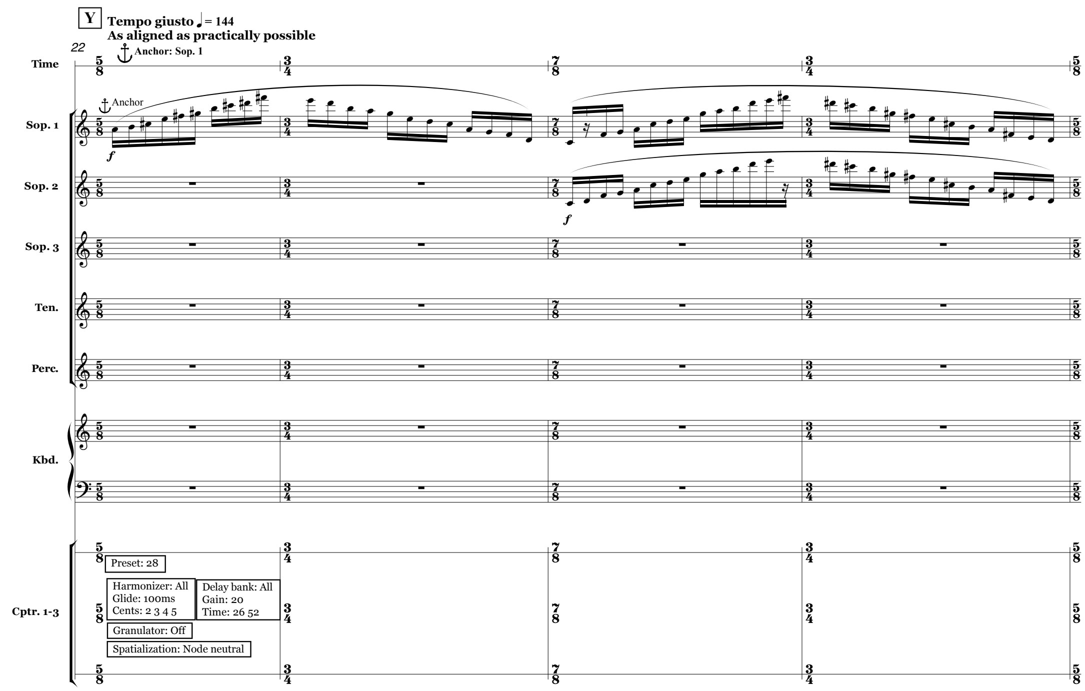
Y
Tempo giusto ♩= 144
As aligned as practically possible
22
Anchor: Sop. 1
Anchor
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
Preset: 28
Harmonizer: All
Glide: 100ms
Cents: 2 3 4 5
Delay bank: All
Gain: 20
Time: 26 52
Granulator: Off
Spatialization: Node neutral

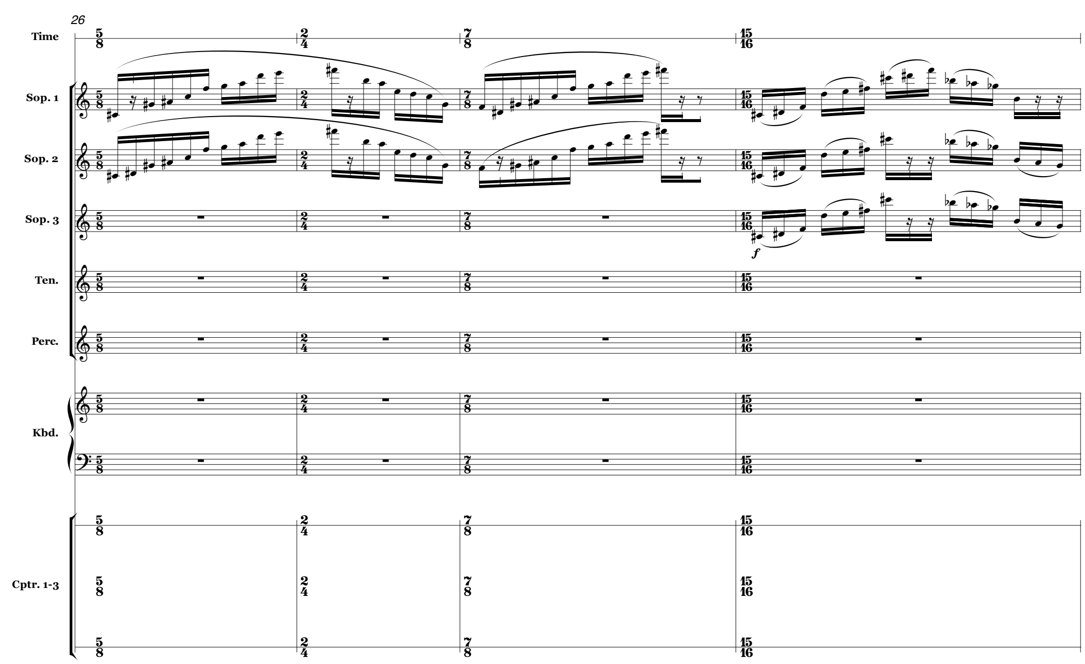
26
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3

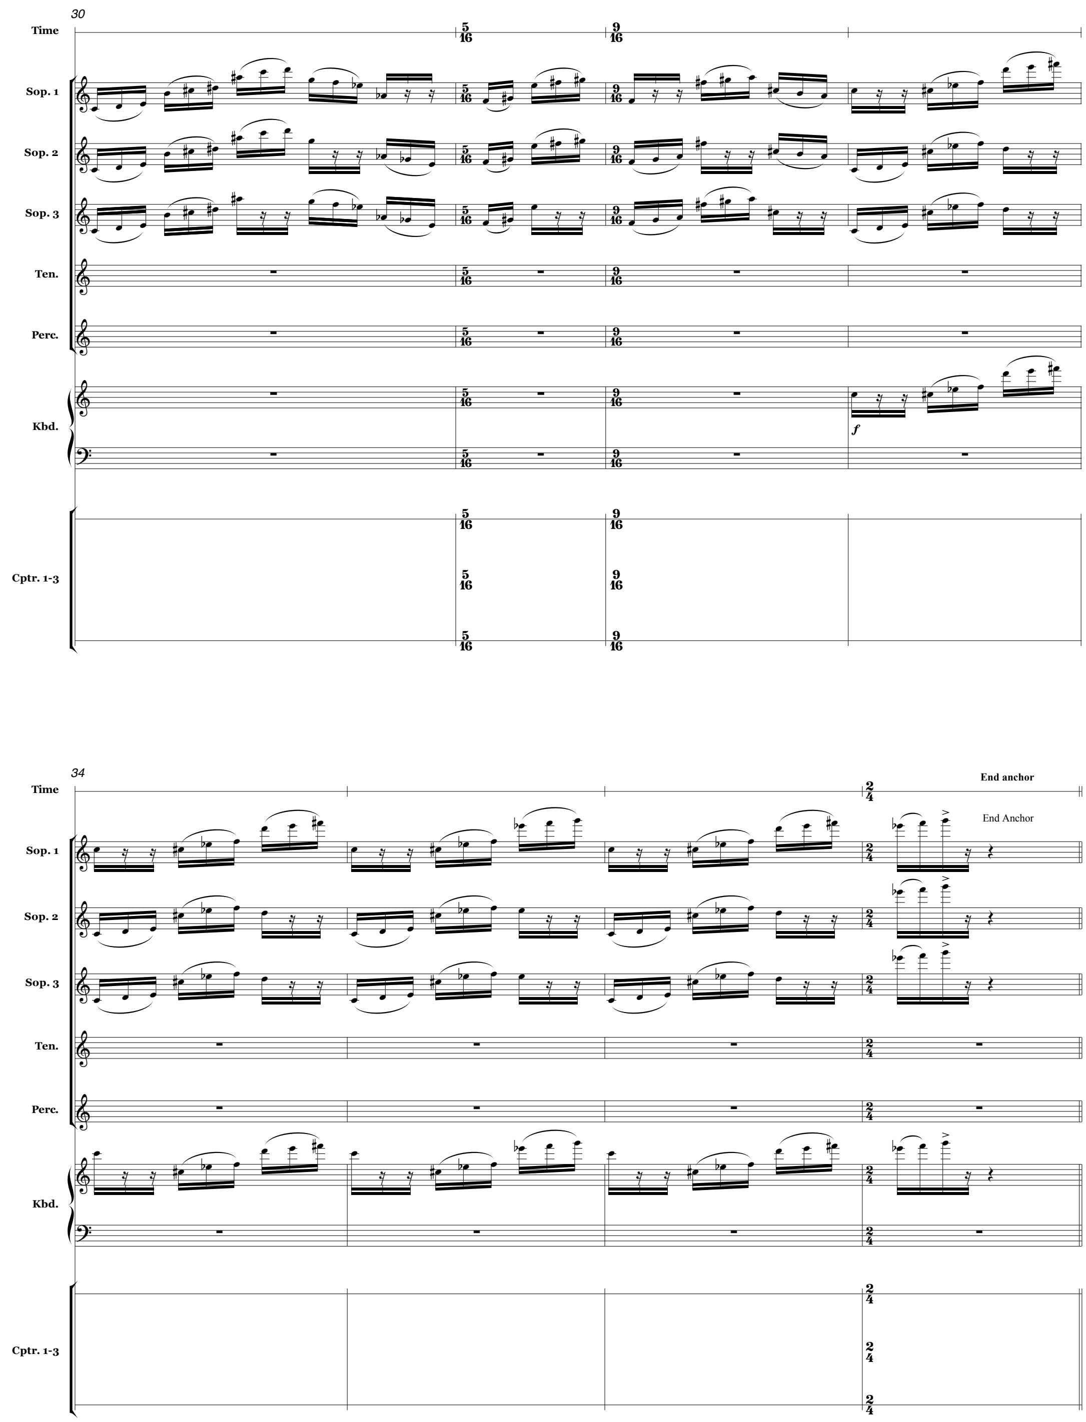
30
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
34
End anchor
End Anchor

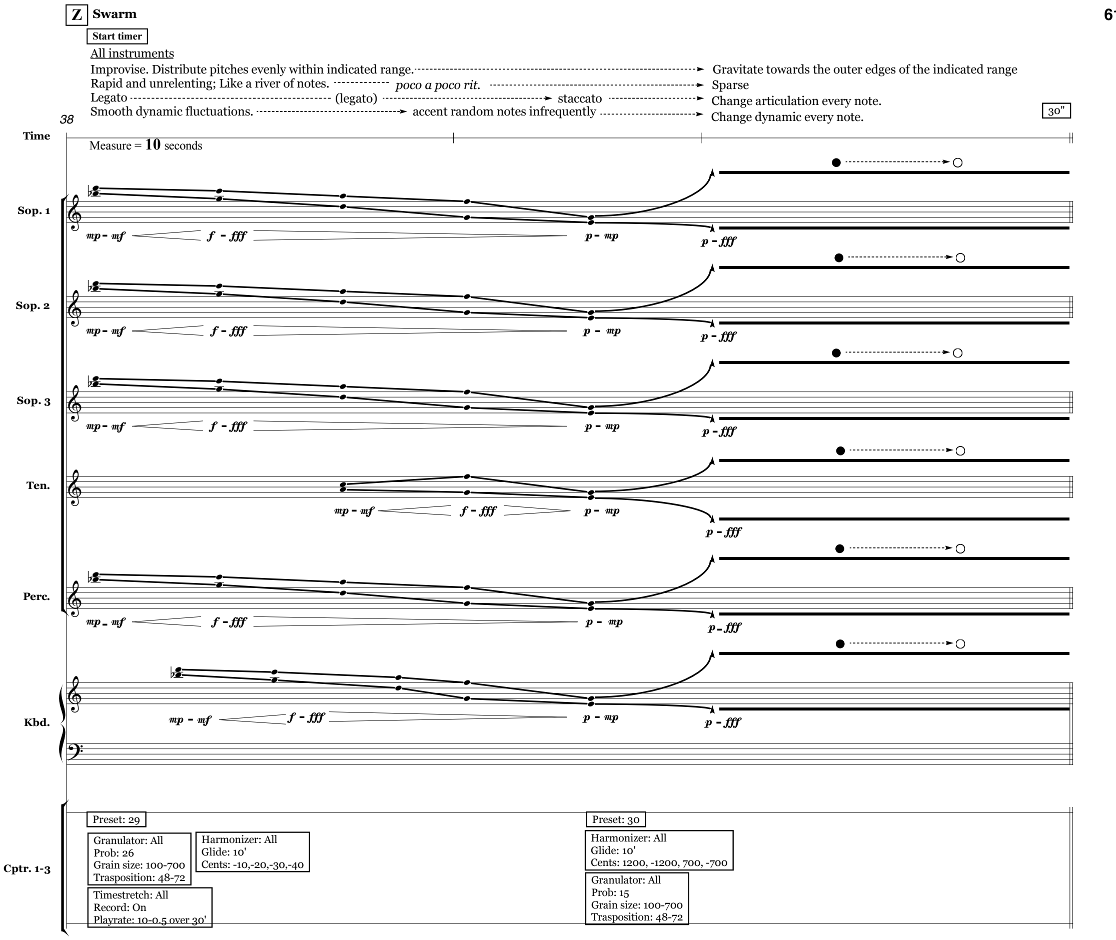
Z Swarm
61
Start timer
All instruments
Improvise. Distribute pitches evenly within indicated range. → Gravitate towards the outer edges of the indicated range
Rapid and unrelenting; Like a river of notes. poco a poco rit. → Sparse
Legato (legato) → staccato → Change articulation every note.
Smooth dynamic fluctuations. → accent random notes infrequently → Change dynamic every note.
30"
38
Time
Measure = 10 seconds
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
mp - mf
f - fff
p - mp
p - fff
Cptr. 1-3
Preset: 29
Granulator: All
Prob: 26
Grain size: 100-700
Trasposition: 48-72
Harmonizer: All
Glide: 10'
Cents: -10,-20,-30,-40
Timestretch: All
Record: On
Playrate: 10-0.5 over 30'
Preset: 30
Harmonizer: All
Glide: 10'
Cents: 1200, -1200, 700, -700
Granulator: All
Prob: 15
Grain size: 100-700
Trasposition: 48-72

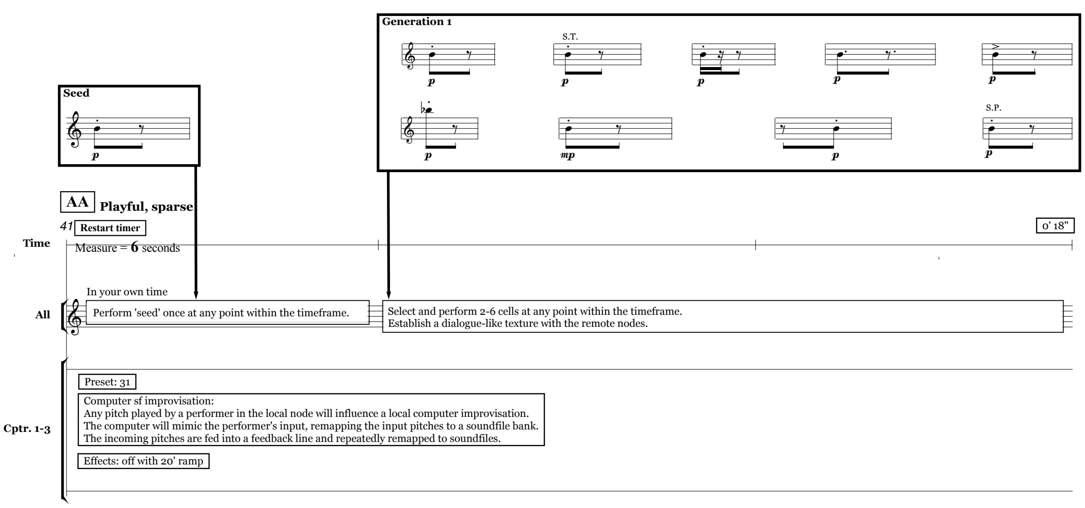
Seed
p
Generation 1
S.T.
S.P.
p
mp
AA Playful, sparse
41
Restart timer
0' 18"
Time
Measure = 6 seconds
In your own time
All
Perform 'seed' once at any point within the timeframe.
Select and perform 2-6 cells at any point within the timeframe.
Establish a dialogue-like texture with the remote nodes.
Cptr. 1-3
Preset: 31
Computer sf improvisation:
Any pitch played by a performer in the local node will influence a local computer improvisation.
The computer will mimic the performer's input, remapping the input pitches to a soundfile bank.
The incoming pitches are fed into a feedback line and repeatedly remapped to soundfiles.
Effects: off with 20' ramp

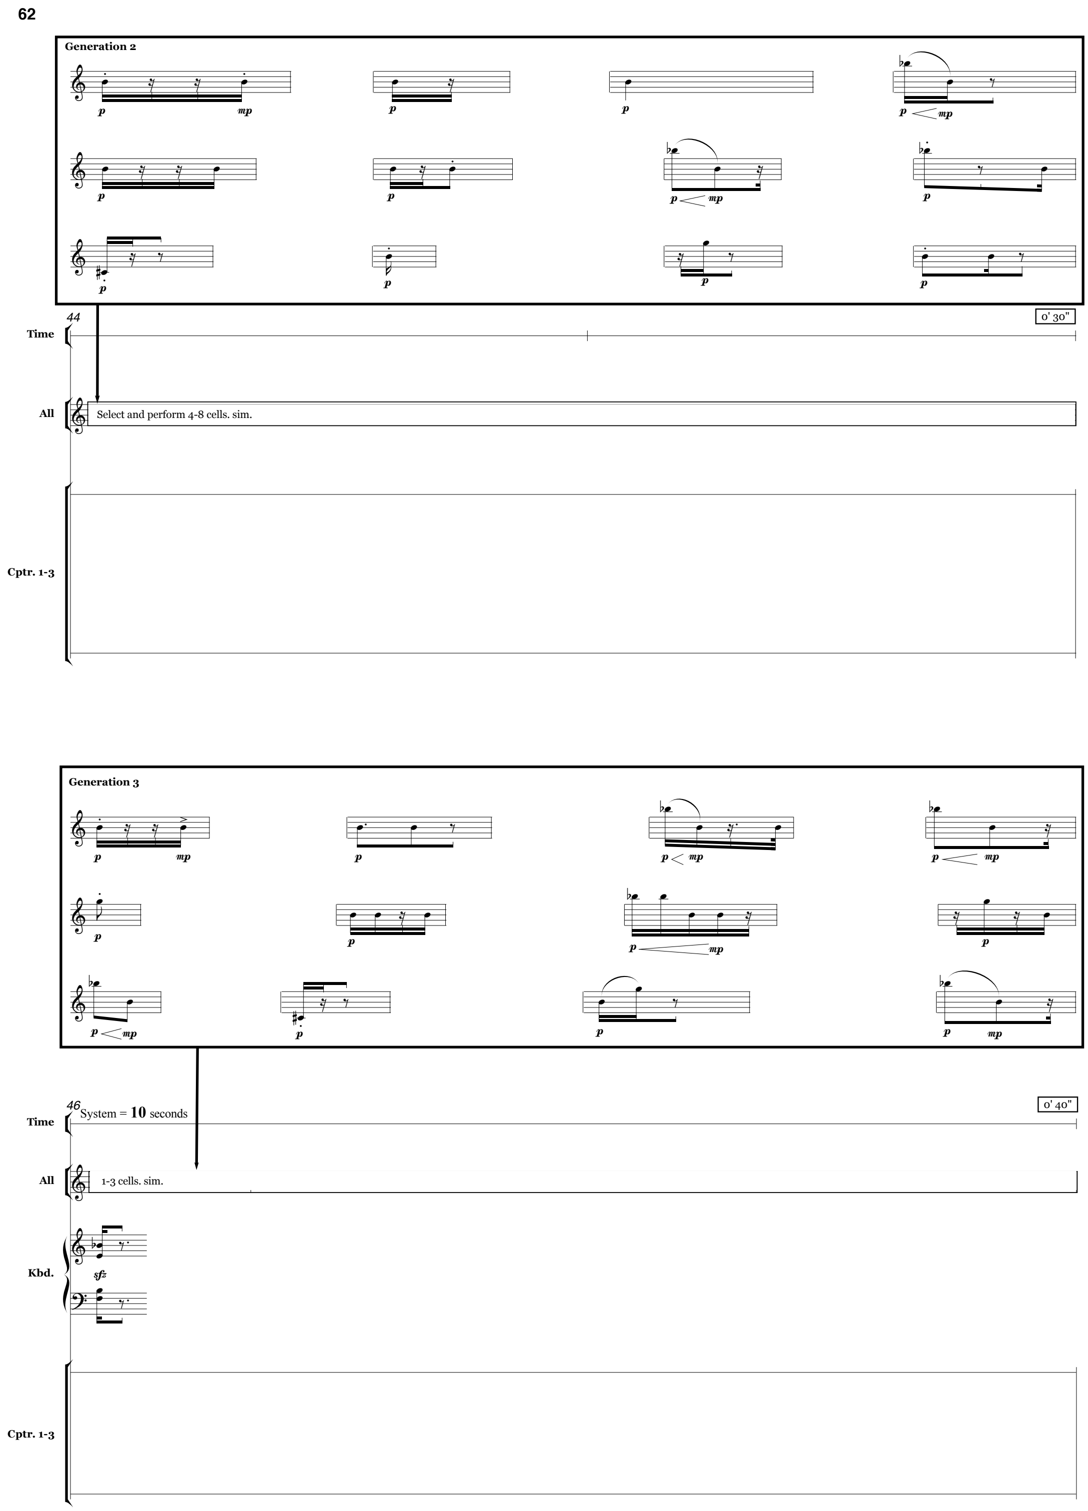
62
Generation 2
44
Time
0' 30"
All
Select and perform 4-8 cells. sim.
Cptr. 1-3
Generation 3
46
System = 10 seconds
0' 40"
Time
All
1-3 cells. sim.
Kbd.
Cptr. 1-3

63

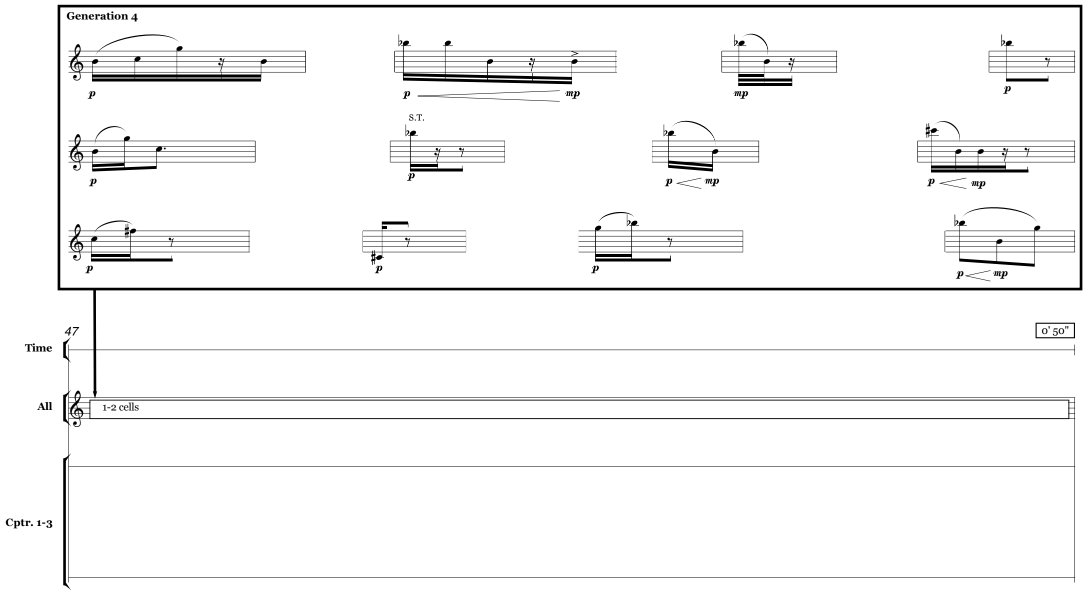
Generation 4
p
p
mp
mp
p
S.T.
p
p
p
mp
p
mp
p
p
p
p
mp
47
0' 50"
Time
All
1-2 cells
Cptr. 1-3

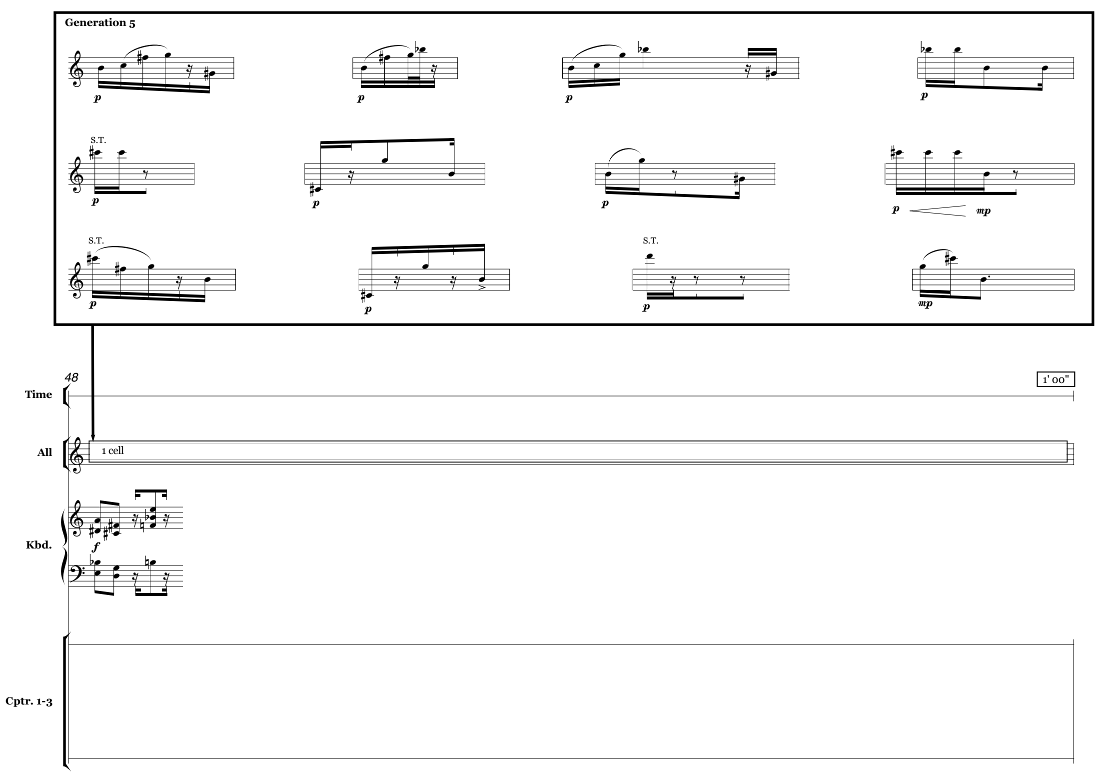
Generation 5
p
p
p
p
S.T.
p
p
p
p
mp
S.T.
p
p
S.T.
p
mp
48
1' 00"
Time
All
1 cell
Kbd.
f
Cptr. 1-3

64

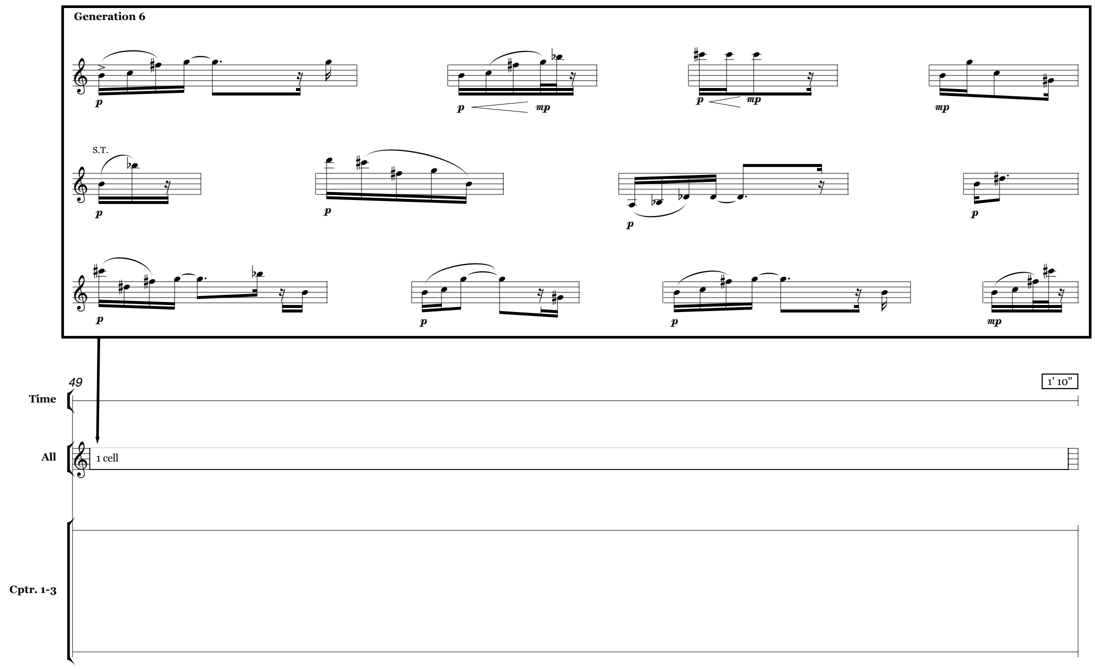
Generation 6
S.T.
p
mp
49
1' 10"
Time
All
1 cell
Cptr. 1-3

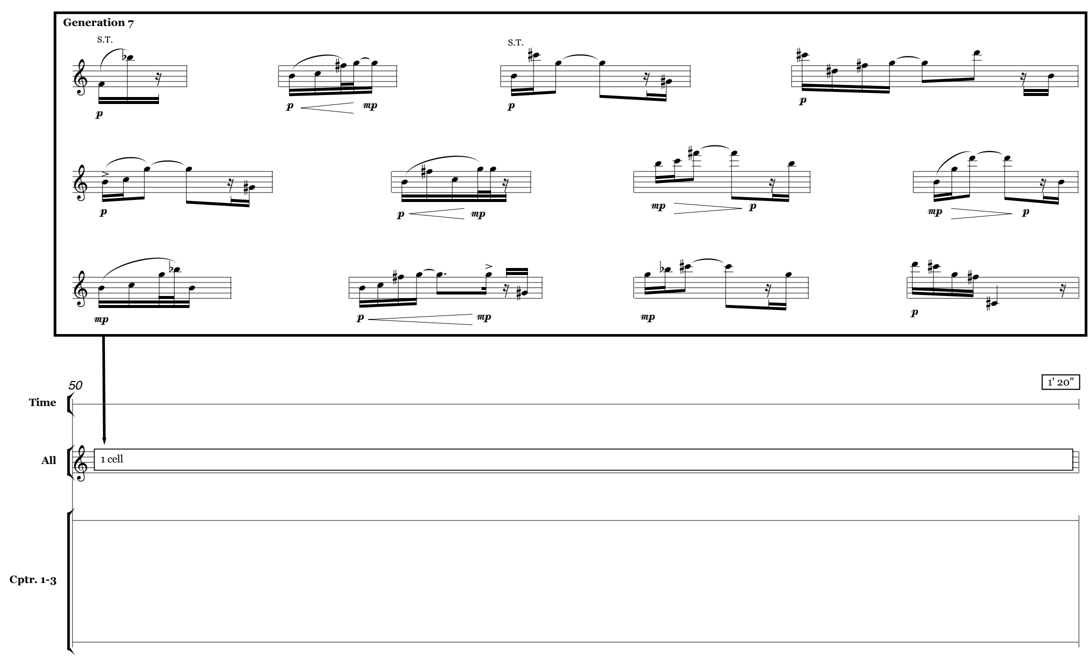
Generation 7
S.T.
p
mp
50
1' 20"
Time
All
1 cell
Cptr. 1-3

65

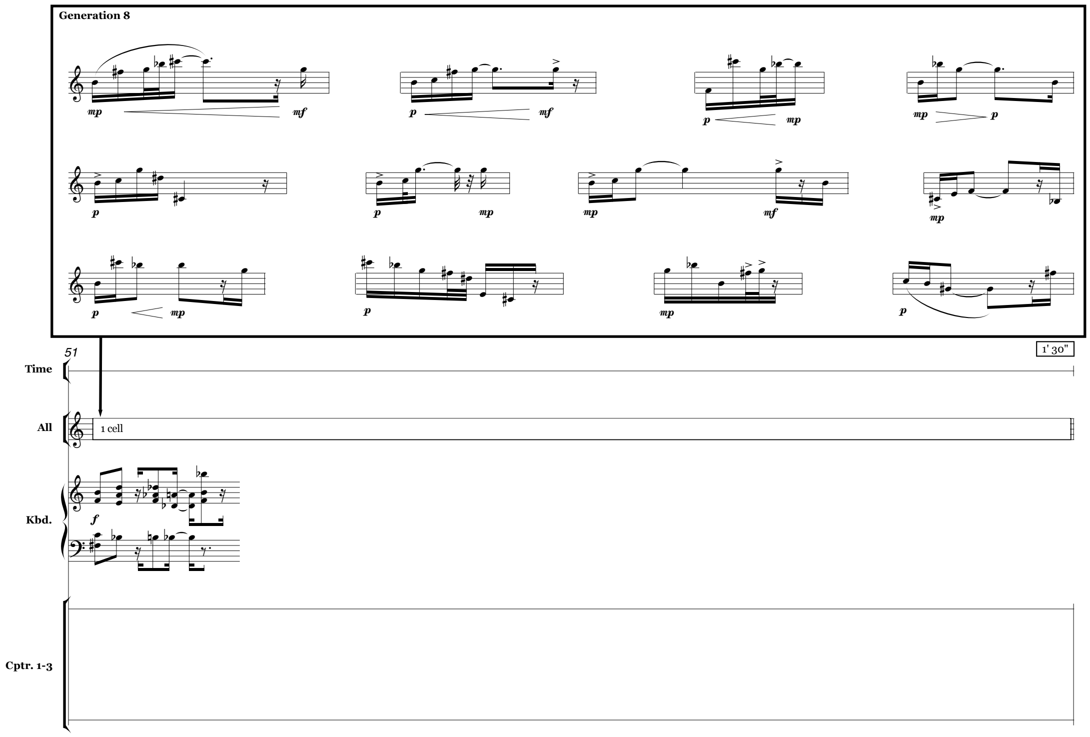
Generation 8
51
1' 30"
Time
All
1 cell
Kbd.
Cptr. 1-3

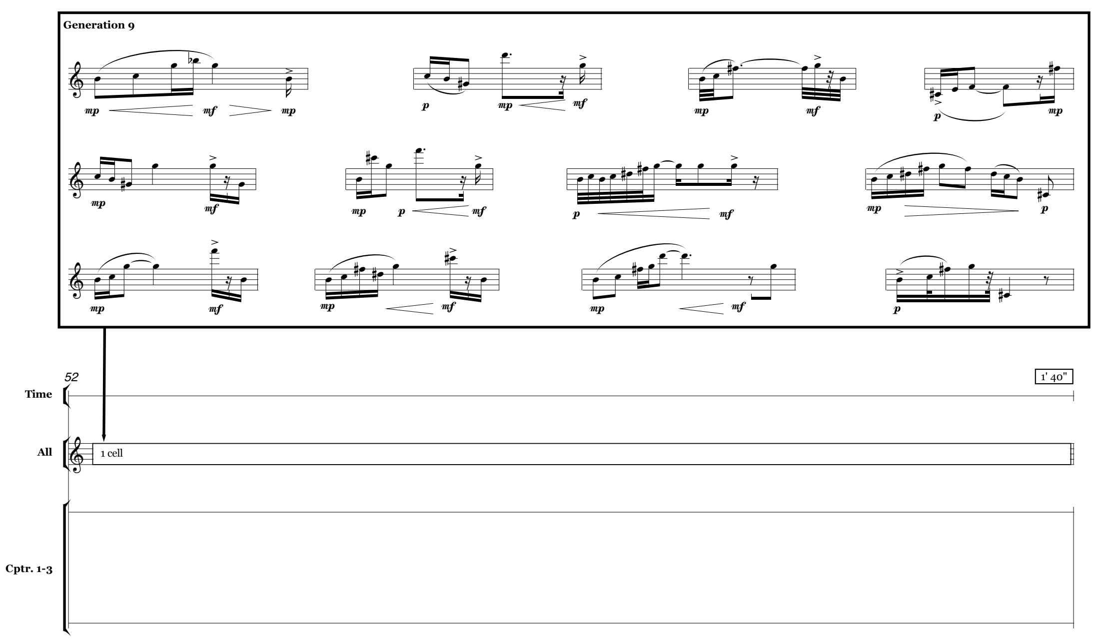
Generation 9
52
1' 40"
Time
All
1 cell
Cptr. 1-3

66

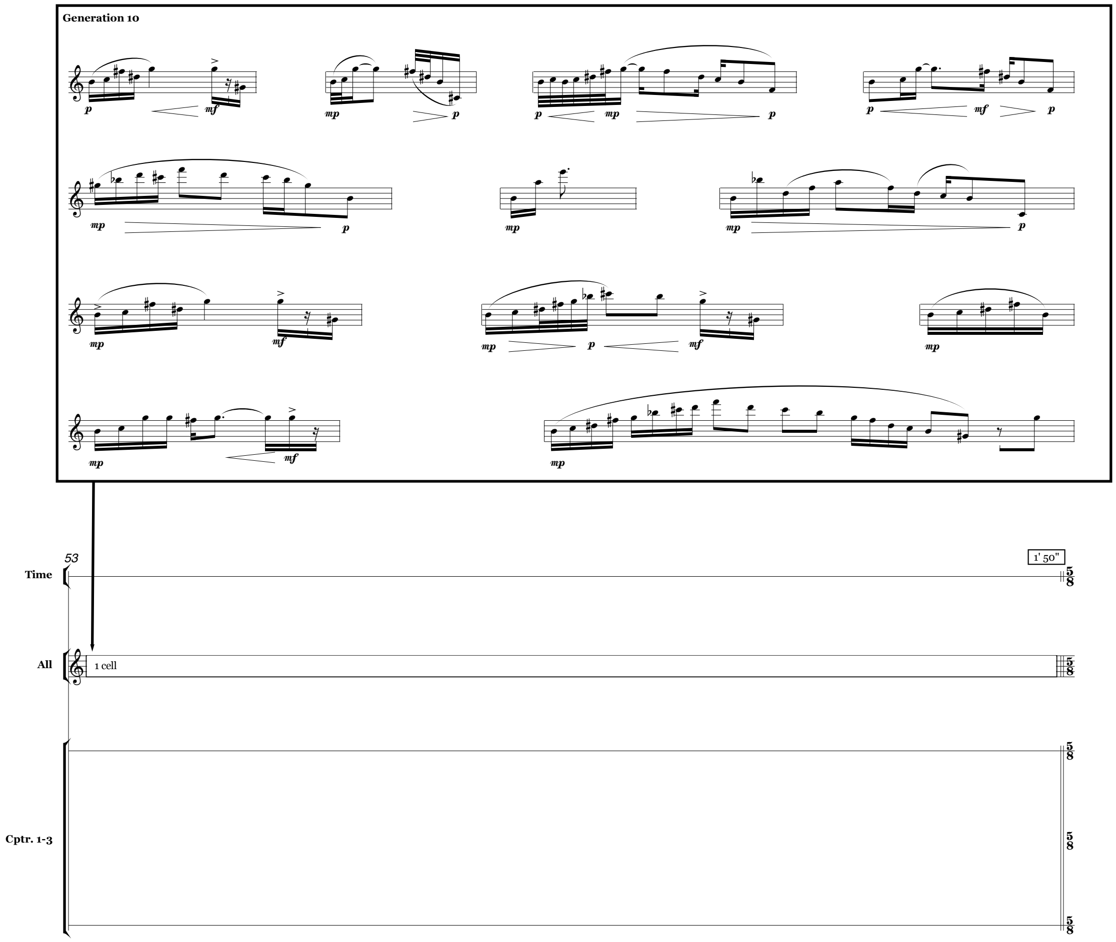
Generation 10
p mf
mp p
p mp p
p mf p
mp p
mp
mp p
mp mf
mp p mf
mp
mp mf
mp
53
1' 50"
Time
All
1 cell
Cptr. 1-3
5/8


BB
Chainfall ♩= 168
54
Time
Anchor: Sop. 1
Anchor: Perc.
Anchor
End anchor
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
Preset: 31
Computer improvisation: Off
59
Anchor: Sop. 2
Anchor: Sop. 1
Anchor: Kbd.

68

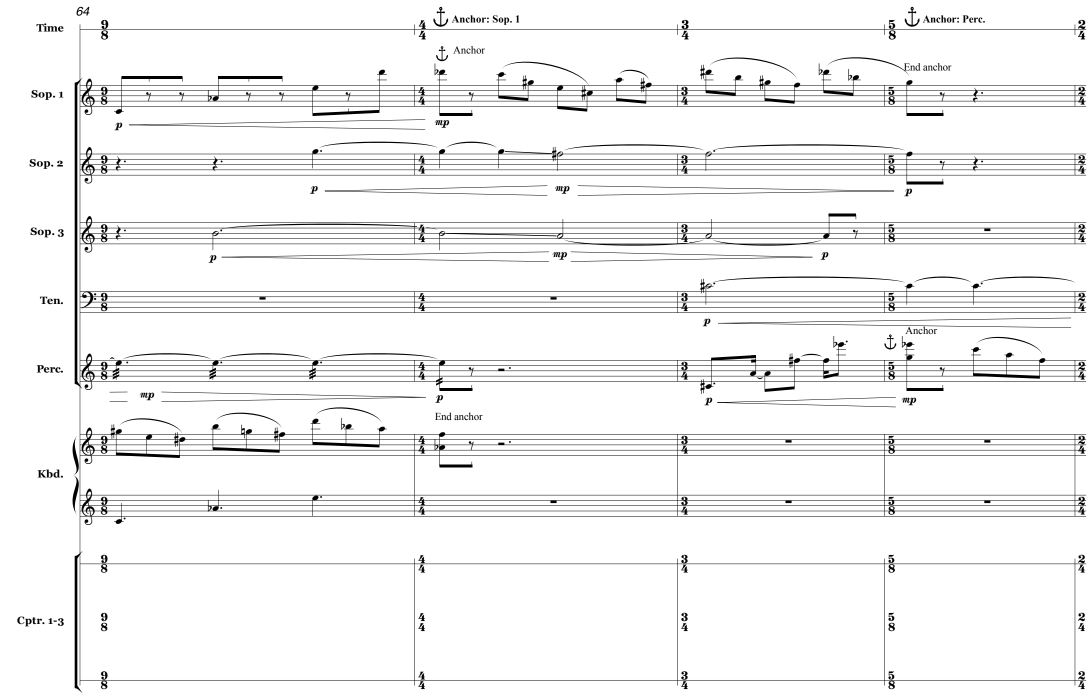
64
Time
Anchor: Sop. 1
Anchor: Perc.
Sop. 1
Anchor
End anchor
Sop. 2
Sop. 3
Ten.
Perc.
Anchor
Kbd.
End anchor
Cptr. 1-3

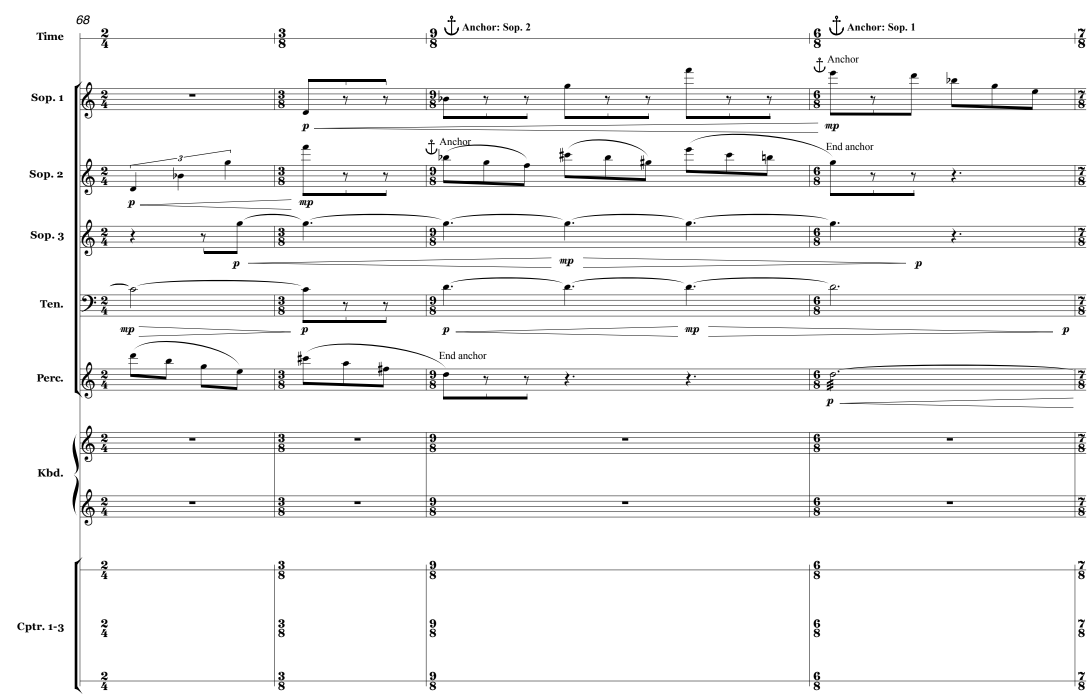
68
Time
Anchor: Sop. 2
Anchor: Sop. 1
Sop. 1
Anchor
Sop. 2
Anchor
End anchor
Sop. 3
Ten.
Perc.
End anchor
Kbd.
Cptr. 1-3

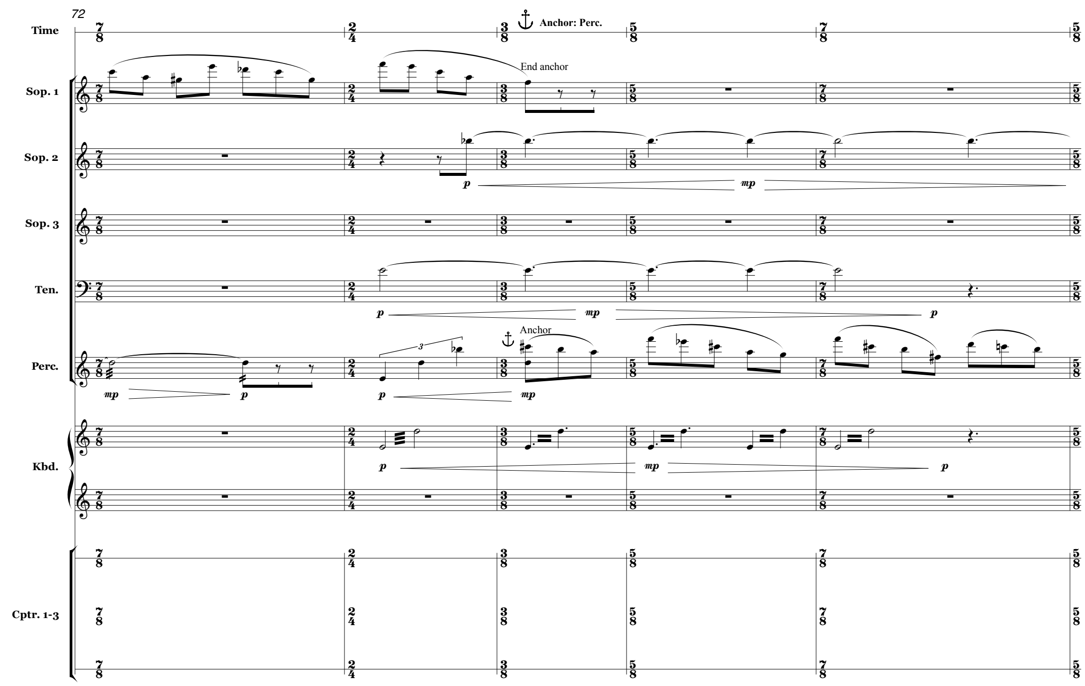
72
Time
Anchor: Perc.
End anchor
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Anchor
Kbd.
Cptr. 1-3

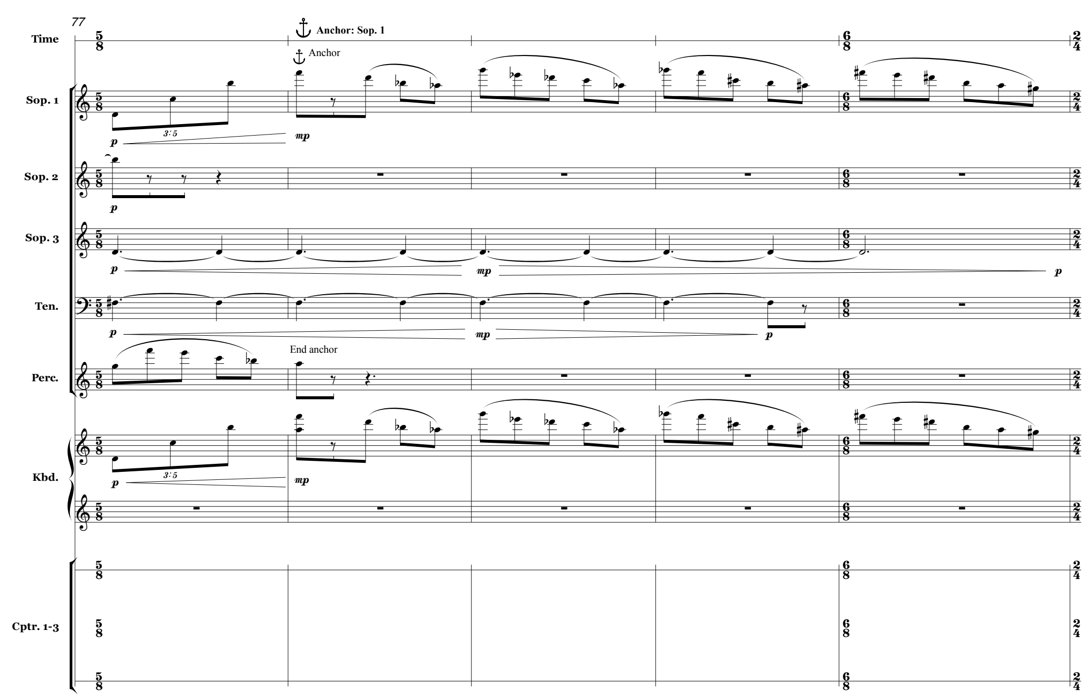
77
Time
Anchor: Sop. 1
Anchor
Sop. 1
Sop. 2
Sop. 3
Ten.
End anchor
Perc.
Kbd.
Cptr. 1-3

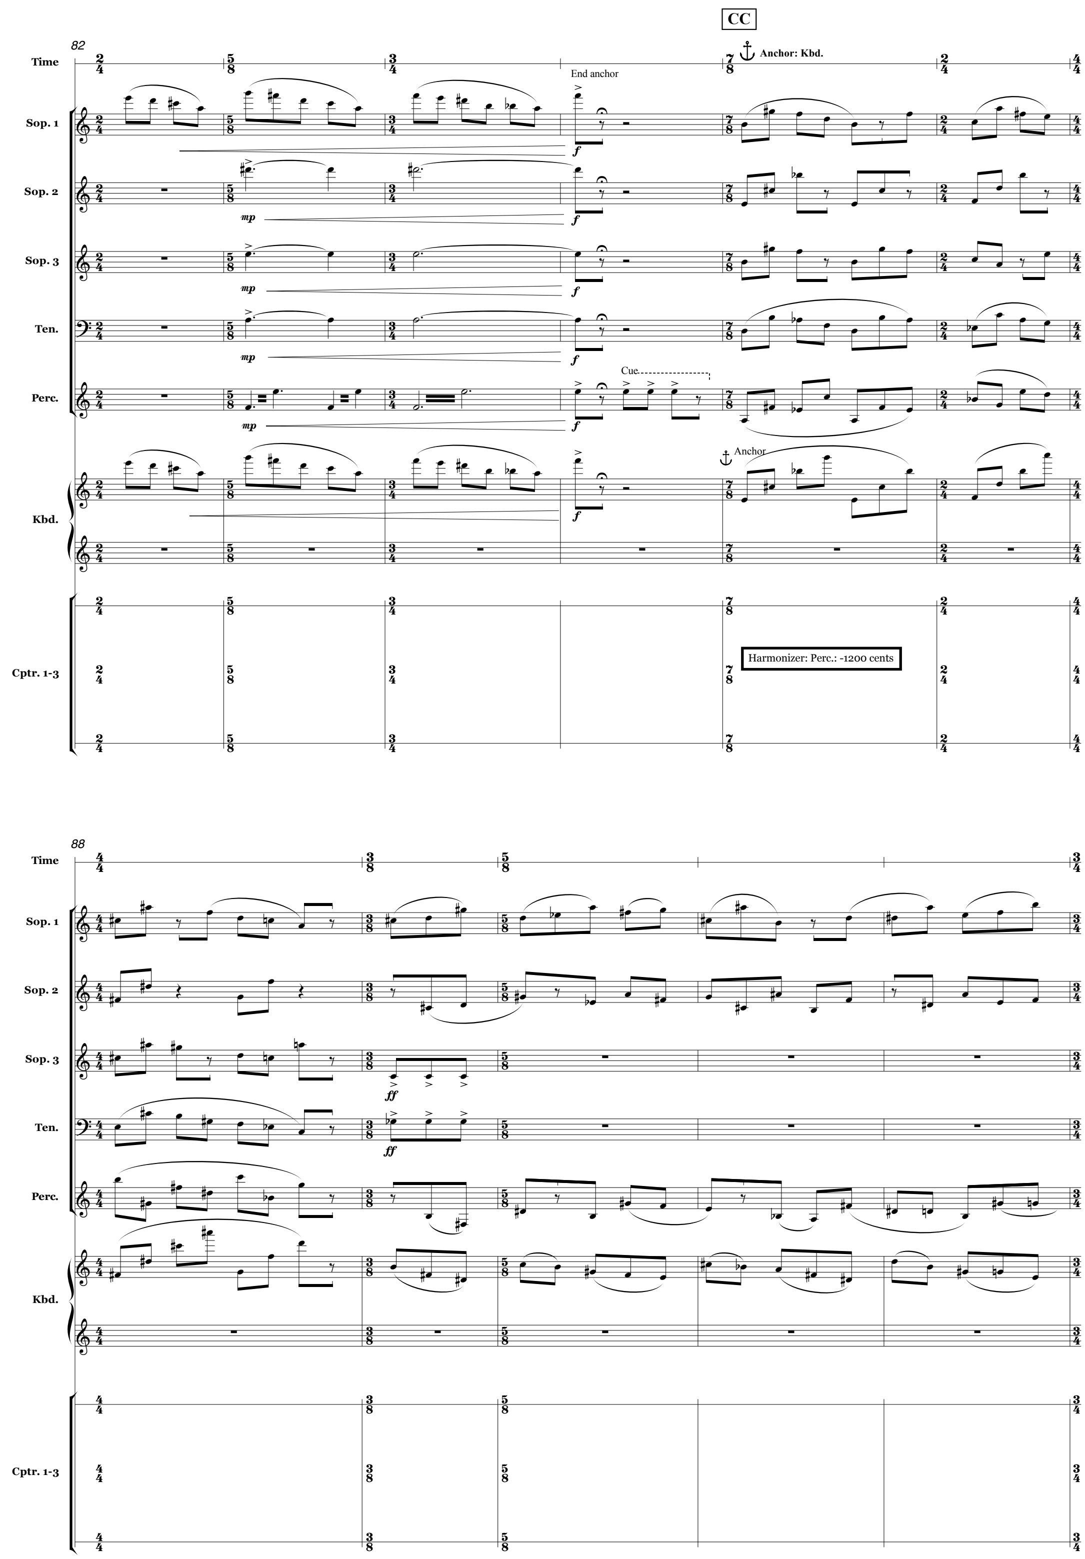
82
CC
Anchor: Kbd.
End anchor
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
mp
f
Cue
Anchor
Harmonizer: Perc.: -1200 cents
88
ff

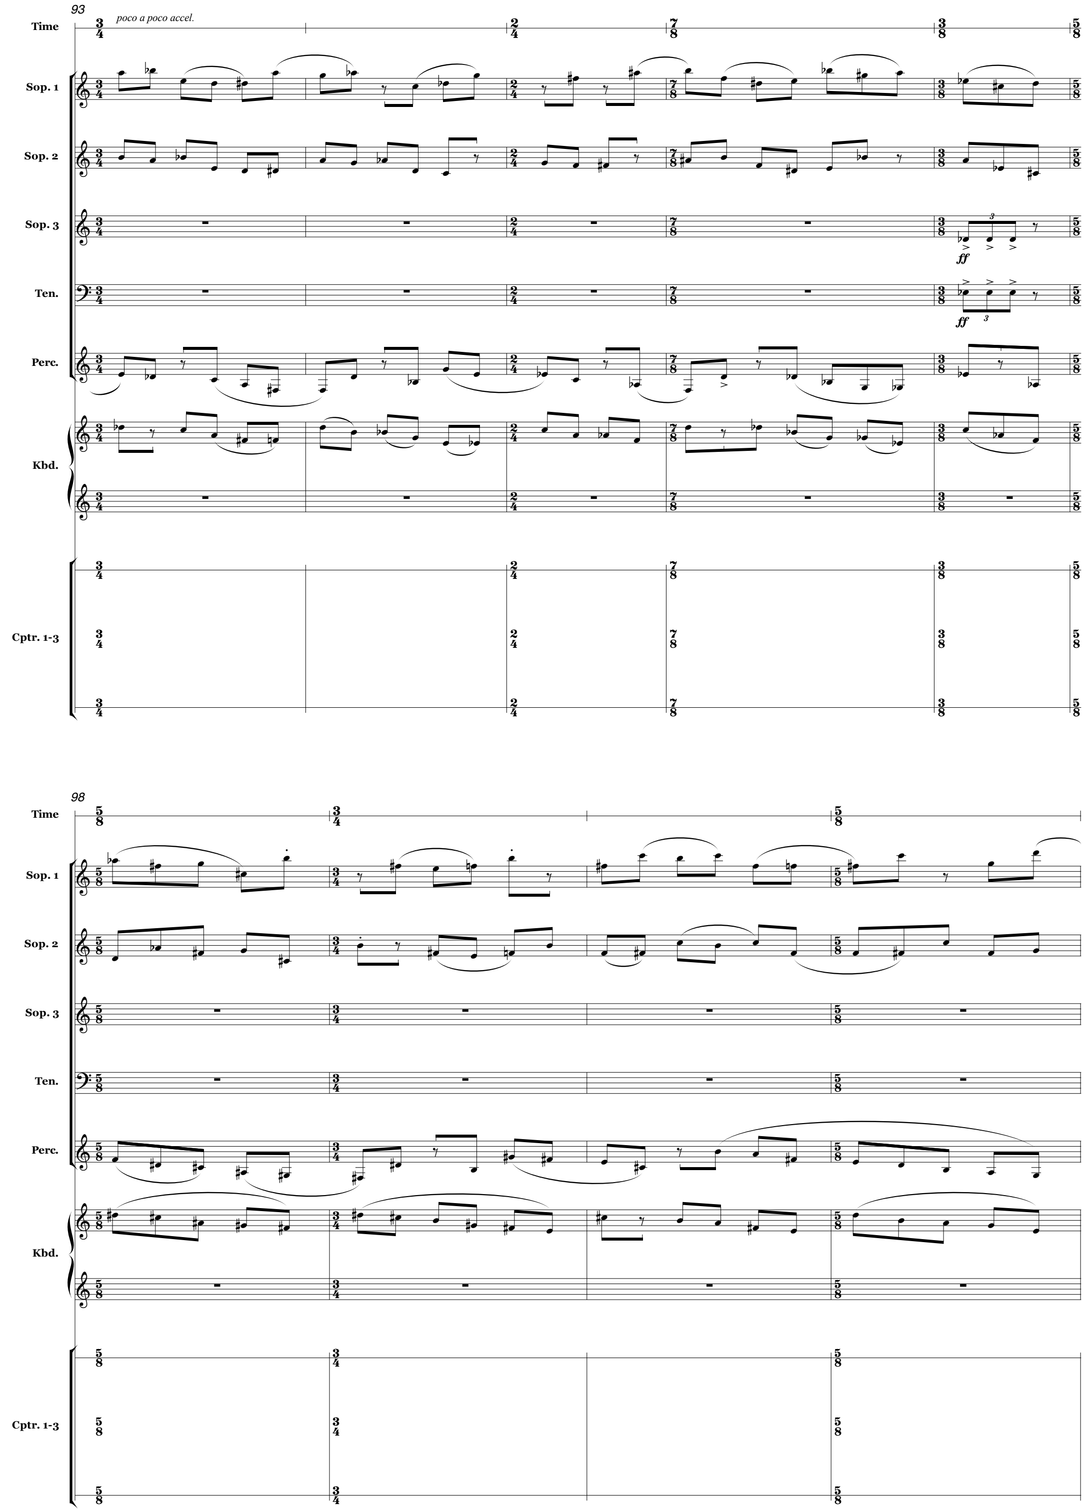
93
Time
poco a poco accel.
Sop. 1
Sop. 2
Sop. 3
ff
Ten.
ff
Perc.
Kbd.
Cptr. 1-3
98
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3

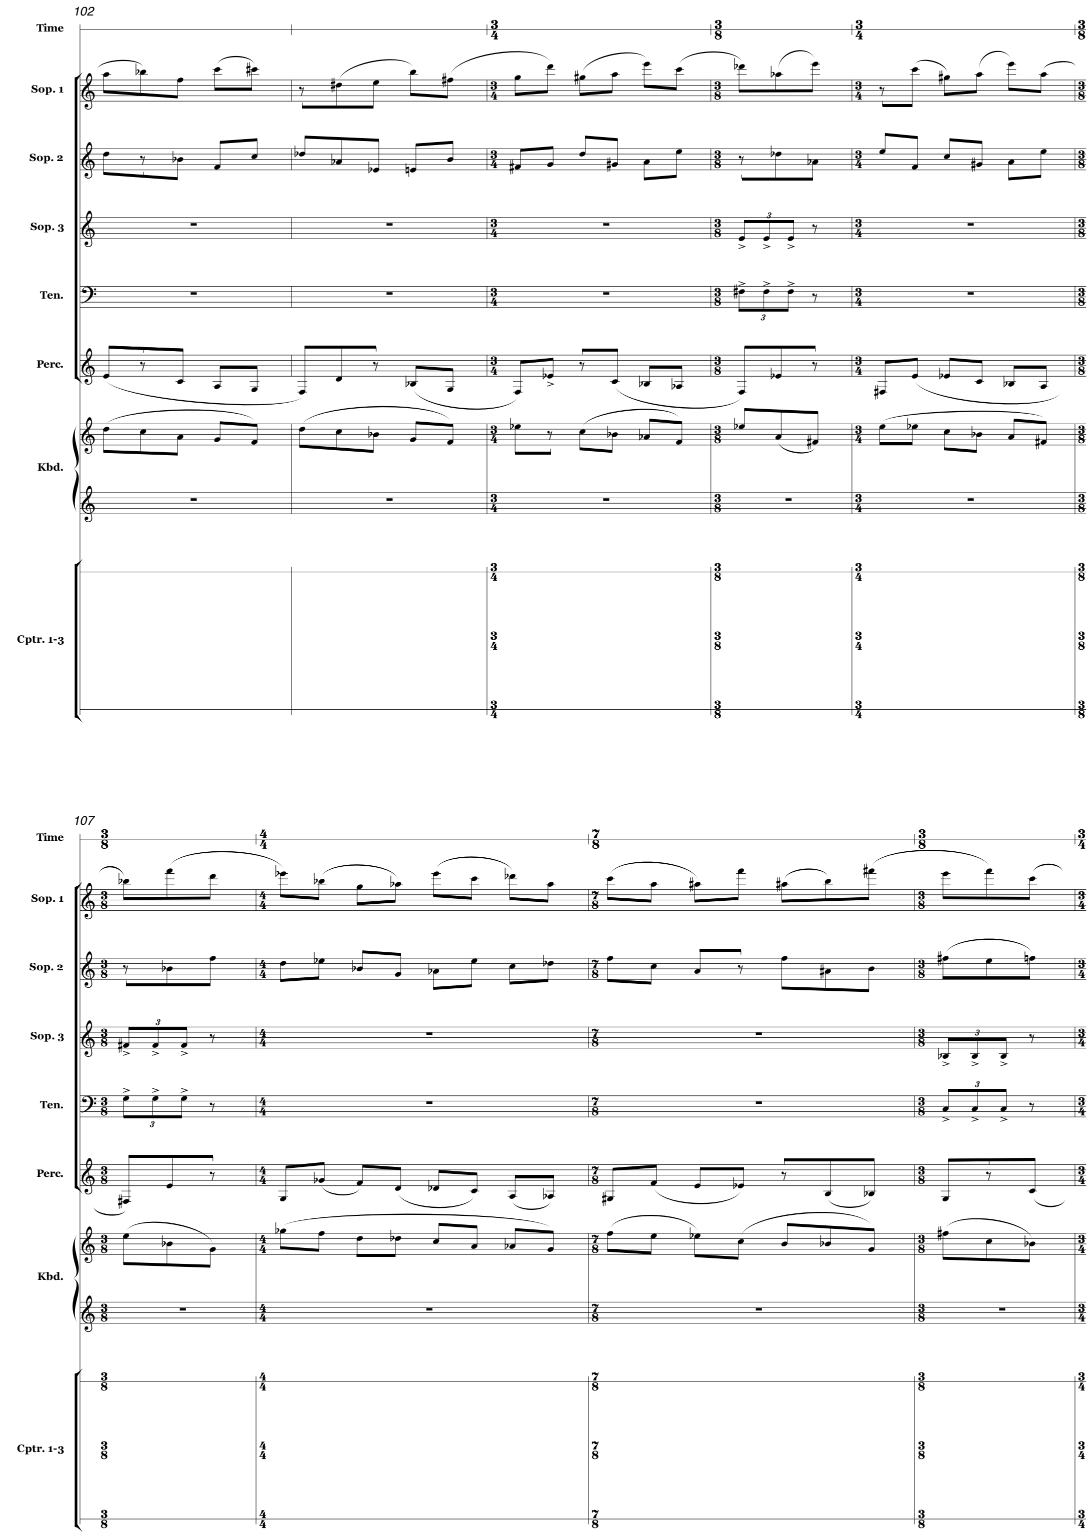
102
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
107
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3


111
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
116
rall.
mf
ff

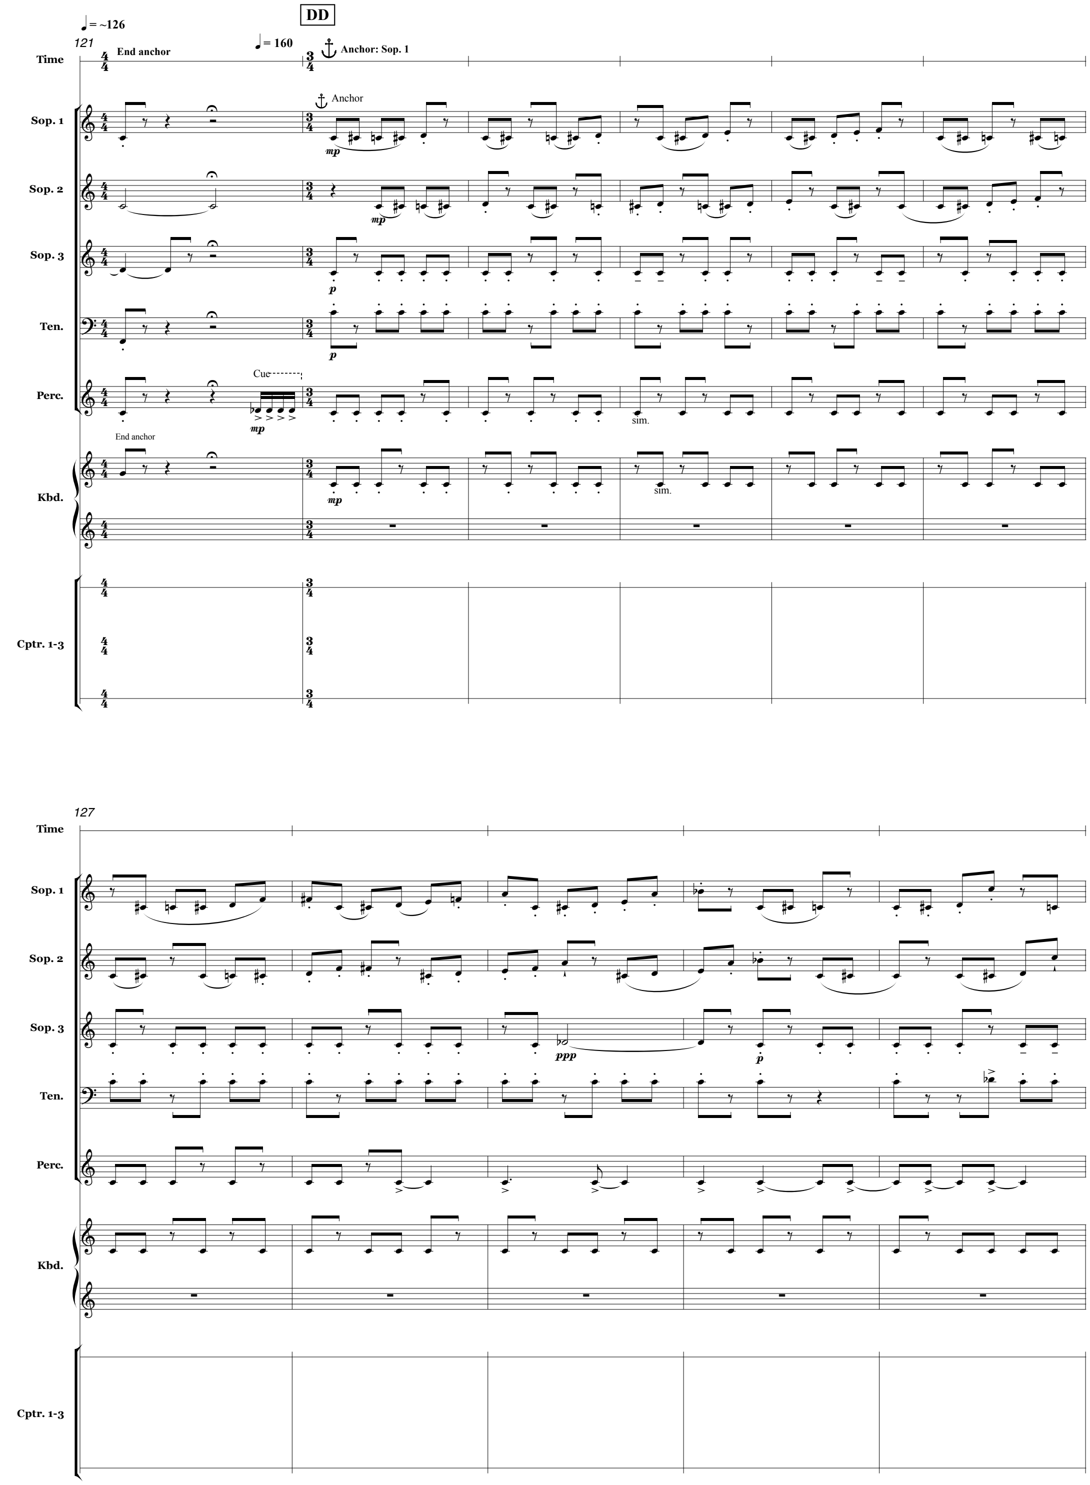
♩= ~126
121
End anchor
♩= 160
DD
Anchor: Sop. 1
Time
Sop. 1
Anchor
mp
Sop. 2
mp
Sop. 3
p
Ten.
p
Perc.
Cue
mp
sim.
End anchor
Kbd.
mp
sim.
Cptr. 1-3
127
Time
Sop. 1
Sop. 2
Sop. 3
ppp
p
Ten.
Perc.
Kbd.
Cptr. 1-3

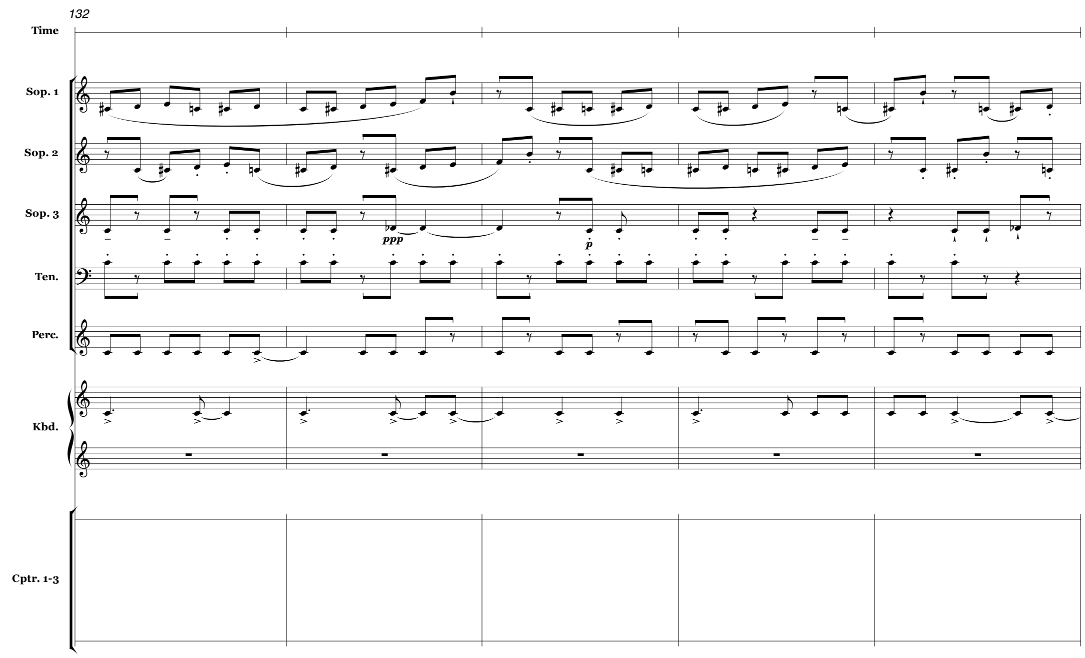
132
Time
Sop. 1
Sop. 2
Sop. 3
ppp
p
Ten.
Perc.
Kbd.
Cptr. 1-3

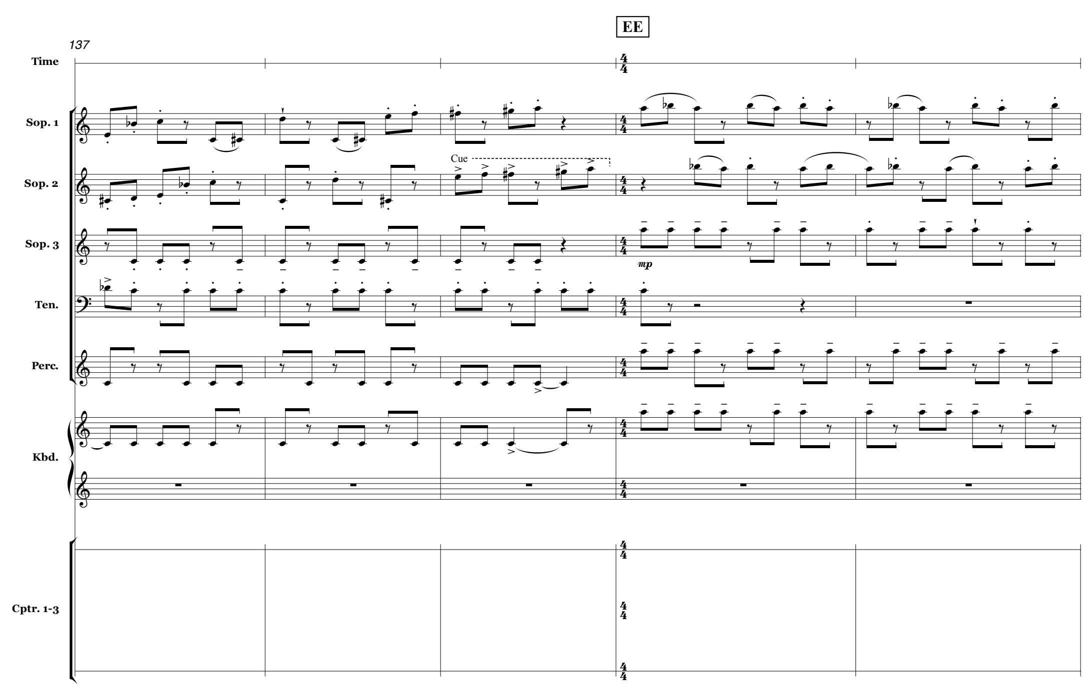
EE
137
Time
Sop. 1
Sop. 2
Cue
Sop. 3
mp
Ten.
Perc.
Kbd.
Cptr. 1-3


142
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Solo
f
9
Perc.
sim.
Kbd.
sim.
Cptr. 1-3

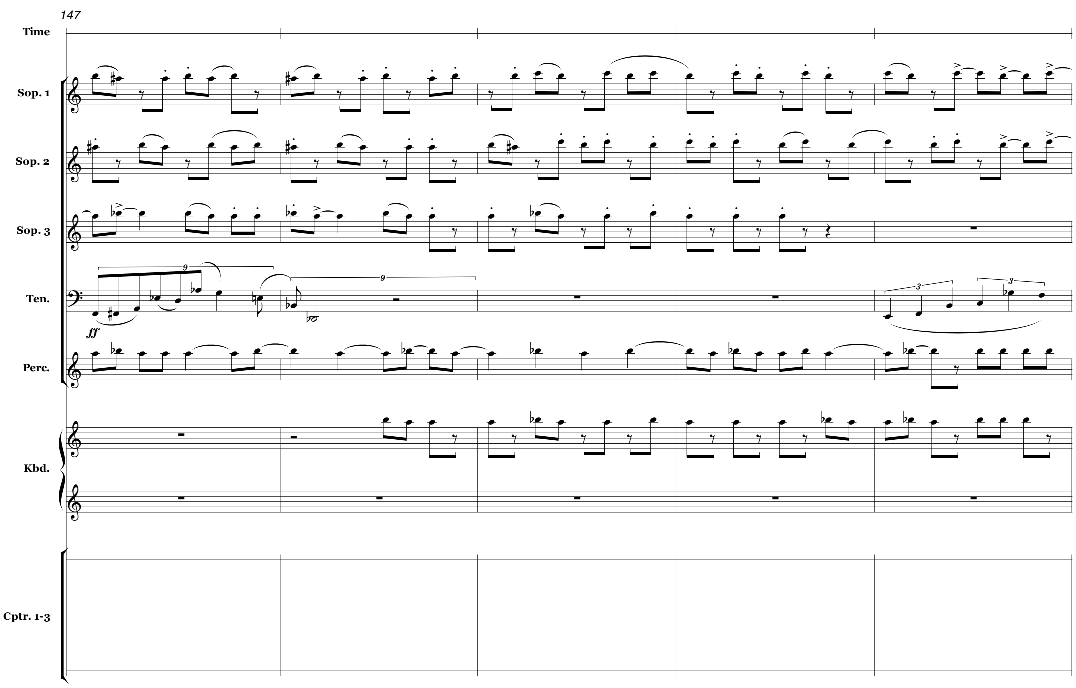
147
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
9
ff
9
3
3
Perc.
Kbd.
Cptr. 1-3

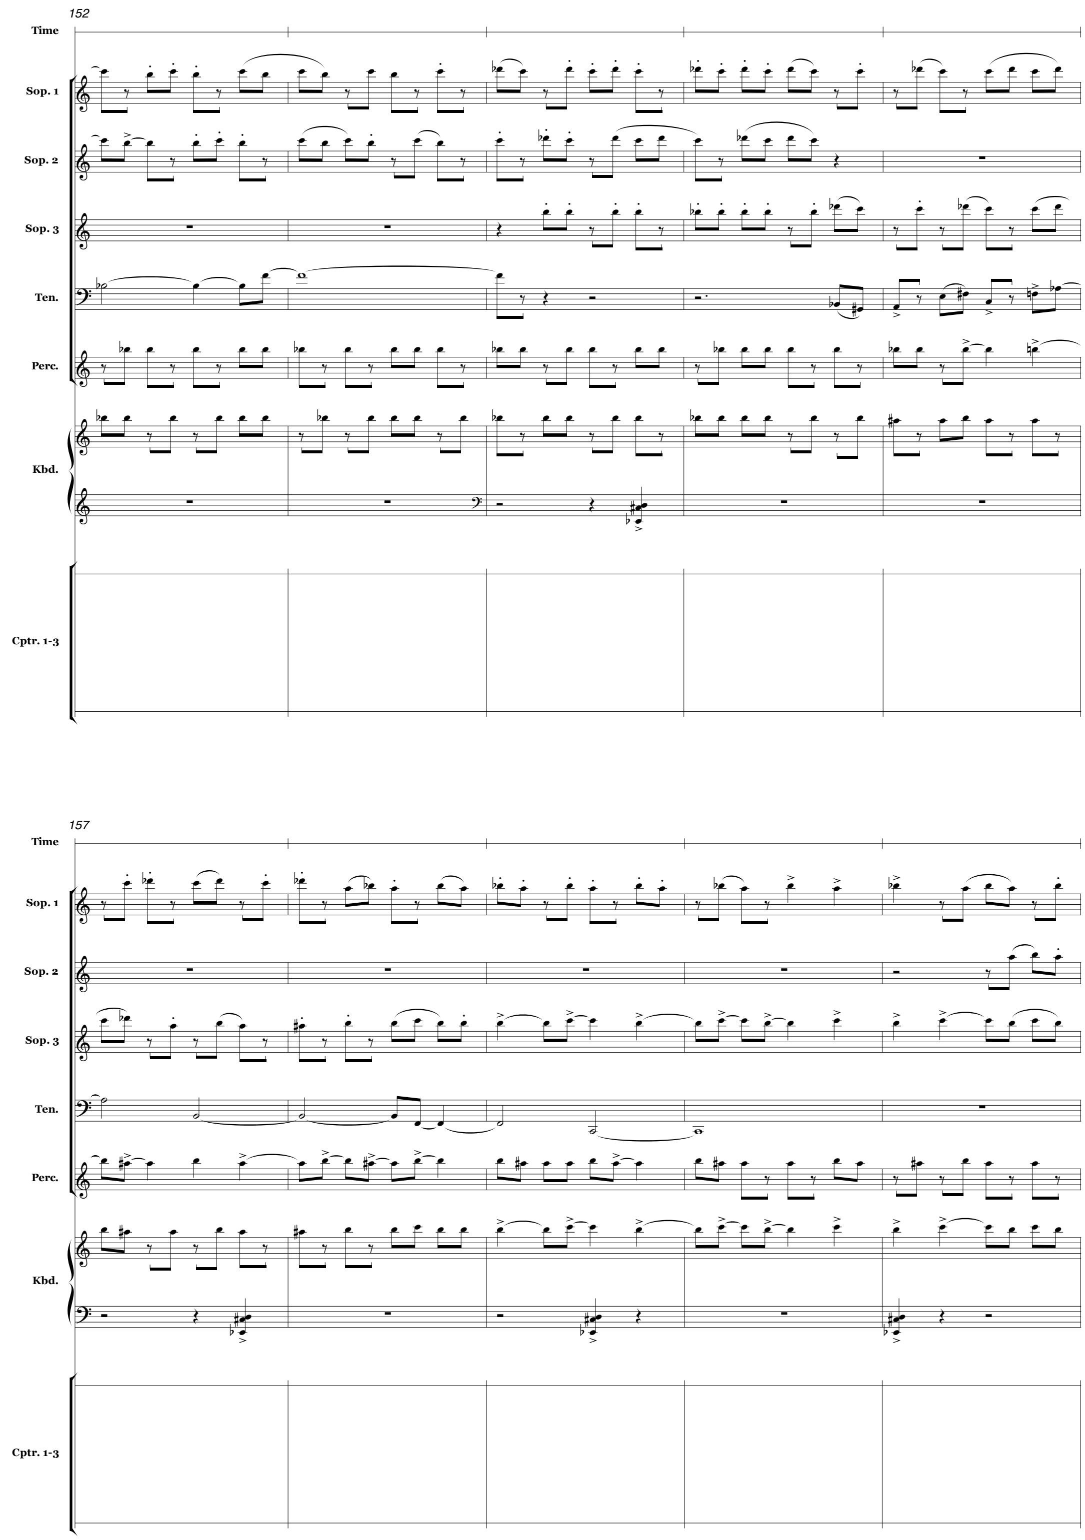
152
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
157
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3

162
End anchor
Time
Sop. 1
End anchor
Sop. 2
Sop. 3
Ten.
End solo
Perc.
Kbd.
fff
Cptr. 1-3

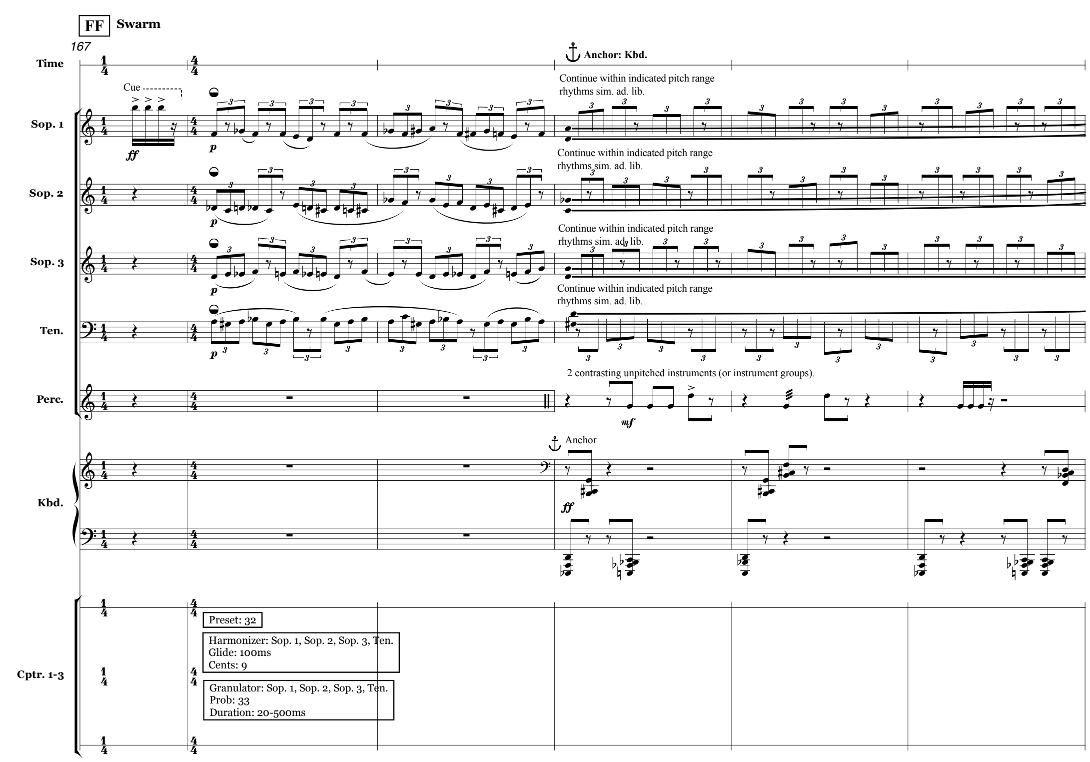

FF
Swarm
167
Time
Anchor: Kbd.
Sop. 1
Cue
ff
p
Continue within indicated pitch range
rhythms sim. ad. lib.
Sop. 2
Continue within indicated pitch range
rhythms sim. ad. lib.
Sop. 3
Continue within indicated pitch range
rhythms sim. ad. lib.
Ten.
Continue within indicated pitch range
rhythms sim. ad. lib.
Perc.
2 contrasting unpitched instruments (or instrument groups).
mf
Kbd.
Anchor
ff
Cptr. 1-3
Preset: 32
Harmonizer: Sop. 1, Sop. 2, Sop. 3, Ten.
Glide: 100ms
Cents: 9
Granulator: Sop. 1, Sop. 2, Sop. 3, Ten.
Prob: 33
Duration: 20-500ms

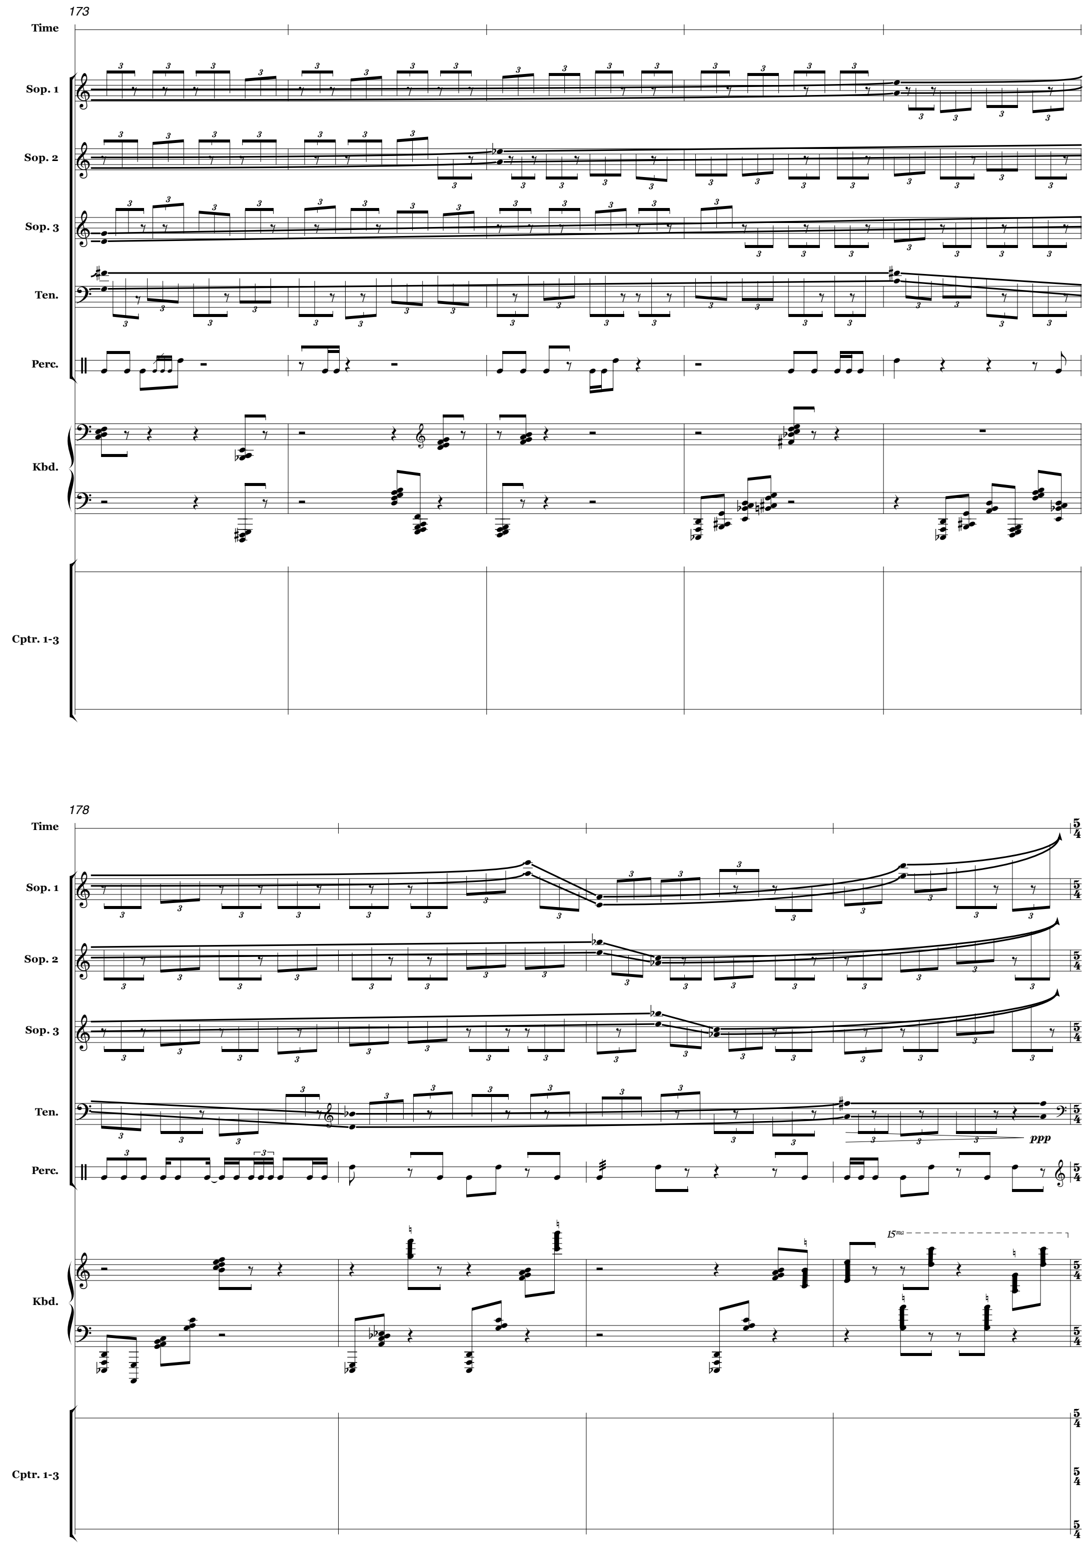
173
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
178
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
ppp
Perc.
15ma
Kbd.
Cptr. 1-3

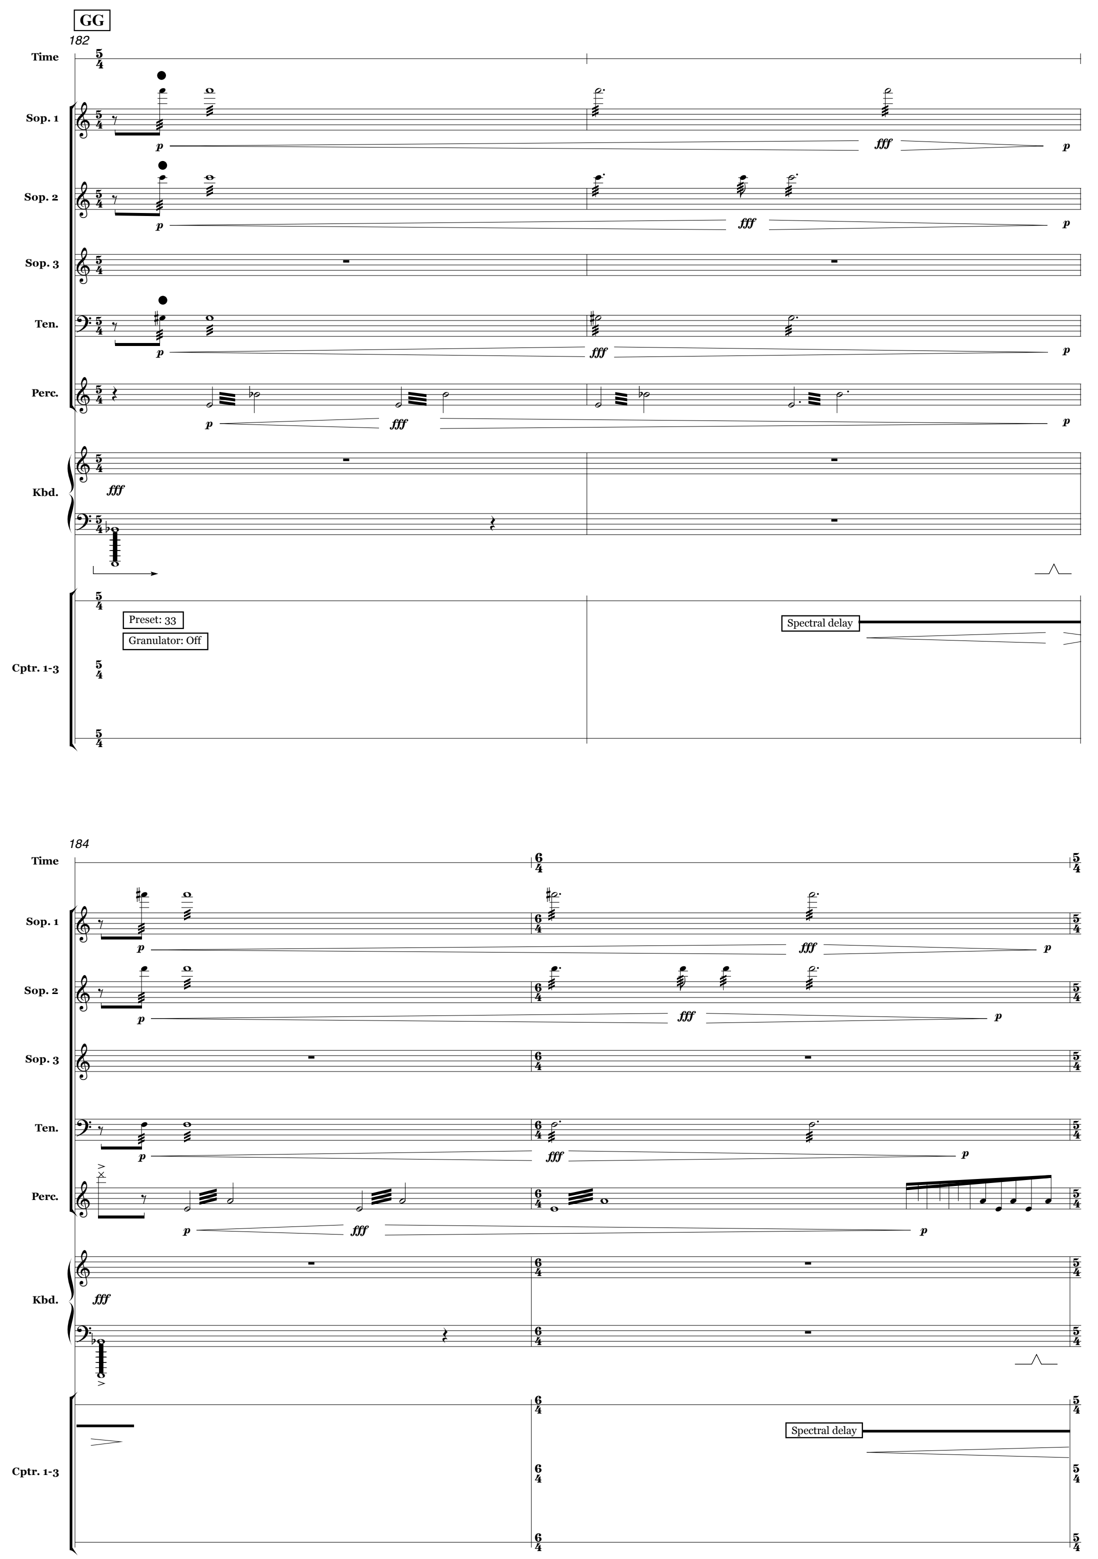
GG
182
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
Preset: 33
Granulator: Off
Spectral delay
184
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
Spectral delay

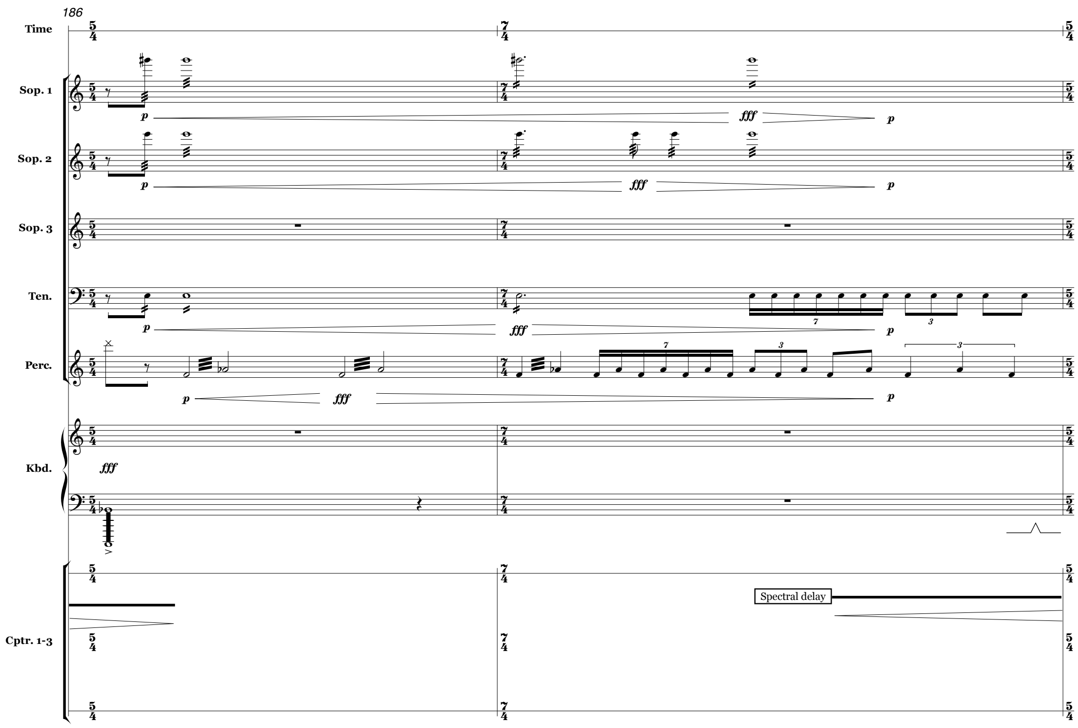
186
Time
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
Spectral delay

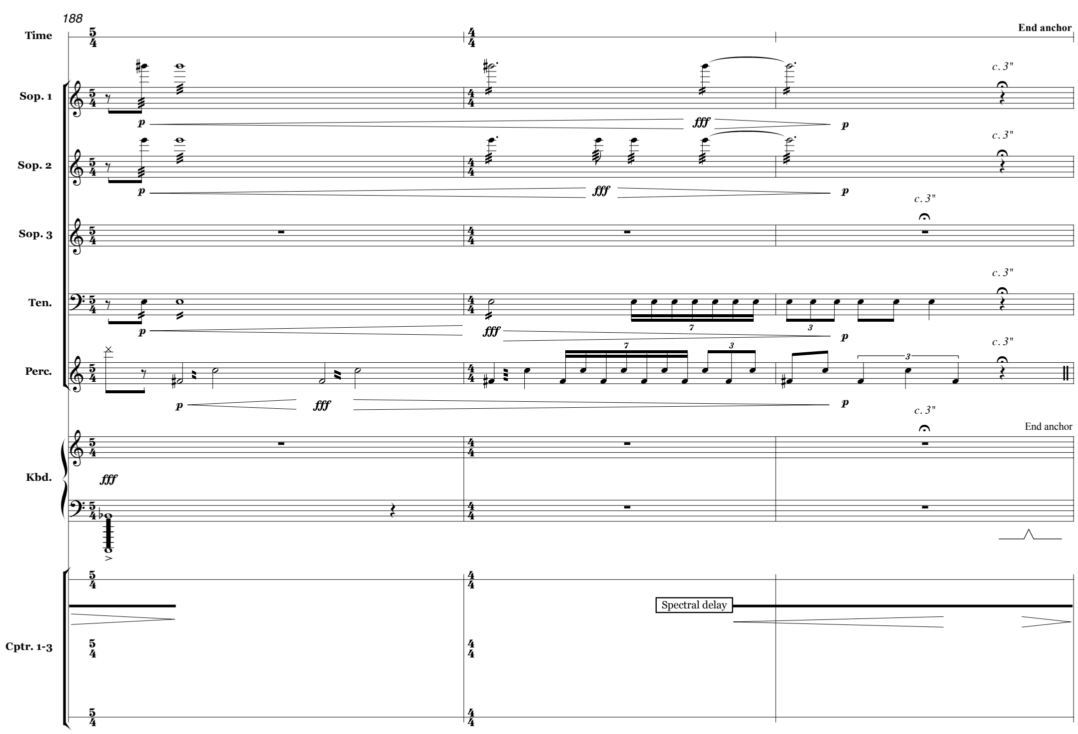
188
Time
End anchor
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
c. 3"
End anchor
Spectral delay

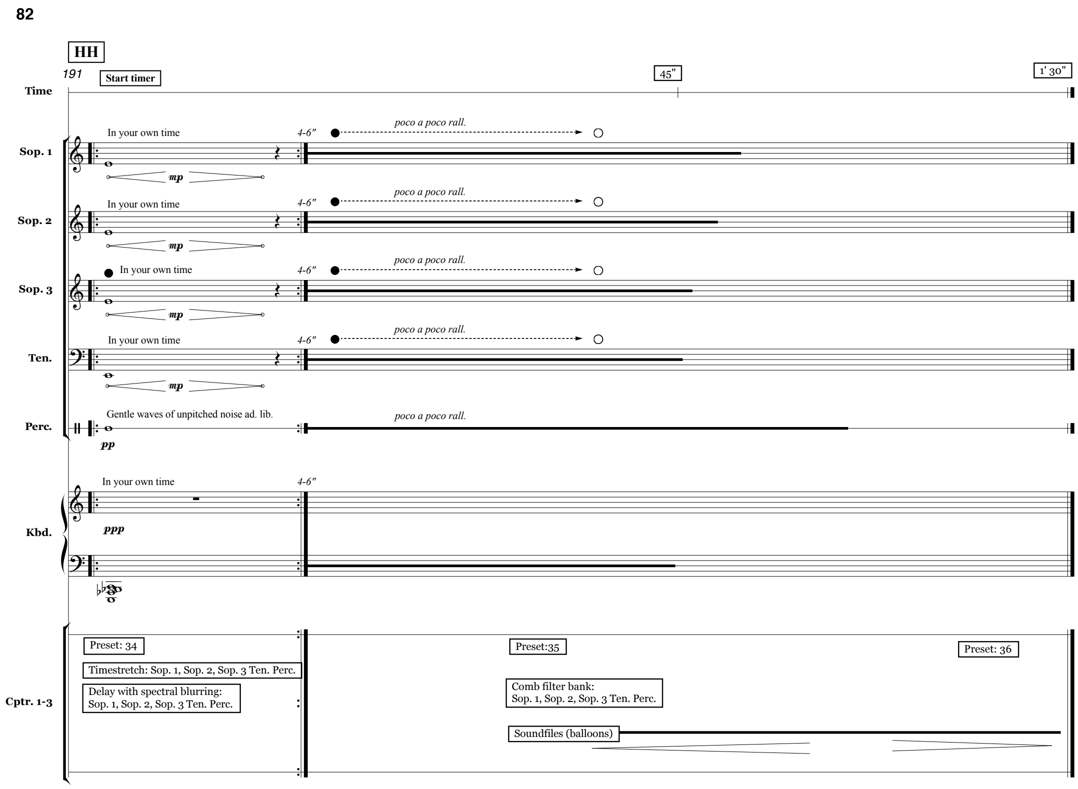
82
HH
191
Start timer
Time
45"
1' 30"
Sop. 1
Sop. 2
Sop. 3
Ten.
Perc.
Kbd.
Cptr. 1-3
In your own time
4-6"
poco a poco rall.
mp
Gentle waves of unpitched noise ad. lib.
pp
ppp
Preset: 34
Timestretch: Sop. 1, Sop. 2, Sop. 3 Ten. Perc.
Delay with spectral blurring:
Sop. 1, Sop. 2, Sop. 3 Ten. Perc.
Preset:35
Comb filter bank:
Sop. 1, Sop. 2, Sop. 3 Ten. Perc.
Soundfiles (balloons)
Preset: 36

APPENDIX B: TRIPLE INTERVAL CYCLE DIAGRAMS

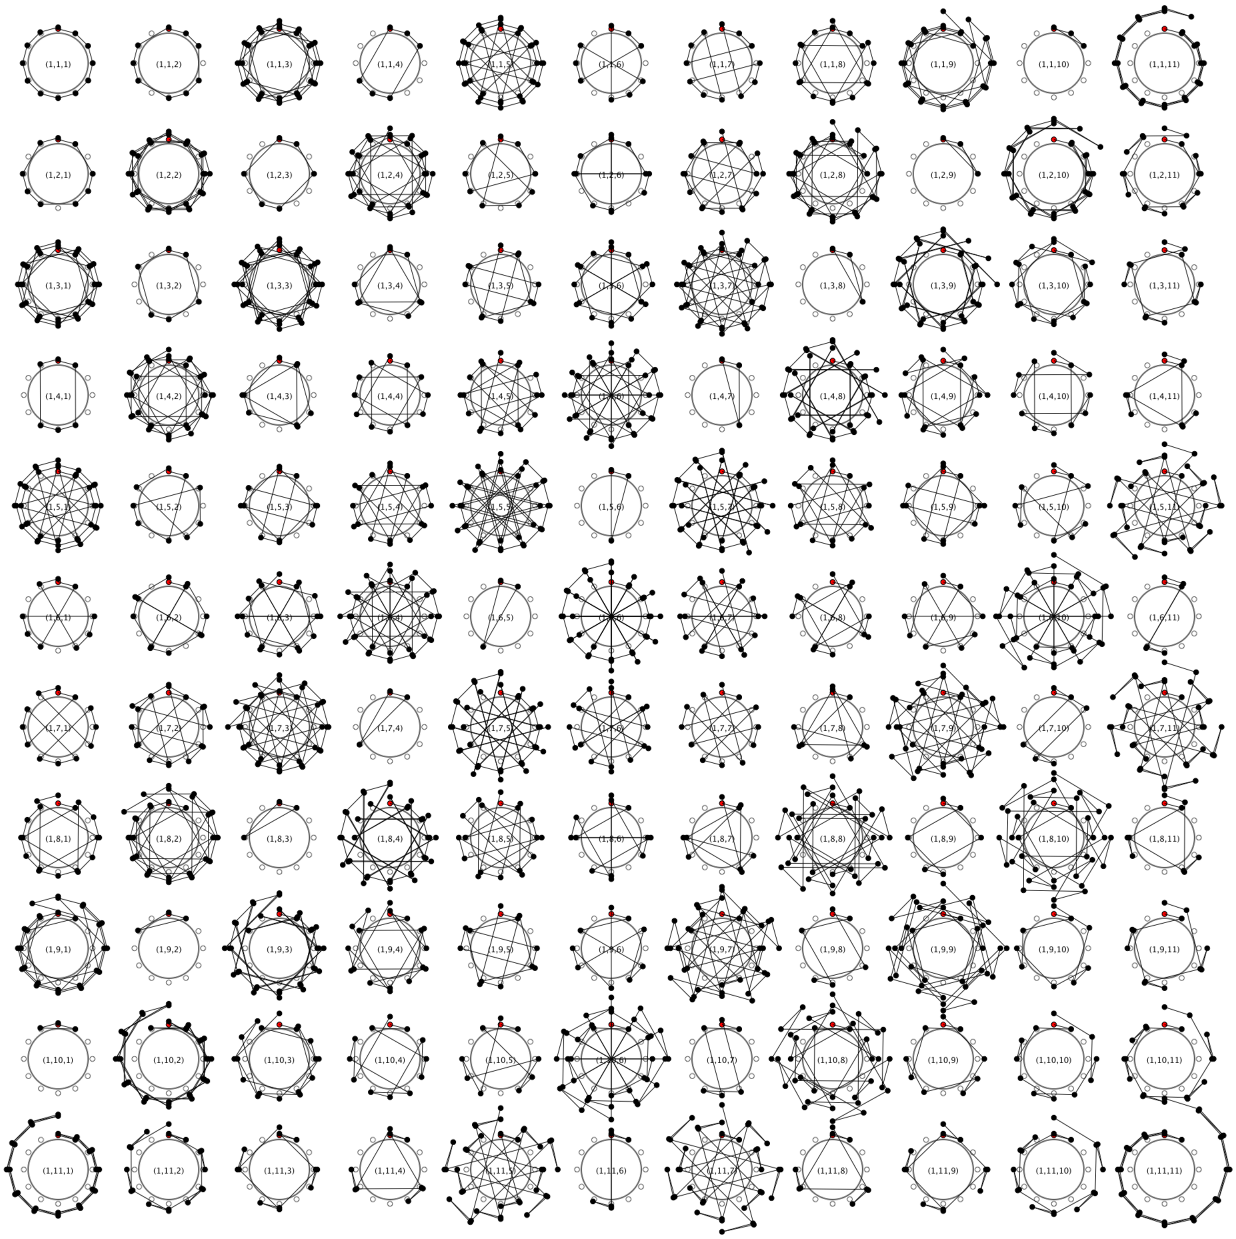

Figure B-1 (1,1,1)-cycle to (1,e,e)-cycle

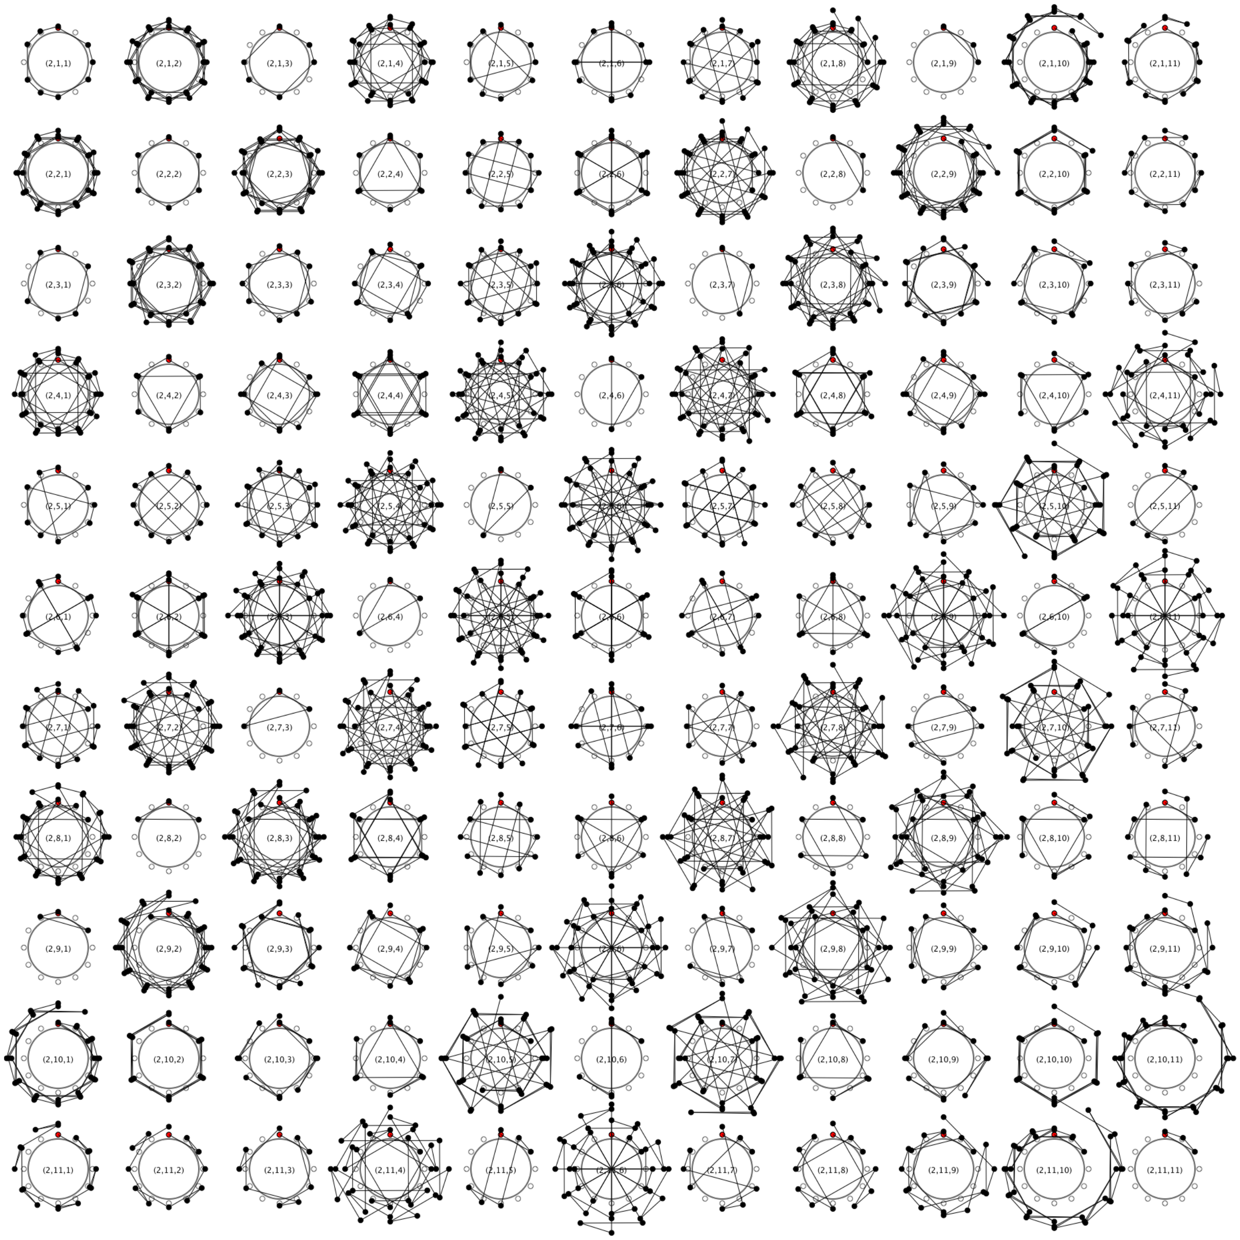

Figure B-2 (2,1,1)-cycle to (2,e,e)-cycle

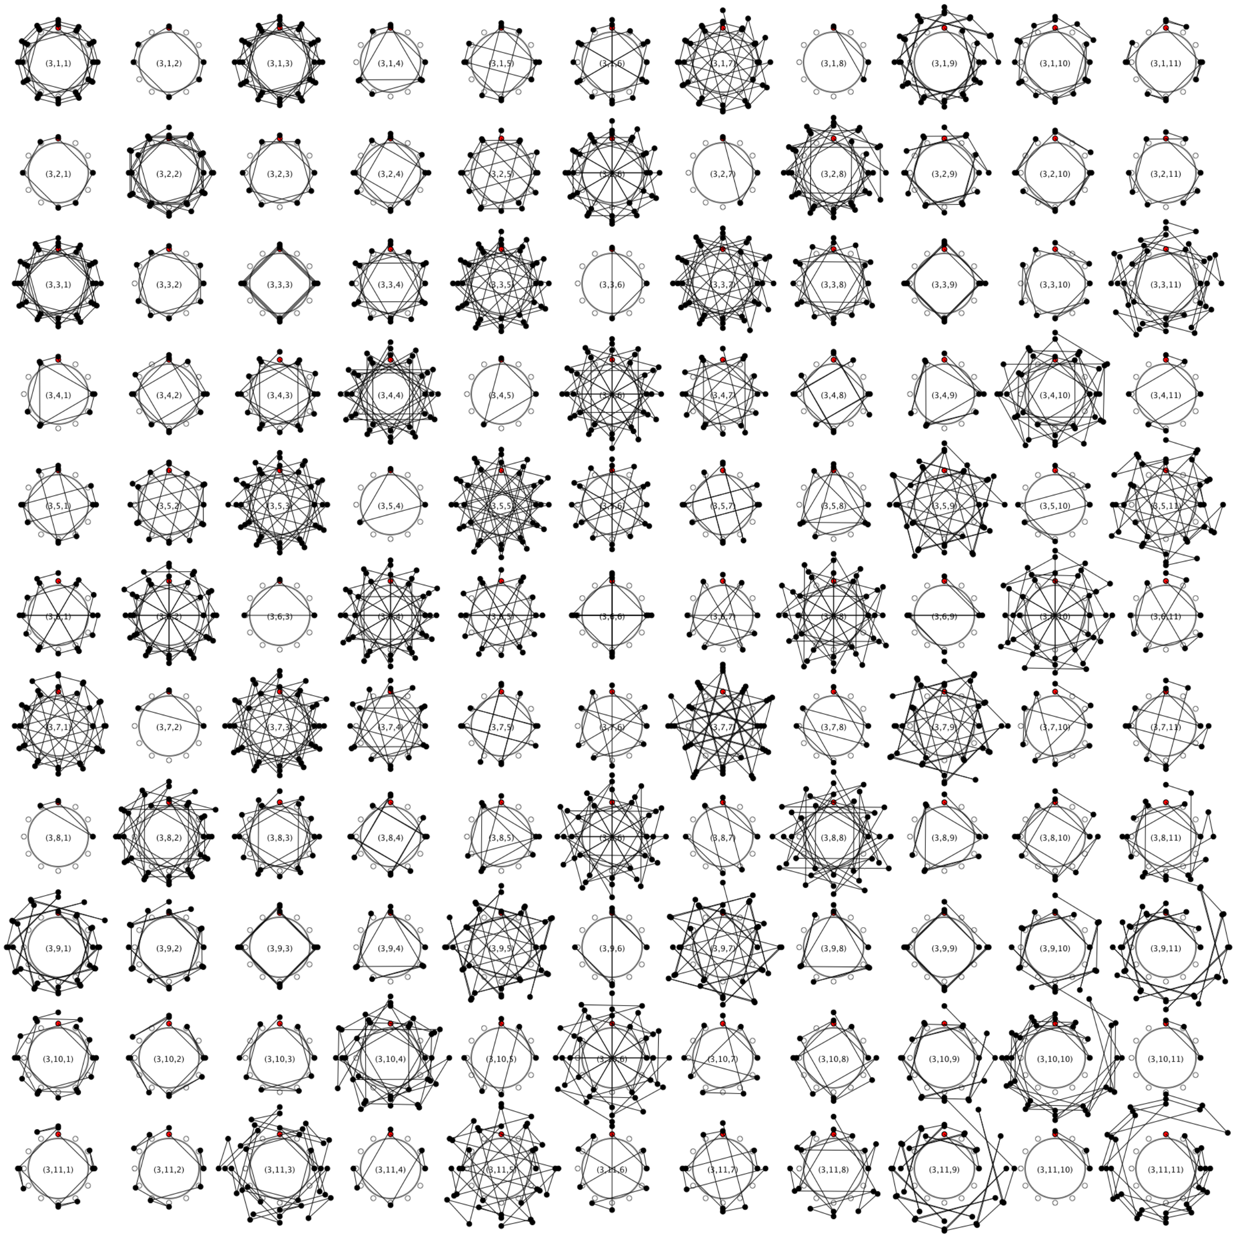

Figure B-3 (3,1,1)-cycle to (3,e,e)-cycle

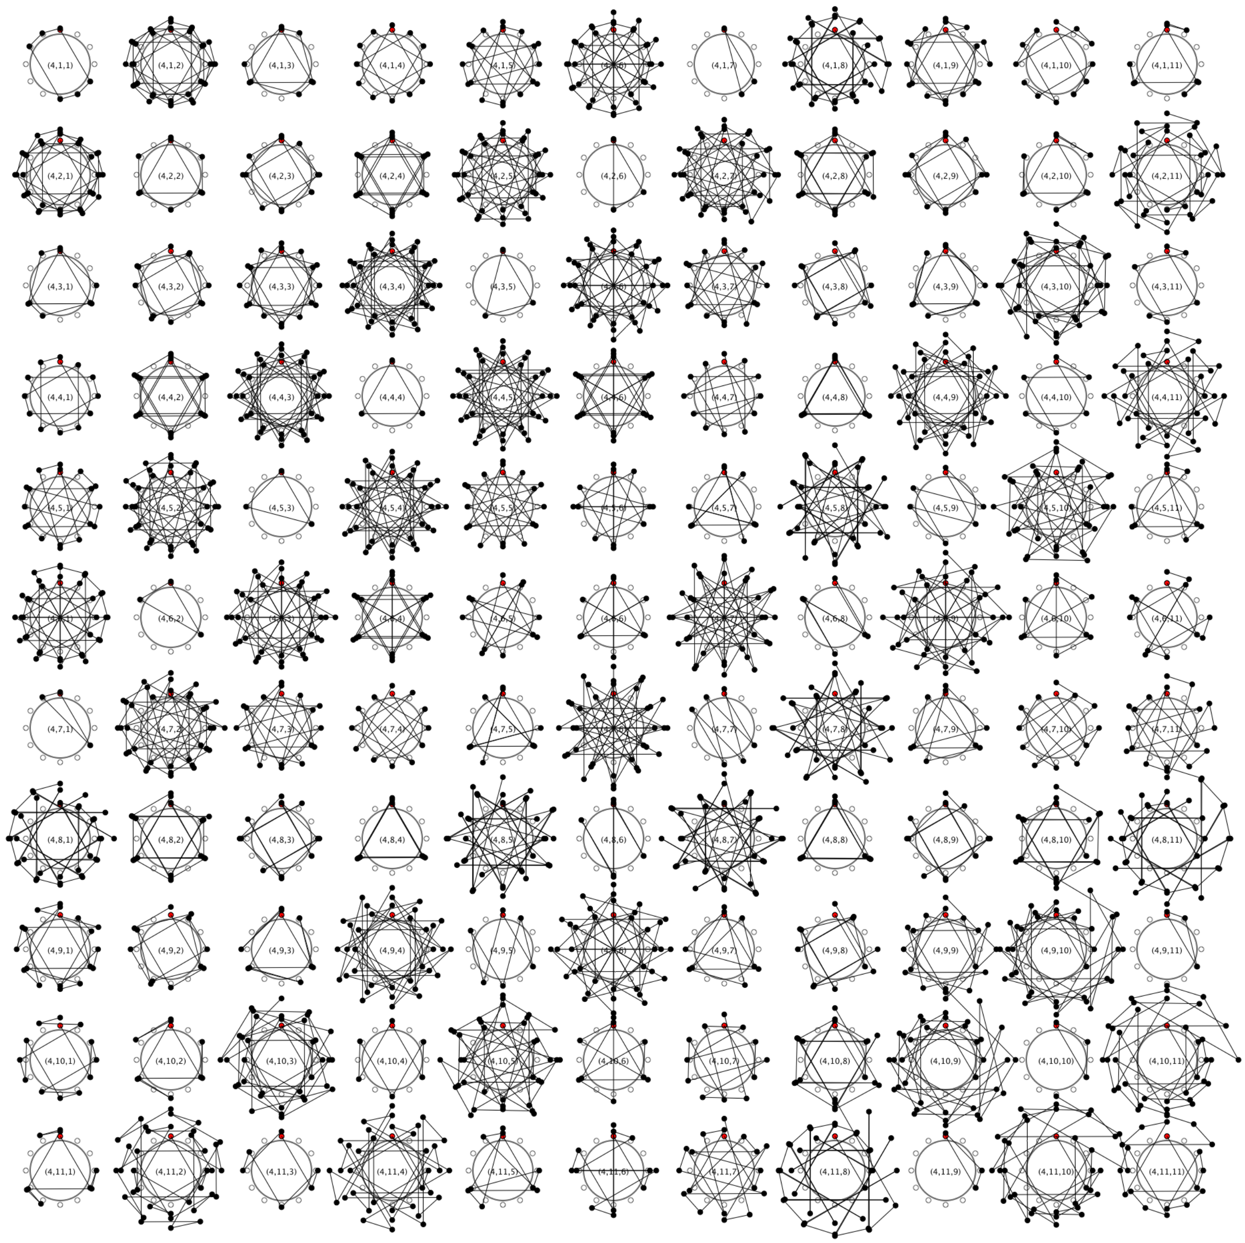

Figure B-4 (4,1,1)-cycle to (4,e,e)-cycle

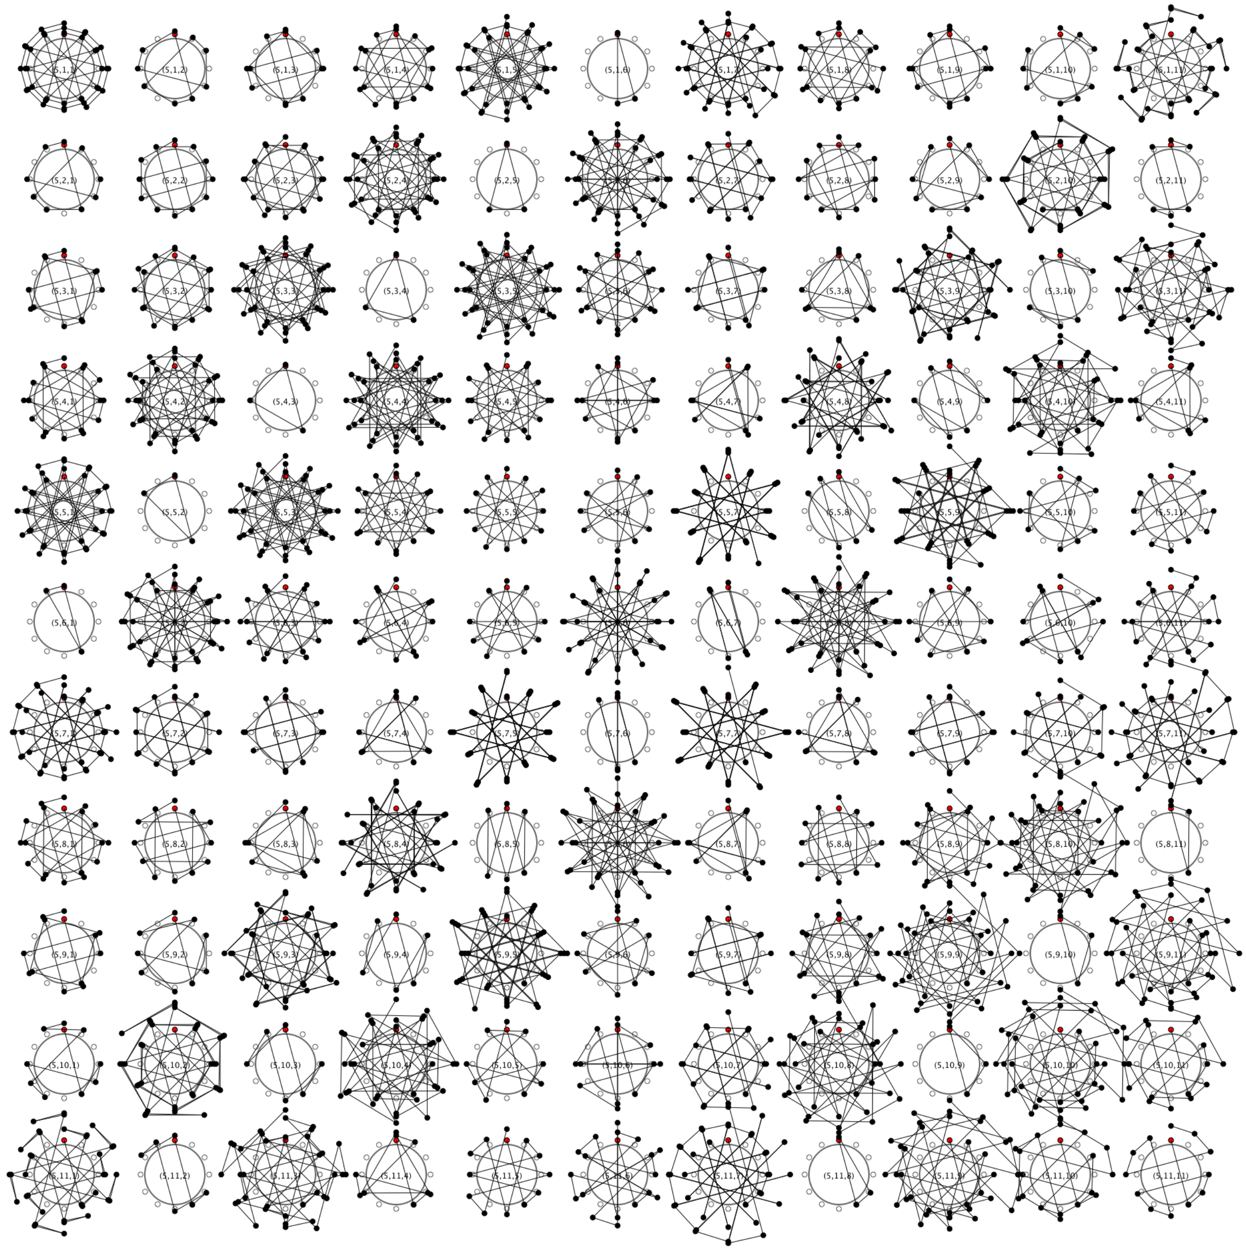

Figure B-5 (5,1,1)-cycle to (5,e,e)-cycle

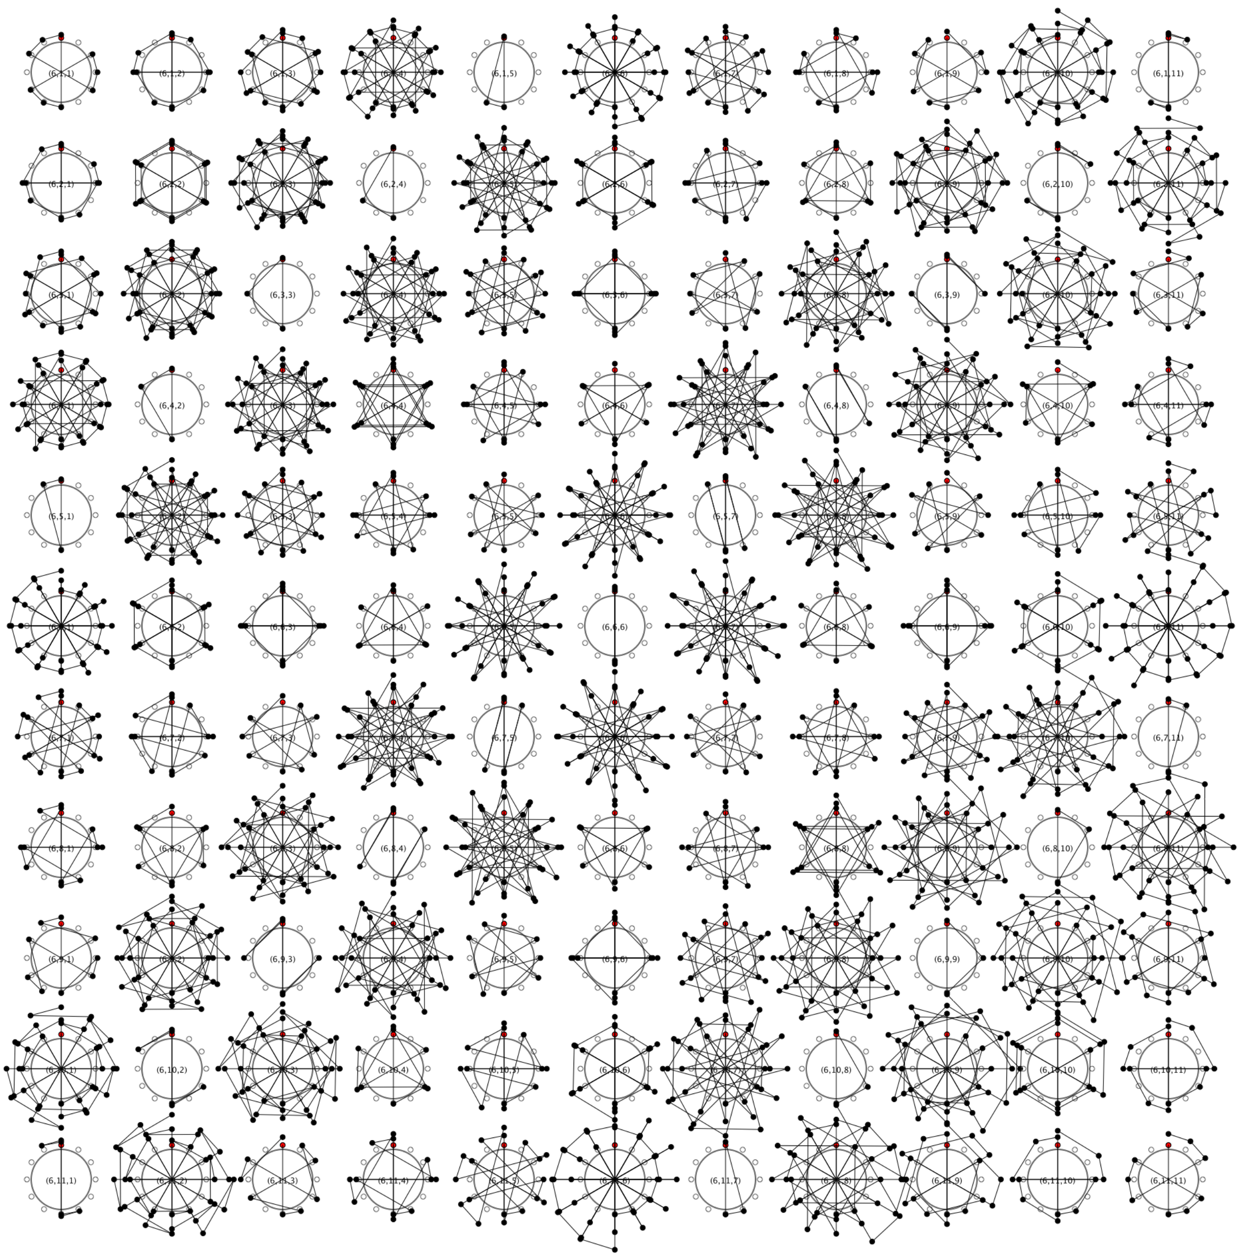

Figure B-6 (6,1,1)-cycle to (6,e,e)-cycle

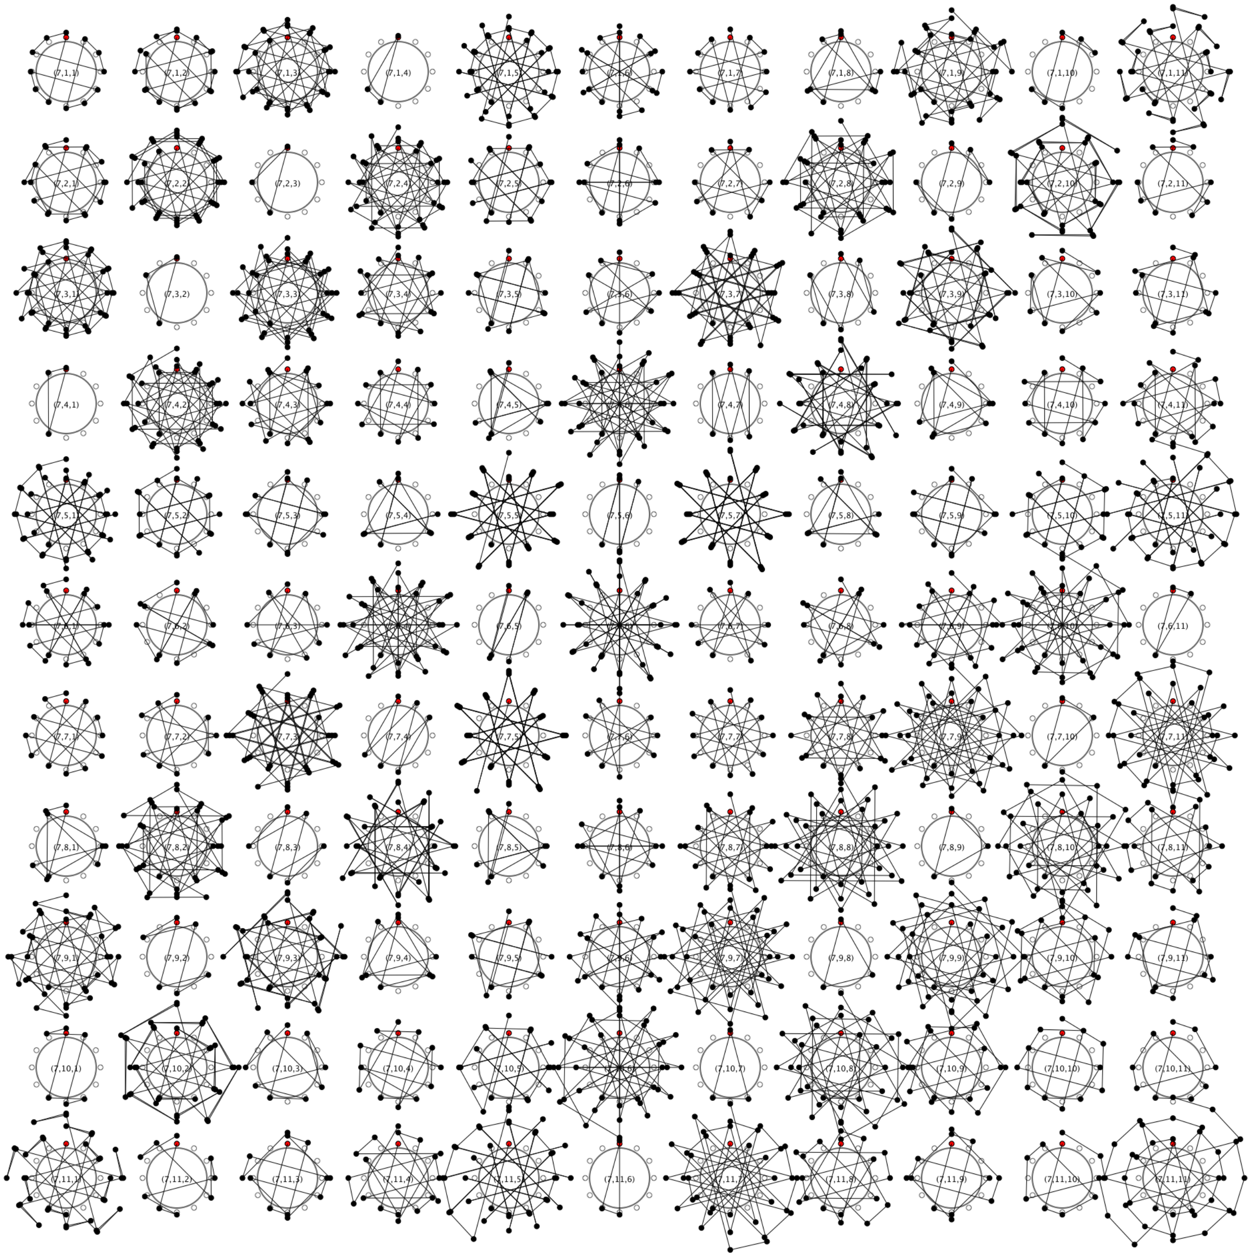

Figure B-7 (7,1,1)-cycle to (7,e,e)-cycle

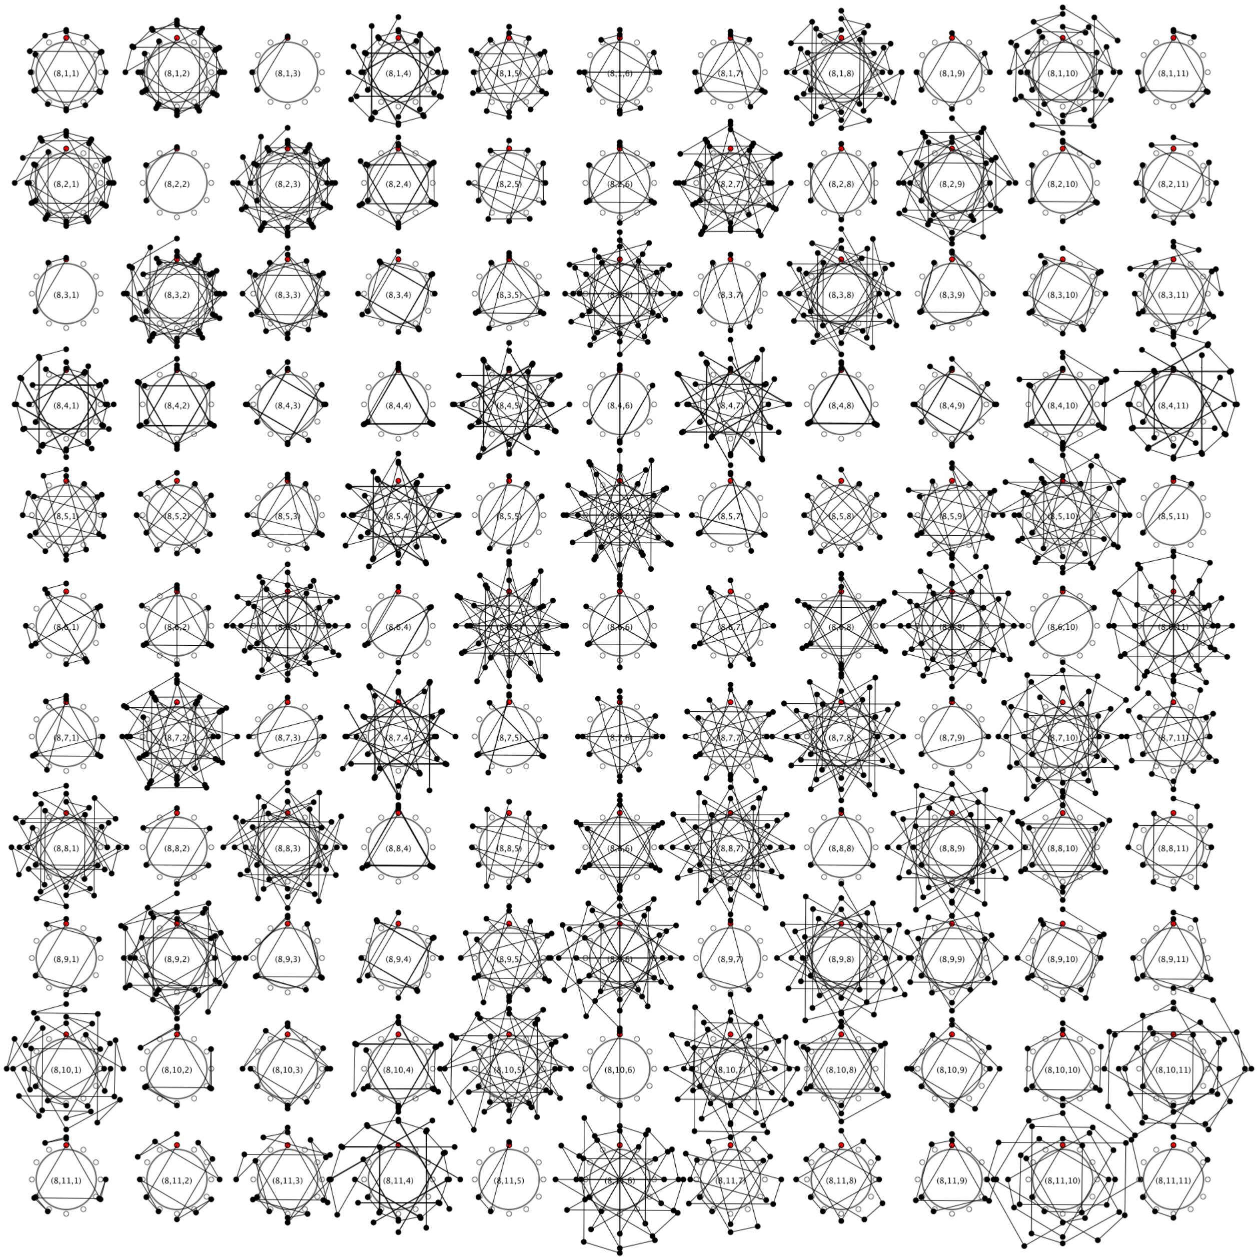

Figure B-8 (8,1,1)-cycle to (8,e,e)-cycle

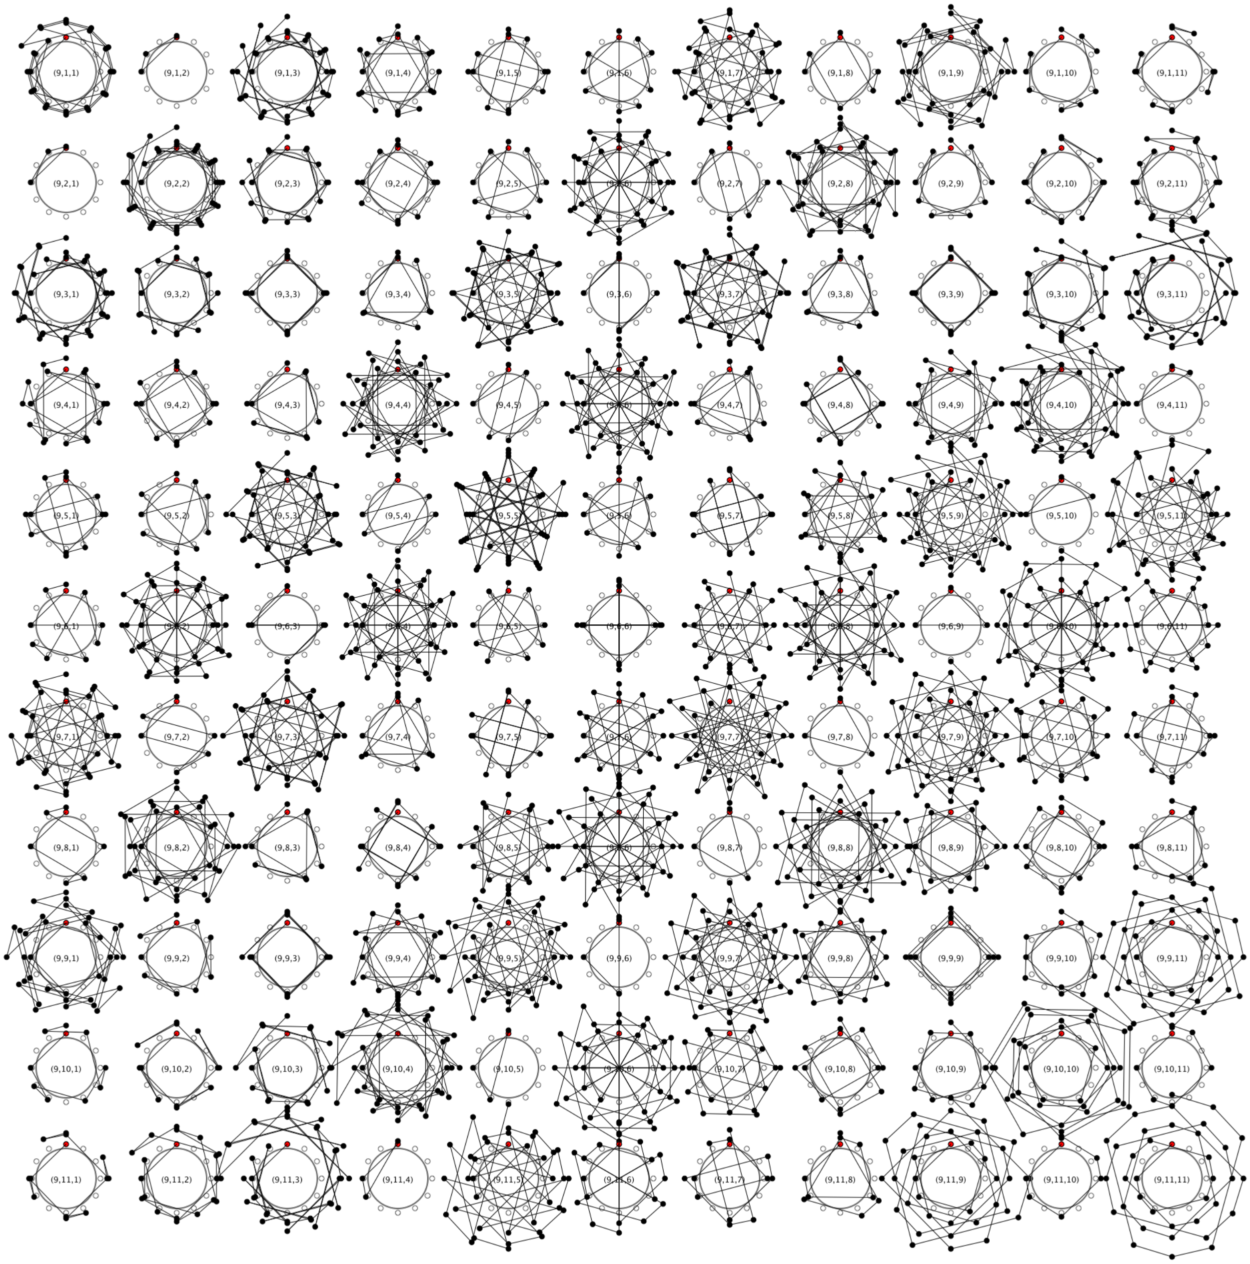

Figure B-9 (9,1,1)-cycle to (9,e,e)-cycle

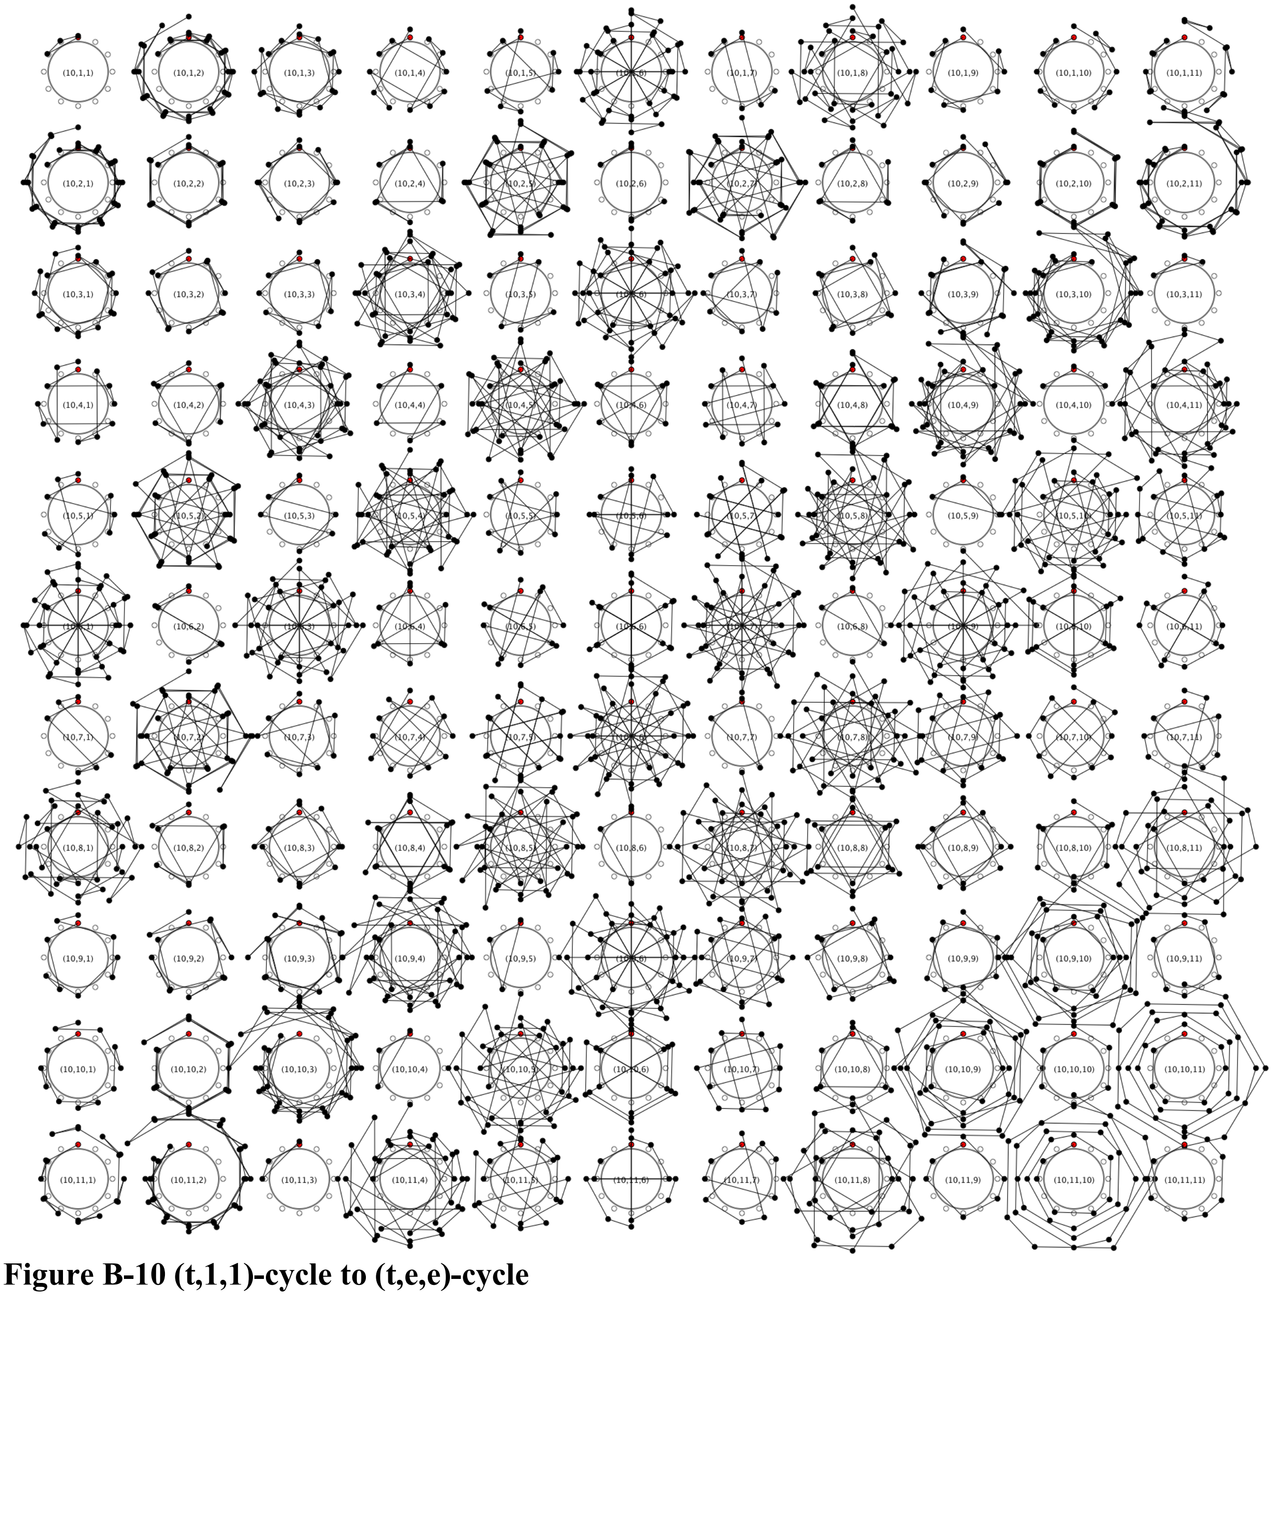

Figure B-10 (t,1,1)-cycle to (t,e,e)-cycle

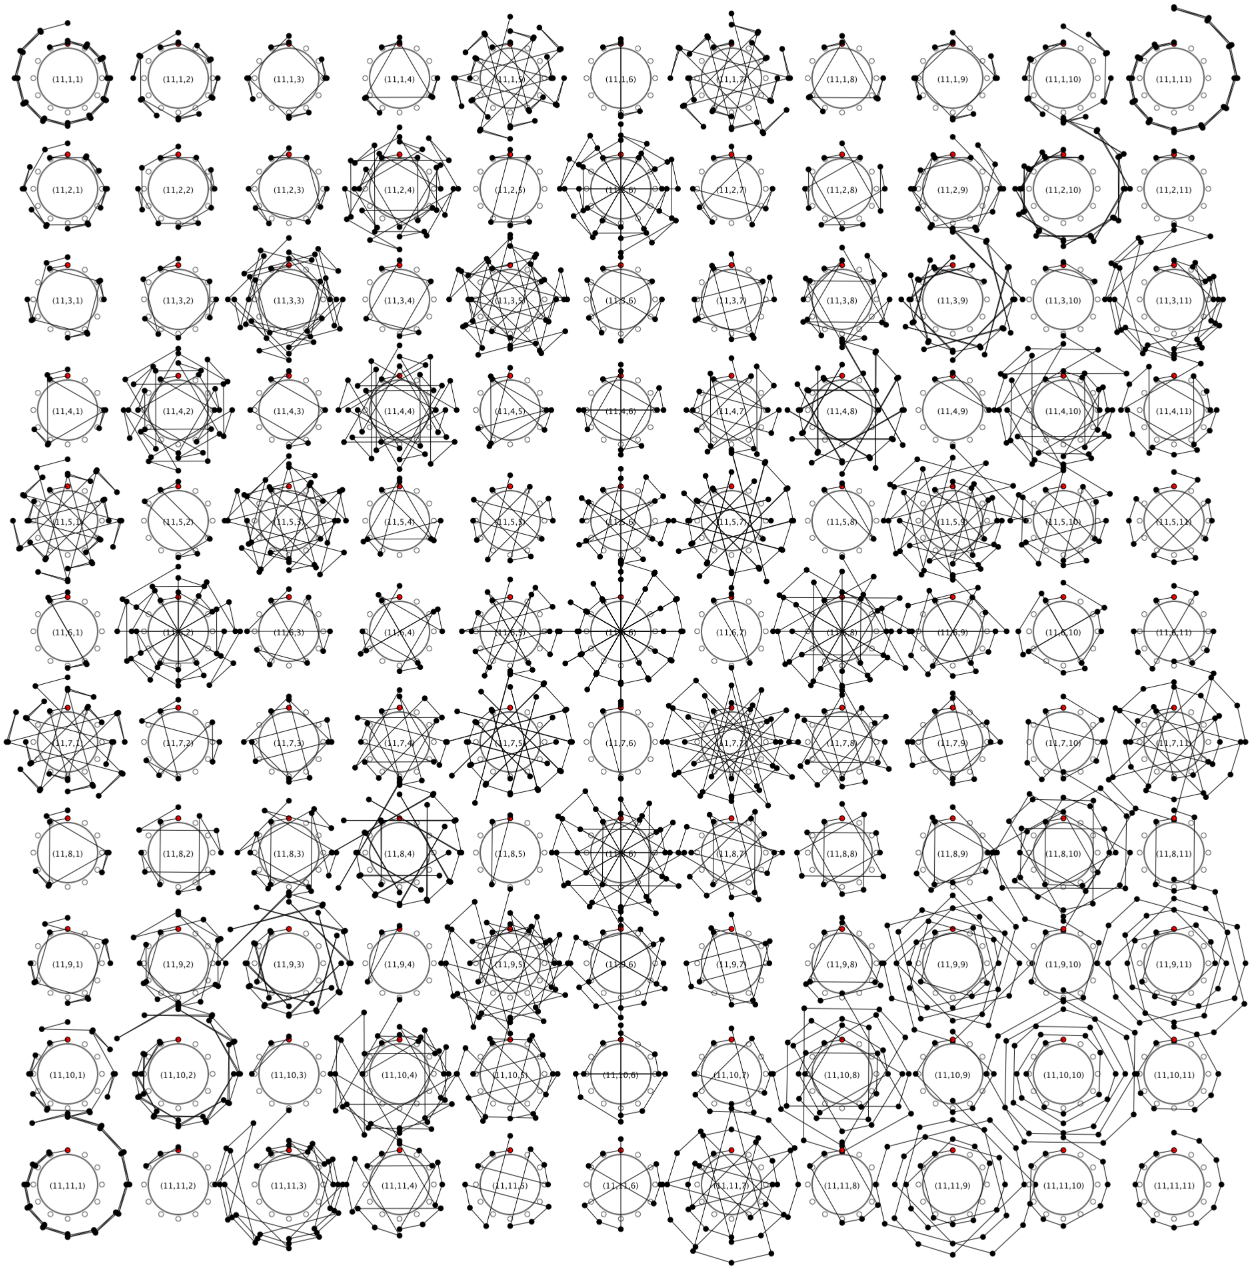

Figure B-11 (e,1,1)-cycle to (e,e,e)-cycle

APPENDIX C: ACCOMPANYING AUDIO AND SOFTWARE (DVD)

Performance recordings of *Through a Window*

The two recordings are stereo mixes taken from the Eckhardt-Gramatté Hall and the Doolittle Studio during the première performance on February 8, 2018 at the *Forms of Sound Festival* at the University of Calgary. The ensemble consisted of August Murphy, soprano 1 (flute); Edmond Agopian, soprano 2 (violin); Chinley Hinacay, soprano 3 (soprano saxophone); Ethan Mitchell, tenor (cello); Tim Borton, percussion; Rachel Kreyner, keyboard (piano); Naithan Bosse, computer 1; Melike Ceylan, computer 2; and Abdullah Soydan, computer 3.

Performance system

A folder containing all of the original software modules and documentation required to perform *Through a Window*. Note that the Max performance software relies on the *UBCToolbox* (Hamel and Pritchard), the *vb.stretch~* external (Böhm), the *bc.yinstats* external (Levy), and the *yin~* external (Schnell).

Sample banks

A library of soundfiles used in the performance of *Through a Window*.