A PORTFOLIO OF ELECTROACOUSTIC AND LIVE ELECTRONIC COMPOSITIONS WITH RESEARCH INTO STOCHASTIC SYNTHESIS

by

ERIC WALTER REGINALD BUMSTEAD

A thesis submitted to the
University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

Department of Music
School of Languages,
Cultures,
Art History and Music
College of Arts and Law
University of Birmingham
January 2011

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

ABSTRACT

This thesis includes a portfolio of electroacoustic and live electronic compositions consisting of multichannel fixed media works and pieces that include instruments, real time computer processing and electronic interfaces. It also details research into stochastic synthesis and its application in music composition. Work carried out in real-time processing of multichannel signals including convolution, spatialisation, as well as decorrelation from one to many channels, sample accurate looping and various sample accurate control systems, is also included in this thesis.

ACKNOWLEDGEMENTS

Firstly I would like to thank my supervisor Dr. Scott Wilson for his detailed explanations and his sense of humour. Also, Prof. Jonty Harrison for his consummate musicianship, Sergio Luque for introducing me to stochastic synthesis and his insight into the algorithms and James Carpenter for his musical and programming advice.

I would also like to thank the Alberta Foundation for the Arts, the University of Birmingham and the Barber Institute for their financial assistance.

Finally, I would like to thank my parents and my wife.

TABLE OF CONTENTS

LIST OF PIECES

LIST OF COMPUTER CODE

CONTENTS OF THE DVDS

DVD 1

C code

stoc.synth~.xcodeproj

Eight Channel Pieces:

|*Kroʊn*|
BlckWnd
Eight Channel Player
Immortal
Lithium
Poly
River

MaxMSP Patches

Audio Rate Pedal
Equal Power Loop
FFT Decorrelation & Convolution
Convolution 8 Ch. to 8 Ch.
Convolution Mono to 8 Ch.
Conv. M. to 8 Ch. Time Stretch
Decorrelation & Convolution
Decorrelation Mono to 8 Ch.
Decorrelation Mono to Stereo
In a Cage
Morse Code Reader
Multi-Delay 8
Multiple Loops
OctPan
RandClick
Spinning Plates
Stochastic Synthesiser
Test Sounds

Stereo Piece & Stereo Mix Down

Fields
Poly Stereo

Videos

In a Cage
Spinning Plates

Commentary & Scores

Commentary
Poly
Spinning Plates

DVD 2

BlckWnd	(10 Channels)
|*Kroʊn*|	(12 Channels)
River	(10 Channels)
Immortal	(9 Channels)
Lithium	(9 Channels)
Lithium	(10 Channels)

Introduction

This commentary gives an overview of my pieces as they fit into the progression of my studies. It also outlines my research into stochastic synthesis.

I came to the University of Birmingham with a background that was focused on live electronics. It had been my interest during my undergraduate degree and was an area I pursued as part of my year at the Banff Centre for the Arts. Some of the artists I met in Banff were building custom electronic interfaces for live performance; this inspired me to take a summer course at Stanford to build up some electronic interface experience. Arriving in Birmingham, I began working with the Lemur (a multi-touch control surface) and the Ethersense (used to pass sensor data into the computer) to create interfaces for live performance. These devices are among the best available but they did not allow me to take advantage of my roots in percussion performance; instead they sent me into the area of composition full of difficult-to-control custom interfaces and static laptop performances. I found inspiration in an unusual place.

The Nintendo Wii had just come out, and I could use its controller to connect my percussion playing with my live electronics composition. The accelerometers and buttons were not, in themselves, a new form of control but here was something inexpensive and wireless that interfaced easily with the computer. Combining my previous work using a theremin as a controller with the wiimote, I had the ingredients for my first piece, *In a Cage*. My work using the two controllers was interesting; however the actual sounds used in the piece were unsatisfactory. I found the control I had over these sounds to be adequate, but they took a lot of time to produce and did not give back much quality in return. |*Kroʊn*| is a good example of the struggle I had with these synthetic textures.

I found the answer to my sound synthesis dilemma when I heard Sergio Luque play one of his stochastic synthesis pieces. Here was something that had the versatility I was looking for. Trying to get organic sounds from FM modulated by noise had proved difficult but possible. With stochastic synthesis, many of the elementary sounds produced matched my aesthetic. I could get the broad bandwidth, thick, edgy, 'full of life' sounds I was looking for with only a few adjustments.

I began in earnest trying to get the Max/MSP language to produce these sounds. Numerous failed attempts to code the algorithm in Max/MSP led me to the Software Development Kit (SDK), which allows composers to write their own Max/MSP objects in the C programming language. I switched my focus to learning how to code in C and to getting sound generators running in Max/MSP using the SDK. By starting with simple generators and gradually building up the complexity, I was able to develop a multichannel stochastic synthesis generator.

Although I now had this very powerful tool, there was no easy way to control it in a live setting, so I abandoned that path in favour of rendering in the studio. I could now produce multiple channels of related sounds; and I had the ability to control the output in real time. As I tested this generator in the multichannel studio, I recorded some of the material, taking further control of the sounds. I used this recorded material to take the obvious next step, writing *BlckWnd* as a fixed media composition.

BlckWnd was influenced directly by the facilities available at the University of Birmingham and the interactions I had with colleagues. The latter I have already mentioned and the former can be seen in my 8 + n channel[1] works, where the speakers in the studios are laid out in a ring of eight with separate stereo monitoring. I was also influenced by my penchant for industrial music. My synthesis-based work is still edited, rearranged and processed in the same way that acousmatic composers edit, rearrange and process, but I found noisy, harsh and gritty textures more interesting than the typical electroacoustic sound which features clean, smooth and harmonious textures. Many composers do wonderful work with synthesis and generated sound but I find it is still often done within the typical acousmatic aesthetic. I wanted to challenge this tradition. I believe some of my interest with the grainy side of synthesis has to do with my affinity for noisy music (after all, I am a percussionist and percussion sounds – even in western art music – tends more toward noise than harmonicity), but I also think that we can learn from the well-produced and well-mastered tracks coming out of the commercial music world. There is an enormous amount of room between electroacoustic music, noise music, techno, pop and sound art that still needs to be explored. I see my work leaning very much toward acousmatic music in the electroacoustic genre but drawing heavily from these other genres.

After I finished *BlckWnd,* the piece using the stochastic synthesis code I had written, I composed a few pieces that are not based on synthesis. *Poly, Immortal* and *Lithium* explored a different area of composition stemming from body percussion and singing respectively. *Spinning Plates* makes use of my live electronic background, while incorporating my stochastic synthesis material. It is more refined than my previous foray into mixed works. *River* and *Fields* complete the portfolio with my last synthesis based pieces.

[1] All of my eight channel pieces use the speaker setup of a ring; see Appendix 1. The + n is a variable number of extra channels from 0 to 4.

PROGRAMMING

Stochastic Synthesis

My research into stochastic synthesis is based on the work of Sergio Luque Ancona, predominantly his work on the algorithms of Iannis Xenakis. Xenakis began his research into using computers to synthesise sound in 1967 during his time at Indiana University in Bloomington and continued for a further five years after he moved to Paris in 1972. He changed his algorithm slightly for his further work in the late 1980s. This causes some confusion, since Xenakis uses the title *Dynamic Stochastic Synthesis* for his algorithms in both these time periods. In the introduction to a 2006 paper entitled *Stochastic Synthesis - Origins and Extensions,* Luque states "The available information about the dynamic stochastic synthesis algorithms is often incomplete, confusing or wrong" (Luque 2006). During one of my conversations with Luque, he relayed his own struggle with the text of *Formalized Music* saying that the algorithms given in the book do not match the sounds Xenakis uses in *GENDY3* (1991). *Formalized Music* was written in 1992 which was concurrent with the composition of *GENDY3* (Xenakis 1992). Luque presumed that the work was obfuscated intentionally. It seems plausible that Xenakis simply wanted to keep the cutting edge new work to himself for the time being. Unfortunately, when he died in 2001 he left no clear record of the later algorithm.

All of my generators and my further work are based on the 1991 implementation of the algorithm as reconstructed by Luque, and follow the same mathematics as Marie-Hélène Serra, Xenakis' assistant for the writing of *GENDY3* (Serra 1993), and Peter Hoffmann, though I have not confirmed that his program produces the same sounds (Hoffmann 2000). I will outline the specific procedure involved in stochastic synthesis here in order to differentiate my algorithm from others which use the title of stochastic synthesis (because they use random processes) but do not use Xenakis' method (Collins 2011, Brown and Jenkins 2004). While the algorithm itself is relatively straightforward and consists of simple processes, much of the published writing on the algorithm is very technical. I am therefore describing it in plain English to assist the reader.

The process relies on the liner interpolation of breakpoints. Each breakpoint has both its amplitude (between -1 and 1) and its duration (number of samples) determined by an individual second-order random walk. The last breakpoint is always linearly interpolated to the first of the next cycle. This ever-evolving waveform is what is sent out of the Digital Analogue Converter (DAC) as sound. The second-order random walk is the key element of the process. A first-order random walk begins with a single number from a ± distribution of random numbers (e.g. -10 to 10). This is used as the step size for the random walk. The random walk is constrained by an upper and lower barrier using a fold function (a fold function reflects the new number back into the range so that 106 in a range with an upper limit of 100 would become 94). A new step size is generated and added to the old output, folding it into the upper and lower barriers if necessary. This means that a new number is derived from the previous number. With a small step size and reasonable barriers (such as ± 10 and 0 to 100) you will get the effect of a wandering number. A second-order random walk uses a first-order random walk to generate its step size. The main difference between first and second-order random walks is therefore that the barriers of the first-order random walk, within the second-order random walk, are ± a number. As in the first-order random walk, the next number is based on the last. The first-order random walk generates a step size and that is added to the last output, folding as necessary. Using the duration as an example the process is started with randomly generating a -1, 0 or 1(Rand.). The barriers of the first-order random walk (FORW) are -5 and 5 and the barriers for the duration (second-order random walk, SORW) are 50 and 100 with the initial values of 0 and 75 respectively:

Rand.	N/A	1	1	0	1	1	1	1	-1	0	-1	1
FORW	0	1	2	2	3	4	5	4	3	3	2	3
SORW	75	76	78	80	83	87	92	96	99	98	100	97

The beauty of this technique is that it relies on a single random number for each of the parameters. The second-order random walks tend to hover around the barriers. This behaviour is due to the first-order random walk's tendency to stay either positive or negative for stretches of time. If the output of the first-order random walk is positive then the second-order section will keep increasing and the fold function will keep it within the constraints of the folded barriers as seen in the example above.

One way of creating multiple voices is to set up similar generators with a different random seed for each. In this execution, each parameter is completely independent. In Luque's paper there is mention of a method to link multiple instances of the generators (Luque 2006, p. 29). While Xenakis appears to have originated the idea, the first implementation I am aware of is Luque's in his yet unpublished SuperCollider UGen (Xenakis 1992, p. 298). Linking multiple generators can be accomplished by having one second-order random walk control the duration of all the generators, and by keeping the amplitude values independent. This technique works on two or more speakers giving the synthesis a single unified sound that shifts between the speakers in stereo or around the space in surround multichannel setups. Extending this approach by interpolating between this linked version and the completely independent version produces some interesting results. Once there are multiple generators sharing the same duration, it is a simple matter of interpolating between their single duration stream and an equal number of parallel duration streams with their own random seeds. The independent voices of multiple generators can coalesce into a single voice and vice versa. Because of the variations in duration between the single duration stream and the independent streams, pitch fluctuations can occur in the time between moving from one to the other that bear no resemblance to the pitches produced by either one. The effect of these pitch fluctuations can be minimised by selecting sounds with properties that do not suffer from pitch bending (for example, noisy timbres).

Stochastic synthesis produces a vast array of interesting sounds. This approach is deficient in one area; because all of the durations are based on numbers of samples, the frequencies produced are limited. It is for this reason I implemented floating-point durations (i.e. breakpoints that fall between samples). The next step is to calculate where the linear interpolation between the break points intersects with each sample. Opening up all the pitches between the sample-limited ones gives rise to other possibilities. If the minimum and maximum number for the duration of the breakpoints is set to the same number, clear pitches can be attained. If the numbers used for the pitches are integer steps in a 12√2 curve, the result is an equally tempered scale. As far as I am aware, my implementation is the only one using both Xenakis' algorithm and sub-sample accurate breakpoints.

I programmed the implementation of the above functions in the C language. There were advantages to coding this process in a lower level language. Compared to Max/MSP, the tool I normally use to assist my composition, I had much finer control at the sample level. This enabled efficient loops with no unnecessary overhead. At this level there was also the capability of dynamically generating the desired number of synthesisers and of tailoring the object to the needs of the user. The communication between the different synthesisers is far easier in C than in Max/MSP. The “poly~” function has some capabilities for creating multi-voiced functions in Max/MSP but all of the internal communication has to be coded from scratch. Simply using the same variable in C was sufficient for my purposes (such as the use of the same duration across all iterations). Given the difficulty of programming sample-accurate Max patches, the sub-sample interpolation to get accurate pitch is something I could not imagine doing in Max/MSP.

There are some possibilities for stochastic synthesis that still merit further research. One was presented in the Luque paper; using one complete cycle from generator A then one from generator B, back to A and so on, controlled by a further random process and not limited to two generators (Luque 2006, p. 33). The amplitudes of the generators could be linked in the same way as the durations are. This may produce a single, possibly oppressive, sound at one end of the spectrum and sound little different than regular decorrelated stochastic synthesis at the other. The reason the duration-linked voices sound unified is that we can perceive the links between the sounds – phase-coherence – even though the sound may be changing thousands of times per second. With both the durations and the amplitudes linked, there will simply be one generator feeding all channels. If the amplitude is linked and the duration is not, then the similar amplitudes may differ widely in the time domain. The use of different probability distributions at the beginning stage in the generation of a second-order random walk or varying the barriers in real-time may also produce interesting results.

Max/MSP Patches

RandClick

Getting sample-accurate control of parameters in Max/MSP is difficult when dealing with indeterminacy. Using indexes and counters can get only so far in terms of control. It is for this reason I developed the RandClick object in Max/MSP's JavaScript object authoring environment. The RandClick object deals with a small area of control involving ramps and timing. The original intent for the object was to generate clicks, a single sample spike with a value of 1, at random intervals and to know when the next click will occur. It is for this reason the first output emits the clicks, the second a running index (counting samples up from 0 to the current sections total), and the third broadcasts the current total for the duration of the gap between clicks. These three signal outputs are sufficient for some processes (triggering sound samples, starting windowed oscillators), but I inevitably found myself connecting this to a 0 to 1 ramp generator to get these processes to execute. However, this technique lacked fine control so I added the last two outlets, the fourth and fifth, to add missing functionality and streamline the workflow. The fourth outlet emits a ramp from 0 to 1 over the total sample count (output 3) and the fifth generates similar ramps but with a defined ramp length in samples; the remainder of the total sample count is 0 padded. These last two outlets act as a random phasor – a function that is lacking in the standard Max/MSP program. The three arguments given to the object are min, max, and duration. Min/max set the minimum/maximum number of samples between clicks/ramps. Duration sets the length of the ramp for the fifth outlet, which must be within the range set by the min/max numbers.

OctPan

OctPan was originally designed to spatialise live performances. It was repurposed to function in n channels, currently from 1 to 8, and became useful in the studio; the movement of the ring in *BlckWnd* to the front left is an example. At its most basic level, the movement of the channels can be controlled with the mouse (click and drag). The black dots represent the sources; mouse movements within the inner circle will be panned around the ring of speakers using an equal power panning law. If clicking and dragging is done outside the inner circle and inside the outer circle, the amplitude of the sources will diminish relative to their distance to the inner circle. Gestures can be looped by toggling the loop button and then clicking and dragging in the UI. Each parameter can be automated to control precisely the movements (specifics of which are left to the user); some examples can be seen in the pre-set movements. These pre-sets use only three of the available parameters namely – the angle between the voices based on the centre of the ring, the centre of the ring in the UI, and the virtual position of the mouse. With these three parameters the user can define the distance between, the radius of, and the virtual focus of the sources. All of the remaining control parameters can be found in the help file.

Multiple loops

This object records and layers loops of exactly the same length. The patch accompanying the commentary will record an initial loop with its length being defined by the user (two triggers; one to start, one to end). The subsequent layers will be recorded in separate tracks and will use the same loop length as the initial loop. The interface is set up to require only one trigger to move to the next layer. This is accomplished by recording all the time the first loop is playing and only adding the loop to the mix if the trigger is hit at some time during that particular loop.

Multi-delay 8

Multi-delay takes a signal and sends it to random speakers in an array (8 in this case) after a delay. There is a pre-delay (ms) function which shifts the whole process later in time (for example, with a pre-delay of 20 ms and a delay of 75 ms, the first repeat will happen at 95 ms, the next at 170, the next at 325 and so on). The delay sets the time between voices (ms) and the feedback (%) sets the level of the next voice compared to the previous. Finally the number of voices (delay lines) can be set into the argument of the poly object. This patch could be used with the piece "Poly".

Morse Code Reader

This patch will detect onsets or triggers and covert them into Morse Code. Converting the code into a trigger-based model requires all sequences of dots and dashes to be prefixed with a dot. This is necessary because the patch measures the time between events to determine weather the input is a dot or a dash so the first ‘dot’ starts the timer. There is a 350 millisecond wait time after each trigger to determine whether the next trigger is a dot or a dash. One full second after the last trigger the output will be sent. So if the next trigger occurs within that 350ms time it is counted as a dot; if it occurs after 350ms but before 1000ms it is sent as a dash.

So the letter "A" (.-) would be tapped out:

Trig1 Trig2 Wait Trig3 Long wait.

Trig1 tells the system that I have started; Trig2 enters the DOT of the A; the Wait is more than a 350 ms pause but less than 1000 ms to tell the system that the next trigger is a DASH; and Trig3 to end the dash. This is a very latent approach to control, but the impetus was to have a wide range of possible commands with the minimum number of inputs.

FFT Decorrelation & Convolution

Convolution algorithm

All of the convolution patches use an algorithm that multiplies the amplitudes and adds the phases of the respective signals. In the patch this is executed in a less cpu-intense fashion by using the real and imaginary numbers directly. It multiplies the real inputs together and multiplies the imaginary inputs together and sends the resulting difference to the real output and the sum of the two to the imaginary output.

Decorrelation algorithm

The decorrelation patches are based on a patch by Andreas Mniestris derived from Gary S. Kendall (Kendall 1995). The algorithm takes the phase bins and offsets each one by a random amount.

Equal Power Loop

This patch records a loop, then overlaps the ending of the loop with the beginning (by a user defined number of samples), crossfading using an equal power panning law.

Audio Rate Pedal

This patch requires a custom pedal that makes and breaks a connection between an analogue output and an analogue input on the audio interface for the computer. The pedal will generate a trigger in time with audio signals which enter the computer from the same interface into which the pedal is plugged. The patch was originally designed to be used with the loop patches.

PIECES

In a Cage

In a Cage was the first piece that I composed for my postgraduate studies. It started as an étude designed to stretch my Max/MSP abilities but was then expanded into a full piece. The work consists of four sections: the first introduces a pedal chord building up across eight channels individually; in the second, the wiimote controls FM noise, then a crackle sound and a mix of the two; the third section introduces of the distorted vibraphone material; and finally the theremin is used to control the harmonised vibraphone. The odd choice of instruments for this piece was no accident – there was a visual element to the piece that guided the instrument selection. As originally conceived, this work would have needed only simple amplification but the sound I composed for the opening lent itself to an eight channel setup, which the venue was capable of supporting, so I expanded the requirements. The audience can make a clear visual connection to what is happening sonically though some of the inner mechanics are not so obvious. I chose a very simple musical language for this piece to take advantage of the properties of the processes. A performance of the work is improvised using only the components of the performed sections outlined above.

This piece relied heavily on my background as a percussionist and improviser. The complete set of materials for some of my previous works consisted of nothing more than a vibraphone, computer processing and me. There was no score as such and the little planning was outlined only in my head. I seem to parallel David Tudor's path to live electronics in that there is a move from playing to improvising to composing (Collins, Nicolas 2004). This piece sits in the vein of improvised performance as Tudor's does or even free jazz, but this does not mean *In a Cage* is devoid of structure. Free jazz has clear structure but the musicians are tethered to it only lightly compared to other more strict forms of jazz. The structure for *In a Cage* is more rigid than free jazz, falling closer to Tudor's loosely outlined electronic pieces. Many of Tudor's works were written for himself to perform; similarly I intended this piece to be preformed by myself and structured it with that in mind, relying on the aforementioned four sections as my 'score'.

All of the performative parts of this piece involve a strong visual element, which clearly connects physical and sonic gestures. After the initial setup of the underlying chord, each section has its own visual language: the wiimote controller has a clear relation to the sound and unlike the subtle movements of some smaller instruments, the audience can see exactly what is happening when the vibraphone is being played. Likewise, the use of the theremin is easily connected to its sound, even though it is being used as a controller. These different movements connected to sound formed the main concept of the piece. I wanted to lead the audience into the sound world with something to which they perceive a concrete causal connection. The difficulty was creating a piece that made sense of all these elements.

This was one of my early forays into eight channel composition. The multichannel pedal chord that runs through most of the work consists of a single filtered noise stream split into eight. This was the origin of the idea of modifying a single instance of a sound to produce an eight channel track, something on which I expanded in later pieces.

Using the wiimote's visual nature, I decided to control sound in a transparent way; the accelerometers in the device make it easy for an onlooker to connect movement with control. The right hand controls the pitch/frequency and the harmonicity ratio of the generators with the pitch and roll respectively, while the left hand controls volume of the both generators (pitch only). The juxtaposition of extremely different sounds was the impetus for the third section: this was achieved by altering the very pure sound of the vibraphone with the crunching, grating nature of the distortion. The last two sections relate directly to the visual connections I wanted to incorporate in the piece. Once the audience sees and hears the first pitch bend, they understand the relationship between the movement and the resulting sound. Ideally this captures the listener by engaging their understanding of the direct connection between the visual and audio information.

This work has been performed twice. On both occasions there were issues with the execution of particular sections. Some of these issues were in the composition and some were in the performance. In the more recent video recording I tried to deal with the performance problems but (because of time limitations) did not attempt to revise the piece. While this means that the compositional issues remain, it nevertheless gives an accurate presentation of my early work for the portfolio. The performance problems originated in the way the piece was composed and were not apparent until I attempted to play the piece, by which point many critical decisions had been made about the structure and content. The first performance challenge was getting mallets into the performer's hands. For the video, the solution was a choreographed mallet switch, which I consider a mixed success. The other issue, the last transition between the two processed vibraphone sounds, was solved by making a clean break between the sections to allow for the switch. On returning to the piece later, I realised that it is essentially some simple ideas created from related sounds, which results in a somewhat blocky form. (A sectional approach to form can also be found in a few other pieces of my portfolio, such as *BlckWnd* and *River*, but with some significant refinements.) Some of the transitions in *In a Cage* are weak in my opinion; for example, the first move to the vibraphone and the move from the distortion to the harmonisation section. On reflection, I could have set these up differently from the outset using different controls, but I had made the decision early on to use the wiimote for as much of the control as possible. This decision stemmed from the desire to have as little direct interaction with the computer as possible, and to keep the visual aspect of the performance clean. I had to break this convention early on in the process to set up the pedal chord at the very beginning of the piece, as the controller did not lend itself to the quick precise control that a simple fader delivered. I now believe that I should have deviated from this constrained ideal sooner and adopted a more functional level of control in order to preserve the musical intention. One example would have been to use a two pedal system, similar to that used in *Spinning Plates*, to step through the sections, tidying up the transition between the distorted vibraphone and the pitch bend section. To improve the transition further, the two processes could be overlapped to create an elision between the sections. Another improvement would be setting up the wiimote to function in a single hand freeing the other to transfer to the mallets.

The use of a simple musical language (only three pitches, A, D and G) was one of the composition's successful elements. The piece was constructed to show off the musical possibilities of select processes on a resonant instrument. The difference-tones produced with the distortion effect occur more frequently with the perfect fourths and fifths, while these pure intervals highlight the otherwise impossible glissandi. Essentially the limited musical language gave a set of rules to the performer. On one hand, this limits the player but on the other it can challenge the musician to explore the confined musical language fully.

Spinning Plates

This piece was the culmination of techniques I had learned to this point in my studies. It combines the live elements from my first year, the visually performative element which I find particularly interesting, and the stochastic synthesis from my research. Creating ambiguity as to the origin of the sound was the plan from the beginning of crafting this work. *Spinning Plates* addressed a number of problems with the way in which the first year piece *In a Cage* had been set up. From the outset, I wanted to compose a portable piece that could easily be passed on to another performer with little or no Max/MSP experience. To this end, I designed the performance interface to be as simple as possible – the setup was pared down to a minimum. Again, improvisation is at the heart of this work, but with more structure than *In a Cage*. There are free sections as well as strictly controlled sections. The piece as a whole is more successful than my first foray into mixed music.

The set-up consists of:

- a suspended metal plate (the dimensions and specific material are up to the performer);
- a contact microphone (an accelerometer was used for the recording) and an instrument microphone (both microphones are set very close to the plate);
- a transducer directly attached to the metal plate connected to an amplifier (allowing the plate to act like a speaker with the sound being filtered by the plate's own resonances);
- various sticks and mallets to excite the plate;
- an in-ear monitor;
- stereo amplification for the concert hall;
- two trigger pedals;
- a computer with an audio interface with at least four audio outputs (two for the amplification, one for the transducer and one for the in-ear monitor) and two audio

inputs (one if only the contact microphone is used), and an interface for the trigger pedals.

Here it is important to outline where the sound sources originate. The direct sound from the two microphones is sent to the stereo speakers at all times. For the video recording of this piece, I added a third microphone to pick up more of the lower frequencies but this microphone was not used in any of the processing. At times, when sound is being sent to the plate directly, the audience will hear the content of that sound "processed" by the plate. The opening, played with fingers, and the brush section use only the plate as the sound source. Next, the generated noise envelope is sent to the plate and, of course, picked up and passed on. The rolls are treated similarly to the first section. The *ritenuto* notes at the end of this section are created by sampling the last physical attack of the mallet on the plate and playing it back through the plate, resampling the result and repeating the process. The chorusing effect on the scraping section is sent out of the front speakers only. There is one instance where the audio sent to the plate is being sent to the speakers as well – during the stick section, when the onset-detector controls the changes between the two timbres, the stochastic synthesis is being sent out to the main speakers. These last two sections were intended to widen the sound by presenting the listener with the only stereo moments in the piece. Ending as it began, only the direct sound from the plate is heard in the final finger section, finishing with the brush scrape.

While a visual performative element in *Spinning Plates* is present, it does not have the same impact as the larger physical gestures of *In a Cage*. The subtle sounds lent themselves to a much more confined visual space. The physical setup of the piece complicates matters of visual clarity. Placing the plate in front would obscure the performer. Orienting it sideways limits the audience's view of the performance if they are behind the performer and the people on the plate side have the same plate-obstructed view of the first scenario. The obvious choice is to have the plate at a forty five degree angle with the performer in front. This had its own problem of some audience members looking at the back of the performer but I deemed this to be the best compromise. The video is shot mostly side-on because there are some moments where it is useful for documentary reasons to see the gestures behind the instrument. Even the subtle movements in this piece have the ability to engage the audience by connecting a physical cause with an aural effect.

There are direct similarities between Stockhausen's *Mikrophonie I* (Stockhausen, 1964) and *Spinning Plates* for example, the percussive semi-improvisatory nature of the piece, the suspended metal resonator, microphones used to pick up the sounds and processing of the signal. Like *Mikrophonie I,* I see this piece fitting in to the genre of processed live instrumental works but the similarity ends at this point. This is a much easier piece to transport: one player opposed to six, small setup, simple patch. *Mikrophonie I* is intended to explore the realm of possibilities to the full; I wanted to expand upon a few simple ideas and distill their musical possibilities. The implements used in *Mikrophonie I* are many and varied whereas the mallet choices in *Spinning Plates* are in any percussionist's collection. Many of my pieces are either difficult to perform or problematic to mount. With *Spinning Plates* I wanted to create something that can be played by someone else with the least amount of technical difficulty.

The portability of the piece was important to me. From the small setup, to the interface, to the simple Max patch, the real constraint was simplicity. There is some room for stretching or squeezing the setup, depending on the available equipment. Examples of this include the possibility of using one or two microphones and, if the switching of the section is handled by a second person, then only one pedal would be needed.

Improvisation runs throughout the whole work. There are sections with very specific gestures and tight rhythmic passages, but the length of sections is flexible to some extent, to accommodate the improvised material. Instructions in the music such as "Repeat until the feedback peaks" are open enough to account for variation in any number of elements in the signal chain, sampling and responsiveness of the plate. At the end of the second section with the brush, the crossover needs to be very precise and is set up with the in-ear monitor by playing the performer a preview of the pulse section that follows. An example of the elasticity of the sections can be seen in the passage of long rolls. Here the performer builds the roll throughout the gritty material's decrescendo – s/he has open control over when to fade the material out completely and when to start the next section.

I consider *Spinning* Plates a more successful mixed work than *In a Cage* for a number of reasons. The set-up for this piece, as compared to *In a Cage,* is far simpler. There are fewer interfaces to manage (pedals as opposed to theremin, wiimote and computer). There was greater focus on the structure of the work. The musical links are stronger in this piece – for example, it ends as it began, with the use of only the performer's fingers on the plate. The minimal nature of the setup benefited the structure by limiting the palette of the performer and spurring him on to explore the constrained musical space fully. The magnification of the minutiae brings into focus the delicate sounds; this draws the listener into the small details, which is the intention of this piece.

Poly

Poly is a fixed media work in this incarnation, but also has a few other guises. All sounds are made with nothing but the performers' hands and bodies. The music fits together by layering ostinati to build up lively and complex rhythms. There are a number of different ways this piece can be realised (solo, ensemble or studio); my versions were realised in the studio and required specialised microphone techniques for some of the sounds.

This work started out as a live looping piece and only later evolved into fixed media work. I assisted with a workshop given by Dr. Scott Wilson and the percussionist Joby Burgess about writing for live electronics and percussion, and the accompanying compositional and technical challenges. I had started planning and recording material before the workshop but seeing the simple rhythms used in new ways inspired some of the remaining work. I avoided strict rhythm in my compositions before this piece. As I am a percussionist, I had wanted to expand my use of other musical material. With the work I had done and the idea of looping and layering sounds gleaned from Joby, *Poly*'s parts fell into place[2].

All of the sounds are made from claps, snaps or body percussion. Again, portability was the reason for this simple approach. For a few of the sounds the microphone technique is important (in an ensemble version substitutions can be made to cover these parts); otherwise the sounds will stand on their own. Basic as the sounds may be, they cover a wide frequency range. Different claps give mid and mid-high sounds; snaps and two-finger-claps fill in the highs; body percussion, like chest hits, give bass notes.

[2] Joby Burges and Scott Wilson, *Composing for Percussion and Live Electronics Workshop* (April-May 2010).

The layering of the parts was an important consideration in this piece. Multiple straight claps get muddy very quickly so separation of the parts was necessary, and was accomplished in several ways. I tried to avoid simultaneous clapping when constructing the separate rhythmic layers. If rhythmic alignment did occur then the separation was created with pitch difference. Finally, space was used to differentiate, like the hocketing at the end of the first section.

The layers are made from ostinati. Between each section, I break from the repeating patterns to allow both the listener and performer to regroup. *Music for Pieces of Wood* by Steve Reich uses a similar construction (Reich 1980); each part is built one note at a time and, once it is complete, the pattern remains unchanged. This process originated form his piece *Drumming* (Reich 1980, p. 64) and is consistent with his idea of music being a process itself (Reich 2008). *Music for Pieces of Wood* also has thinner moments in the texture which are necessary to temper the tension that has been created by the converging voices.

The versions of this piece that I am including in the portfolio are for fixed media. This piece works well in stereo or eight channel because the separation of the layering is only partially based on spatial separation. In the stereo field there is sufficient room to accommodate all of the voices, but the added space of an eight channel ring allows for clearer separation between parts. *Poly* is unique among my pieces in that the stereo version is as effective as the eight channel. As the piece was initially intended for live performance it can also be realised in a number of live settings. A single player could set this up for loop pedal or use a software based solution with Ableton Live, SuperCollider or Max/MSP, or a percussion ensemble ranging from 5 to 9 players could perform this work.

There are special recording techniques used for some of the parts. All of the claps, snaps and body percussion were close-miked for the submitted versions. One sound in particular, the bass hit from the third section onwards, was recorded by holding my fists as in boxing, keeping one stationary and striking the other against it, with the microphone on the far side. This causes problems in the ensemble version as having 5 to 9 microphones will make this piece prohibitively complicated for some ensembles. Substitutions are marked in the score; the above example could be replaced with a standard kick drum.

|*Kroʊn*|

This piece was a quantum leap beyond anything I had attempted previously in terms of material and scale. I recorded voice, which I had not used before, and combined it with more refined synthesis algorithms. The speaker layout is notable as the number (twelve, 8 + 2 + 1 + 1) is unusual and their functions are idiosyncratic. The eight channel part is approached in a new way and represents a large step forward from my undergraduate work, with the other channels adding flexibility for performance. Interjections separate the work into manageable sections and moderate the pacing.

The majority of the material in |*Kroʊn*| *is* based on the poem *The Cremation of Sam McGee* by Robert Service (Service 2007). I wanted to use the rhythmic character of the poem by incorporating the speech patterns directly into the piece. The actual words did not fit the aesthetic that I was trying to achieve, and would have distracted the listener during the performance. I therefore disguised the words by ring-modulating them with white noise; this removed all sense of meaning from the words but kept their rhythm intact.

Consisting of twelve channels, this piece was my first 8 + n channel fixed media piece. In this case the n was a stereo component plus two mono tracks. One mono channel carried the central vocal line. The other was a "special," designed to be played at a distance from the audience, very loudly, and preferably reflecting off a wall or ceiling. I wanted a central channel to project the vocal material from a single point and not from a virtual central position, panned between the two main speakers, in order to increase both the directionality and the clarity. It was important that the whole audience get the sense that the central channel was clearly distinct from the others, as it carried all the rhythmic material for the second section. The stereo track holds material intended for normal diffusion. While the eight channel, or ring tracks hold the majority of the content for this piece, I wanted the ability to take advantage of the stereo-centric way the BEAST (Birmingham ElectroAcoustic Sound Theatre) system is most often set up. The material in these two tracks relates to the ring content in such a way that it is spatially independent. This allows the diffuser to have the stereo content in the same plane as the ring or to move it independently in relation to the ring.

This piece opens with a crackling sound that transfers smoothly from one speaker to another. Unlike the opening of *In A Cage*, there are eight separate sound generators running in parallel from the beginning. While the sense of consistency in the opening of *In a Cage* stemmed from the single source and simple frequency ratios of the filtered noise, here the underlying structure was reversed, using the eight different instances of the same algorithm combined to make a single homogenous initial sound. The effect was created by sending all the generators to the front two speakers then fading out half of the voices from the left speaker and the other half from the right to create a split sound. I continued this effect by bringing in another speaker with the material from an adjacent speaker and repeating the voice splitting process until each speaker is playing a single voice. Once the process reaches all eight speakers, the voice rhythm enters in the centre speaker and begins its monologue with various granular textures and synthetic interjections. After the silence, a high pitch on the opposite side of the ring from a grainy texture sets up the voice layering of the remainder of the poem, rotates a quarter turn, stops, then moves another quarter. This layering takes advantage of the space by spreading the material out over the whole array. The piece ends with a chirpy noise mimicking the earlier vocal rhythms and is punctuated with the last stanza of the poem; which is, incidentally, the same as the first.

With this piece I tried to expand my use of eight channels from my previous works. The setup for those pieces was either a set of stereo speakers or a diffuse version of a single sound fulfilling the same function as in *In a Cage*. In |*Kroʊn*|, each voice is treated as a distinct member of the array and not as part of a smaller array or a fraction of the whole. Each voice at the end of the opening section occupied its own space but fitted in with the remaining voices by virtue of being generated in a similar manner. Similarly, the layering section has, quite literally, a separate voice in each of the ring of eight speakers; again the consistency is achieved by carefully mixing in the related vocal clips. These techniques create a sound-field in and around which the stereo material can thicken the texture depending on how the material is diffused. The idea of using a immersive sound-field as a component part to a multichannel piece was expanded in the pieces *BlckWnd* and *River*.

There are two major interjections in |*Kroʊn*|: a long drone in the first voice section and the silence immediately preceding the high/grainy section. My intention was to break up the different sections in a manner befitting the harshness of the larger sections. Smooth transitions or gradual changes would not have portrayed the jagged character I was trying to convey.

BlckWnd

This work is set out in an 8 + 2 format and is made up of many related but distinct sections. The + 2 is similar to the stereo tracks in |*Kroʊn*| and is treated in a different manner from that of the ring. There are a number of refinements to the structure of this piece compared with previous works, particularly my approach to the drones in *BlckWnd,* which is very different from my other works. I wanted to explore movement within the performance space and the sense of movement produced by the stochastic synthesis facilitated this nicely; however this sense of space contributed to making the problems of mixing this piece down into a more portable format insurmountable.

As in |*Kroʊn*| there is a separate stereo track intended for diffusion. In this case the separation lent itself to the “stem”-based manner of composing. Keeping material of similar function in stems allows for flexibility later on in the production process (e.g. the mastering stage). Performing with these two different layers of the piece adds flexibility. The balance can be adjusted to the room and the bodies of the audience dampening the space. Compensating for the speaker setup is also possible with this approach. In performance, diffusion of the stereo part was treated differently from that of |*Kroʊn*|. Here the intention was to suggest the sound of a very simple drum kit (snare and bass drum only) so the sounds were kept stationary and projected from the front.

I incorporated a number of refinements based on issues with the block-like structure of *In a Cage*. The sections in *BlckWnd* are not simply stitched together but progress in a continuous manner. All of the material in this piece was derived from stochastic synthesis with the two exceptions of the snare drum sample and the short clip of commercial music at the end[3]. While the timbres of the stochastic synthesis change throughout, the language is consistent, and this consistency of language creates continuity and helps bind the piece together.

I tried in *BlckWnd* to develop the idea of presenting a ring of sound as a cohesive whole by linking the generators directly using stochastic synthesis. This differs from the approach to the ring in the opening of |*Kroʊn*| where the homogenous sound was rendered by virtue of the 8 voices running separate iterations of the same synthesiser. The linking was achieved by using a single evolving array for the durations between break points and letting the eight instances generate their own amplitudes, and can be heard in the opening gesture where a single sound slowly evolves into eight distinct but related parts.

The second part of the piece sets up the recurring snare and base drum interjections and hints at the motive that fully develops just before the end. These interjections expand into the only complete continuous phrase in the piece and set up the climax, which is an extremely short (2.5 second) clip from *Hyperpower!* (Reznor 2007). This sample fits into the sound aesthetic and communicates the last big push more elegantly than would have been possible otherwise. The fractured nature of the piece serves to draw the listener in more by delivering a flowing foreground section at the end without any of the breaks that pepper the piece up to this point. There are a number of long drones in the work which change over time. Some of these are layered generators that move forward by their nature as controlled random processes, always evolving because of their unpredictable interactions. Other long phrases take advantage of the control built into the stochastic synthesiser. Using the many possible timbres and the incredibly fine parameter control, creating seamless organic transitions from one timbre to the

[3] Permission for this clip was still in process at time of first submission

next is possible. Pitch movement was the last of the subtle time-based changes to the background material. In Xenakis' work, all of the break points occur on samples (Xenakis 1992). With my system I generate break points between samples, interpolating the resulting stream. This is used in the last 1:42 of the piece to modulate the frequency of the drone that runs the length of this section. This enormous level of control I have with stochastic synthesis can be seen in shorter sections of the piece such as the metamorphosis from 2:07 to 2:24. The transition from a lively active texture to a thin static sound is a gesture that moves the piece forward and shows off the versatility of the algorithm.

I designed this piece for a large array of speakers in a concert situation, and as a trade off, it cannot be presented entirely successfully outside such a setting. I tried reducing the ten tracks to a stereo mix but this bears only a pale resemblance to the full production. All of the elements are present in the mixed down version, but there is none of the sense of space I tried to achieve with the ring. The sense of physical volume is created in concert by using many distinct sources for the eight voices (twenty-four channel, for example). The voices themselves create density by producing similar, but not identical sounds. With even a small system the diffuser can fill the space: this cannot be compared with the capabilities of the stereo version. Sometimes pieces can be modified compositionally to allow for an effective stereo reduction. The subject of a piece's capabilities came up during a conversation with Prof. Jonty Harrison. In his piece *Hot Air*, a section fades away into the distance (Harrison 1996). In a concert setting, the material can be played in speakers that directly relate to the distance away from the audience fading from close to middle distant to far speakers. For the CD copy of the work, Harrison felt that this would "not work in living rooms"[4]. For this reason the fade is shorter in the published copy of the work. Because the gesture was time-

[4] Conversation with Prof. Jonty Harrison (Dec 14, 2010).

based, a modification was able to be made that preserved the integrity of the work. Modification of this kind was not possible with *BlckWnd.* In the stereo reduction of *BlckWnd* the eight channel section collapses, becoming unintelligible, and the stereo track is crammed into the same space, thickening the already dense texture even further. *BlckWnd* does not survive this reduction.

In some of the eight channel material, there is a sense of movement which is an interesting side effect generated when the synthesisers are linked. Using the same durations for the eight stochastic synthesisers synchronises them in the time domain, but their amplitudes are generated independently. This means that the precedence effect takes over and the sound seems to move around the space, because the ear will localise the loudest generator and follow the sound as the highest amplitude moves randomly from one speaker to the next. The jittering nature of the sound is simply a result of the various second-order random walks interacting with each other.

The use of more coherent multichannel signals expanded the sound-field concept originally used in |*Kroʊn*|. The linked synthesisers create this coherence and deepen the illusion of immersion. The signals are hard panned to the individual channels and only from 1:26 to 1:43 do I deviate from this practice. This is because the phase-coherence is so important to this technique. The above gesture, a movement of all 8 voices to the front left, is accomplished with equal-power panning, inverse proportional volume reduction and a simple reverberation algorithm. I felt that the use of a more complicated way of creating a sound-field such as ambisonics would have an adverse effect on the phase-coherence of the signals and negatively impact the sound and flow of the piece.

River

The programme note reads, “Synthetic flowing textures dotted with sharp interjections form the basis of this piece. These textures and interjections move the listener through the different sections – from sound droplets to roaring rapids to booming falls to a quiet estuary.” The structure was planned from the beginning of the composition to mirror the path of a river from rain to ocean. This piece is unique in my portfolio in that it draws inspiration from a programmatic approach to music. It uses field recordings to which I applied little or no processing. This way of working led to a more acousmatic approach to the work.

The programmatic nature of this work suggested some sonic shapes that would connect the sound of flowing water to the textures in the music. The shapes heard in the piece are drops, water deflecting off obstacles, running water and waterfalls, to name just a few. I wanted to engage the listener with the longer sections. I used stochastic synthesis to create the continuous stream of sound needed to evoke a noisy river. Many decorrelated, random but similar sounds allude to the rushing sounds of quickly moving water. There are similar musical and structural ideas here to those in *BlckWnd*, but the execution is different, particularly in the smoother, longer phrases.

The dry recordings of river rock that I incorporated into this piece add a very different sound to the synthetic texture. This departure from my usual methodology was necessary to provide an articulate voice throughout the work. Stochastic synthesis is good at creating transitions from one generated sound to another and can do so smoothly, even in a short duration of time. The sharp cuts from one stream to the next needed causality to make musical sense. The real transients of the rocks give some cause and effect to the changing sounds.

There is a much stronger acousmatic element in this work. |*Kroʊn*| touched on this idea of presenting recorded sound from behind the "screen" of a speaker. There, it was used in a tenuous way due to the obliteration of the text. In *River* I wanted the connection to river rocks to be obvious. The recordings ground the piece with sounds that would be heard by a riverside and the audience connects these sounds to the physical place.

As *BlckWnd* used refinements based on *In a Cage*'s structure, so *River* makes use of further refinements drawn from the lessons of *BlckWnd*. There is a basic consistency of language but the recorded articulations bring the transitions into focus, sharpening the character of the music. While there are numerous interjections in both pieces, I felt more comfortable letting the stochastic synthesis run in *River*. This piece did not need to use little sounds to move the piece forward; the synthesis was able to carry a phrase on its own.

Fields

Fields is the only piece in my portfolio that was conceived from beginning to end as a stereo piece. I had been considering the idea of doing an acousmatic stereo piece for a while and only conceived the final idea for this project after finishing my multichannel pieces. Despite being composed in stereo, conveying a sense of space to the listener was still important in this piece as the sense of open areas in the title *Fields* implies. There was also some detailed work done with the synthesis in this piece that I had not explored before. The emphasis on recorded material is in stark contrast to the majority of my work but was intended to be in line with this being an acousmatic project. The two separate sound-worlds presented here connect with varying degrees of cohesiveness and contrast.

The title *Fields* is meant to convey a concept of open space and green and gold expanses of grasses and crops. Translating this into sound requires some compositional conceits to instill the sense of space in the listener. There is no reverberation in the piece; the audio processes were restrained to make sure the original sounds were recognisable, and there is the sound of wind in one form or another throughout the piece. This was done to convey to the listener the kind of acoustic events – and those events' reactions with open spaces – that they would hear outdoors.

The opening of the piece is completely synthetic until the entry of the sound of wind in leaves at 0:55. Designed to mimic the sound of wind on a microphone, the synthetic sound blooms from the virtual centre into a dynamic stereo image. Though this is stochastic synthesis there is still a great amount of control present in the rendering. This allowed me to sculpt the indeterminate nature of the sound to present a sound stream that blends well with the field recordings.

Hildegard Westerkamp pioneered the idea of a soundwalk, in which one listens intently to one's present environment (Westerkamp, 2001). She extended this idea first by simply presenting a contrived soundwalk (*Kits Beach Soundwalk)*, and then producing pieces such as *Beneath The Forest Floor* (Westerkamp, 1996). *Fields* is intended to run along the same lines as these pieces, albeit with the addition of synthesis not present in the Westerkamp works. It is for this reason that there are several environmental recordings presented in the piece with little alteration.

There are two drastically different sound worlds presented in this piece. The first is the natural sounds of wind in leaves and the little sounds of leaves and grass crackling. The second is the pitched and unpitched synthetic material. In the beginning these two sounds blend but throughout the piece they diverge until there is a clear distinction between what is 'real' and what is not.

Immortal and *Lithium*

I wanted to stretch myself by composing something very different from the other eight-channel fixed media pieces in this portfolio. Voice has been used throughout the history of electroacoustic music; Stockhausen's *Gesang der Jünglinge* (1956) is one of the earliest uses of recorded singing voice. *Immortal* and *Lithium* are a juxtaposition of more recent trends in acousmatic music to use unprocessed sounds for large sections of pieces, as in Martin Clarke's *Voyager* (2006), and the infinitude of possibilities derived from heavily processed sounds, as in Christian Calon's *Les corps éblouis* (1998). The material for *Immortal* is pure, pitched and tonal containing all natural sounds with minimal processing as opposed to the rather noisy synthetic textures of the other fixed media pieces (|*Kroʊn*| and *BlckWnd)*. The form was set out from the beginning while my other tape pieces grew organically. *Immortal* and *Lithium* each consist of one long gesture made up of flowing phrases moulded into continuous streams that blend into one another. Voice is the source sound for both, but as will be explained below, *Immortal* would not function as a choral piece. From the outset, I wanted to create something 'beautiful'.

A single singing voice, unprocessed, *senza vibrato* and close-miked was the material for this project. I chose one or two takes of each pitch and, with few exceptions, used the whole length of the take. The start of each sound was edited automatically using an amplitude threshold. All of these things strip away many of the characteristics that are usually associated with a choir: vibrato, multiple voices and the sound of a reverberant performance space are not present in this piece. It could only be a fixed-media piece. The single voice creates a unity of sound and generates other phenomena like chorusing when two copies of the same take are layered and slightly offset. The absence of processing in *Immortal* draws direct attention to these phenomena. The close proximity of the microphone produces a delicacy that will be kept when played over loudspeakers. The magnified sound makes the clicks, rasps and minor pitch variations audible, giving the sounds and the piece this delicacy. A slightly inhuman sound is derived from the lack of vibrato which creates smooth phrases and accentuates changes in chord density. It is not humanly possible to produce these sounds in realtime in performance; it requires the material to be treated electroacoustically. While using full takes of notes shaped the piece in unexpected ways, it meant that the ends of notes were natural; having all of the transient part of the notes aligned automatically allows for very crisp starts of chords.

Creating something beautiful is a difficult balancing act. I personally find moments of beauty in my other pieces, but I would not describe any of them as a 'thing of beauty.' *Immortal* has the ingredients for this aesthetic. The form is simple and elegant, the sounds are pure from beginning to end, and there is just enough dissonance mixed in to avoid the charge of being saccharine. With these elements and a little craft, I hope I have created something that could indeed be considered 'beautiful'.

Space, however, remained important in this piece. In the versions of *Immortal* and *Lithium* accompanying this commentary the sounds come from nine positions, a ring of eight speakers and a panned centre position between the front two speakers. The voices are static but pitches do move from one position to another through the work. This means that, depending on how the pieces are rendered, the speaker setup does not have to be a circle. It could be an arc across the front or randomly dotted throughout the audience with the caveat that each "voice" must remain individual and distinguishable in space. While this approach limited the construction of the piece a great deal, it meant that, with a little fine tuning, this set of pieces could be performed on almost any number of speakers.

The pitch selection was based on the range of the vocalist. The work begins and ends with a pitch in the centre of her range. The middle chord is a F major chord with C major seven above to create the climax of the piece, which spans the complete range. In the second half, a B natural is layered to create a chorusing effect and was chosen for its inherent stability. The pitch selection for the remainder of the work was determined by a hierarchy; foremost by the planned structure; secondarily by the consonance/dissonance of the note in the texture; and finally by consideration for the aural effect it had on the adjacent notes (chorusing, effect on the space, etc).

Lithium is based on *Immortal*, using the piece as a whole for the foundation. The work is a palimpsest. Approaching *Lithium* in this manner imposes strict limits: voice position, pitch and timing were unchangeable, leaving only volume and processing as my creative tools. While very little processing is present in my other pieces, *Lithium* is based primarily on processed sound.

A palimpsest is a manuscript that has been reused by imperfectly scraping away the old writing and replacing it with new script, leaving a remnant of the original behind, often only partially legible. The flow, form, timing and shape of this piece echo the original but the two pieces are quite different. The silky textures of *Immortal* give way to sandy ones in *Lithium*. Rough passages take the place of harmonious passages. The surface of the music has been replaced with synthetic textures.

I work well with strict limits – as can be seen in other pieces in my portfolio. Stravinsky writes: “the more art is controlled, limited, worked over, the more it is free” (Stravinsky 1979, p. 85). He then goes on to outline the scale, chromatic intervals, and accents as the concrete elements on which he imposes his narrow frame for each one of his undertakings, the details of which he does not disclose in this part of the text. Finally he states: “the more constraints one imposes, the more one frees one’s self of the chains that shackle the spirit” (Stravinsky 1979, p. 86). The fixed elements of *Lithium* were placement of the voices in space (in this version the nine positions on the ring); the pitches; onset times; and the durational structure. These limits did not, however, detract from the potential of the piece. Because of the clean way I approached *Immortal* (clear edits, no overlapping voices, and dry recordings), I had a solid base and plenty of room to manipulate levels and process the sounds to create something new within the narrow frame I had constructed for myself.

This work was my first venture into certain techniques commonly used in electroacoustic music composition. Previously, if I required a plug in, I would code it myself in an attempt to understand the effects and process that I was using. This was of great benefit because it expanded my knowledge of digital signal processing (DSP). This time consuming approach to DSP merited the investment, for with this understanding I can now use pre-existing tools more effectively. In *Lithium* I could go to a single vocal track in the piece, select appropriate plug ins and settings, and produce good sounds quickly.

Appendix 1

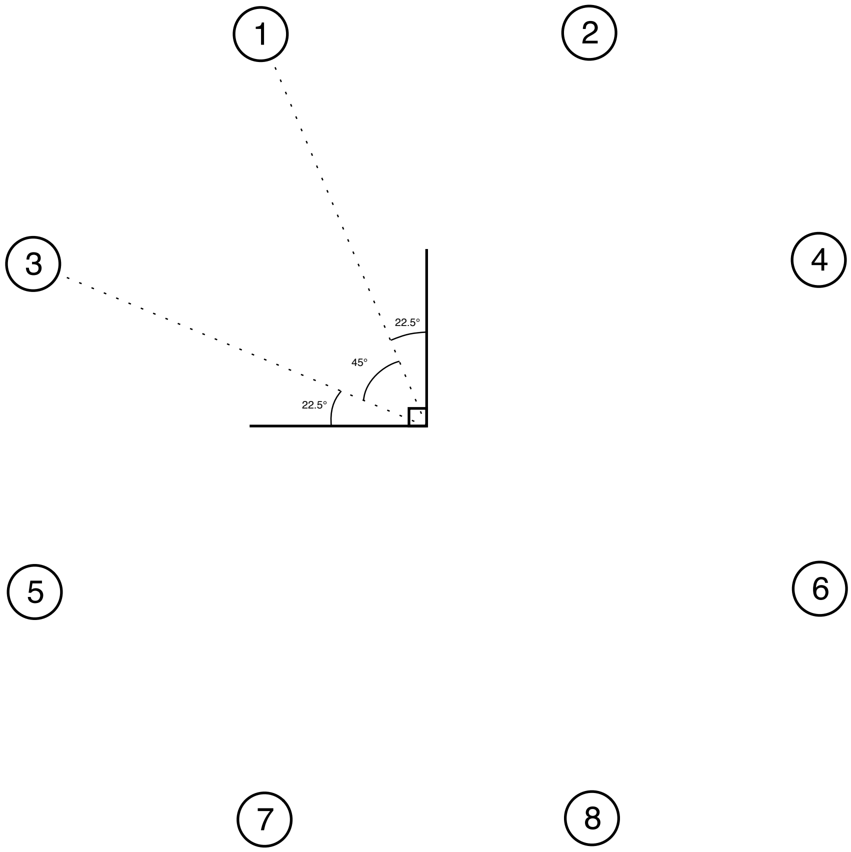

Layout of speakers and channel numbers.

Bibliography

BROWN, ANDREW R. and JENKINS, GREG, 'The Interactive Dynamic Stochastic Synthesizer', *Proceedings Australasian Computer Music Conference*, (2004), 18-22.

COLLINS, NICK, *Gendy1* and *Gendy2*, http://supercollider.git.sourceforge.net/git/gitweb.cgi?p=supercollider/supercollider;a=blob_plain;f=Help/UGens/Oscillators/Gendy1.html, accessed 14 June 2011.

COLLINS, NICOLAS, 'Composers Inside Electronics: Music After David Tudor', *Leonardo Music Journal*, Vol. 14. (2004), 1-3.

HARRISON, JONTY, *Hot Air*, Articles indéfinis (empreintes DIGITALes, IMED 9627, 1996).

HARRISON, JONTY, 'Sound, space, sculpture: some thoughts on the 'what', 'how' and 'why' of sound diffusion', *Organised Sound*, Vol. 3, No. 2. (1998), 117-128.

HOFFMANN, PETER, 'The new GENDY Program', *Computer Music Journal*, Vol. 24, No. 2. (2000), 31-38.

KENDALL, GARY S., 'The Decorrelation of audio signals and its impact on spatial imagery', *Computer Music Journal*, Vol. 19, No. 4. (1995), 71-87.

LUQUE ANCONA, SERGIO, *Stochastic Synthesis - Origins and Extensions*, (published online 2006) http://www.sergioluque.com/texts/luque_stochastic.pdf, accessed 2 June 2007.

REICH, STEVE, 'Music as a Gradual Process', in Christoph Cox and Daniel Warner (ed.), *Audio Culture: Readings in Modern Music*, (New York, NY: Continuum International Publishing Group Inc, 2008).

REICH, STEVE, *Music for Pieces of Wood* (London: Universal Edition, 1980).

REICH, STEVE, Writings on Music (New York, NY: Oxford University Press, 2002)

REZNOR, TRENT MICHAEL, *Hyperpower!*, Nine Inch Nails (Leaving Hope Music Inc, 0602517324297, 2007).

SERRA, MARIE-HÉLÈNE, 'Stochastic Composition and Stochastic Timbre: Gendy3 by Iannis Xenakis', *Perspectives of New Music*, Vol. 31, No. 1, (1993), 236-237.

SERVICE, ROBERT, *The Cremation of Sam McGee*, http://www.blupete.com/Literature/Poetry/ServiceCremation.htm, accessed 30 March 2007.

STOCKHAUSEN, KARLHEINZ, *Mikrophonie* (London: Universal Edition, 1964).

STRAVINSKY, IGOR, *Poetics of Music in the Form of Six Lesons*, Arthur Knodel and Ingolf

Dahl (Trans.) (Cambridge MA: Harvard University Press, 1970).

WESTERKAMP, HILDEGARD, *Soundwalking*, (published online 2001) http://cec.concordia.ca/econtact/soundwalk/Soundwalking.htm, accessed 5 June 2011.

WESTERKAMP, HILDEGARD, *Beneath The Forest Floor* and *Kits Beach Soundwalk*, Transformations (empreintes DIGITALes, IMED 9631, 1996).

XENAKIS, IANNIS, *Formalized Music*, rev. ed. (Stuyvesant, NY: Pendragon Press,1992).

Spinning Plates

By

Eric Bumstead

Spinning Plates
By
Eric Bumstead

Duration: 10:00

Instruments and Setup

- a suspended metal plate (the dimensions and specific material of which are up to the performer);
- a contact microphone (an accelerometer was used for the recording) and an instrument microphone (both microphones set very close to the plate);
- a transducer directly attached to the metal plate connected to an amplifier (allowing the plate to act like a speaker with the sound being filtered by the plate's own resonances);
- various sticks and mallets to excite the plate (a brush, snare sticks, soft mallets);
- an in-ear monitor;
- stereo amplification for the concert hall;
- two trigger pedals; and,
- a computer with an audio interface with at least four audio outputs (two for the amplification, one for the transducer and one for the in-ear monitor) and two audio inputs (one if only the contact microphone is used), and an interface for the trigger pedals.

Notes

The "IMPROV:" sections are suggestion and may be expanded upon.

Fan - Hold the brush underneath and parallel to the plate with the bristles to the right. Rotate the brush counterclockwise splitting the bristles in half (i.e. half in front half behind).

Sweep - Move the bristles across the plate or in the case of the sweep after the fan quick left right motions.

Rattle - With the bristles split by the plate, "shake" the brush front to back.

Pull to sweep - Place bass of bristles on top edge of plate. Pull brush out until the tips of the bristles are on the plate. Sweep down then scrape the very tips off the bottom.

The numbers in boxes above the percussion line are the different section in the Max/MSP patch and are changed with pedal 1. The other pedal (2) is a trigger and is notated as "ped." below the percussion line.

Spinning Plates

Eric Bumstead

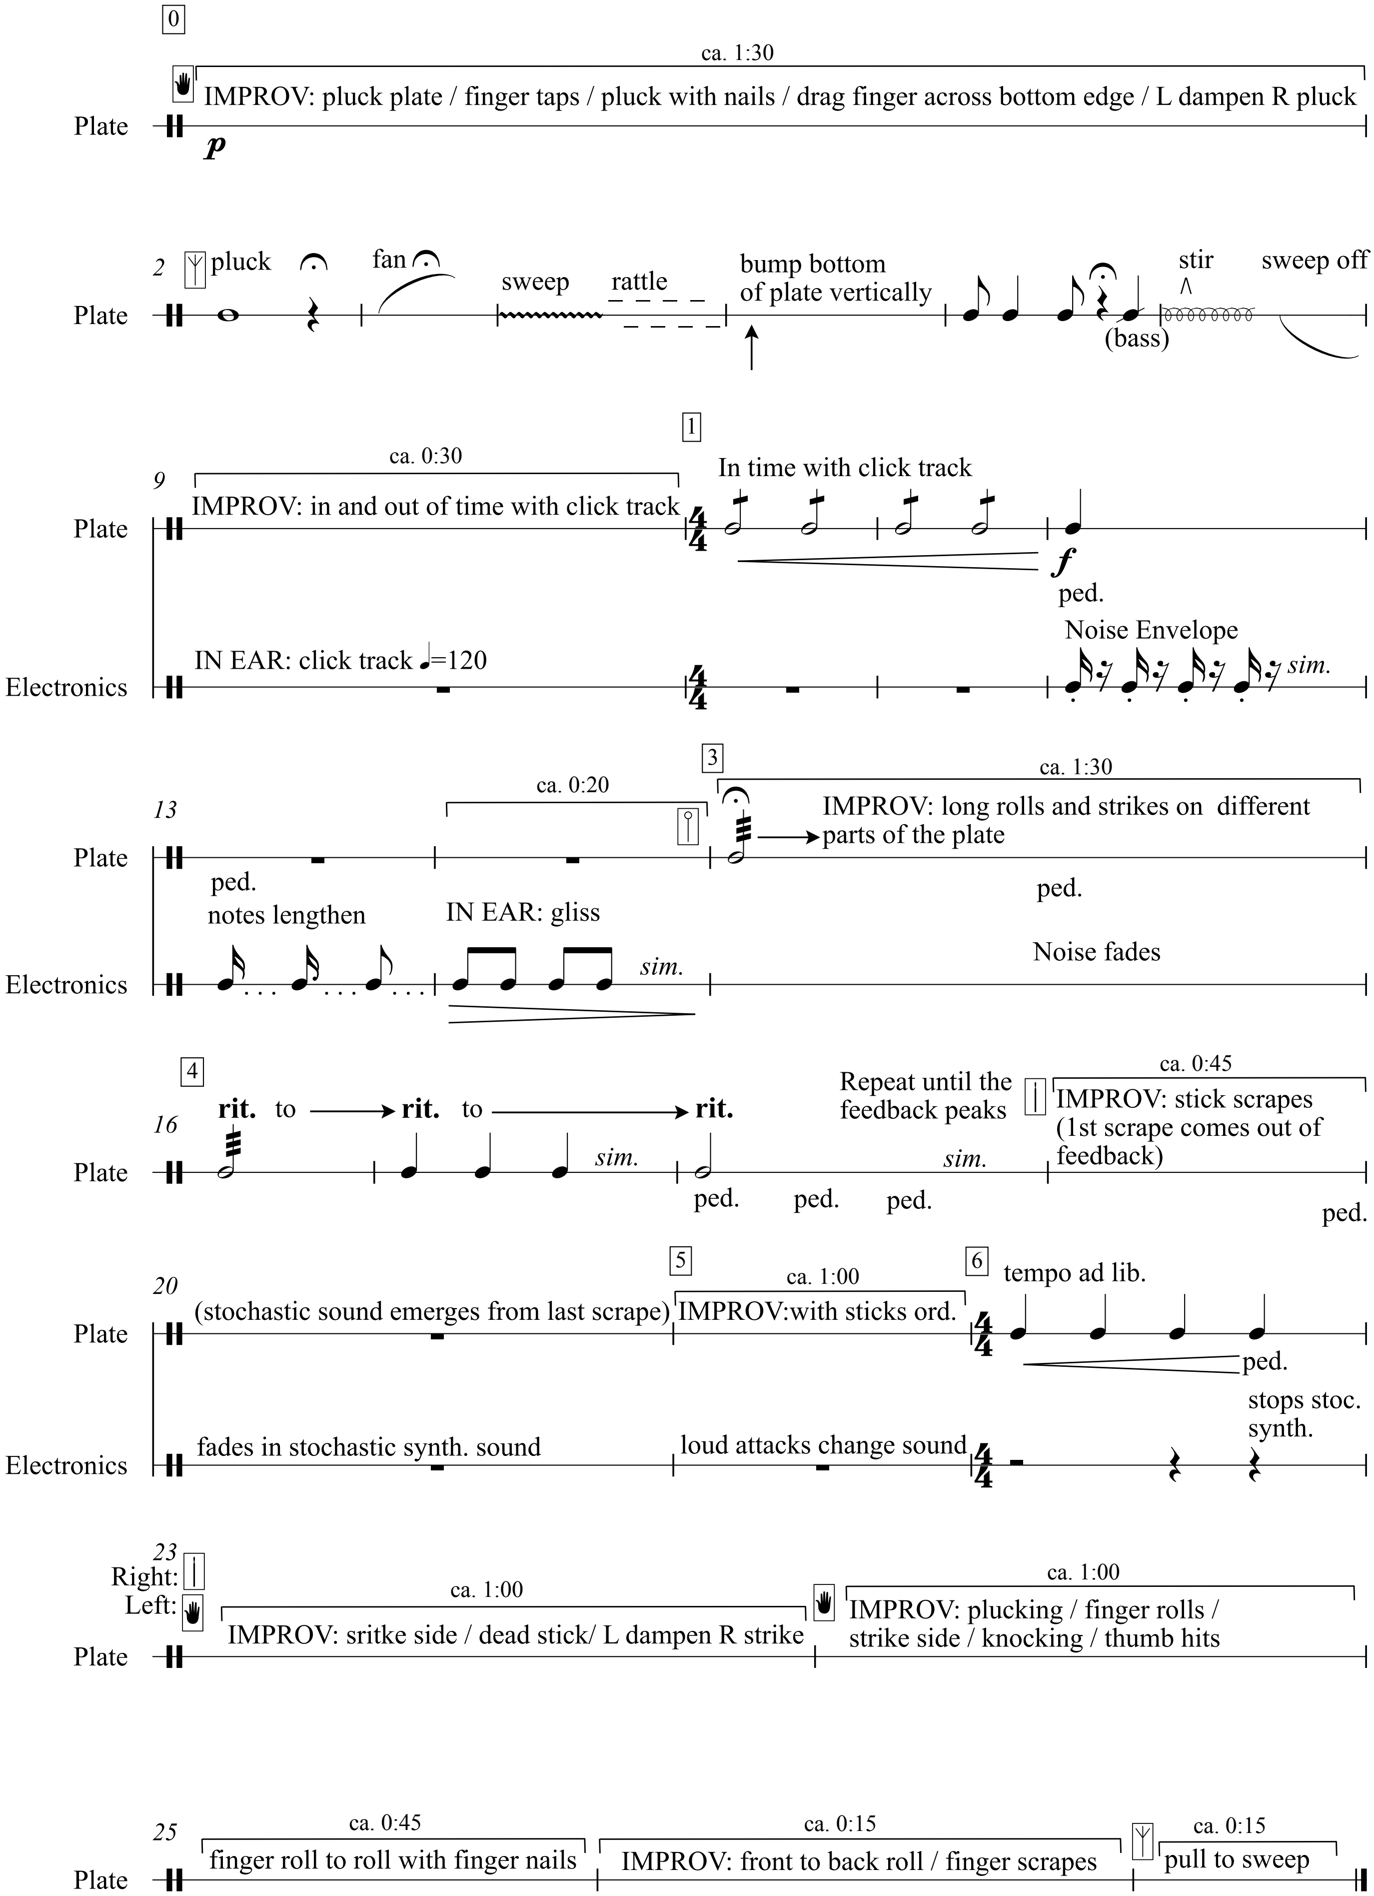

Poly

By

Eric Bumstead

Poly
By
Eric Bumstead

Duration: 5:00

Instruments

Clap	All fingers of one hand closed together, strike palm of other.
Echo	All notes with a X head are a clap through the echo effect (detailed below) .
R on L chest	Slightly cup the right hand and strike upper left part of chest.
Snap	Snap or click fingers.
Chest	Right hand on upper right chest, left hand on upper left chest, strike with flat hands.
2 finger clap	Index and middle finger of one hand closed together strike palm of other.
Finger clap	All fingers of one hand closed together strike all fingers of other hand closed together.
Chest thump	Fist, knuckles and heel, strike and hold on upper middle chest.
Palm clap	Cupped palms strike and hold together.
Boxer thump	Hold fists as in boxing, keeping one stationary and striking the other against it, with the microphone on the far side.

Solo version

This version needs some looping device or computer program connected to a microphone for recording and playing back the loops. The echo effect can be achieved using any standard echo effect set to 90 milliseconds with 75% feedback level.

Ensemble version

5 to 9 players. Echo effect as above. Microphones for all players preferred but depending on the concert space and the material available, not all players need to be miked. The part doublings and substations for 5 players are as follows:

From 23 to 29: Clap and Finger clap can be combined
From 27 to 35: Chest thump can be a kick drum.

From 40 to 54: Boxer thump 1 and Boxer thump 2 can be combined on a kick drum.
From 44 to 54: Clap Snap can be combined.
From 46 to 54: 2 finger clap can be dropped out at this point.

All parts should blend so that every one is equally present. The new voices should enter louder than the others and then drop back into the texture before the next change.

Poly

Eric Bumstead


17
(echo)
(R on L chest)
(Chest)
x2
R on L chest
Chest
Snap
Chest
2 finger clap
21
x2
R on L chest
Chest
Snap
Chest
2 finger clap
Clap
23
x2
R on L chest
Chest
Snap
Chest
2 finger clap
Clap
Finger clap

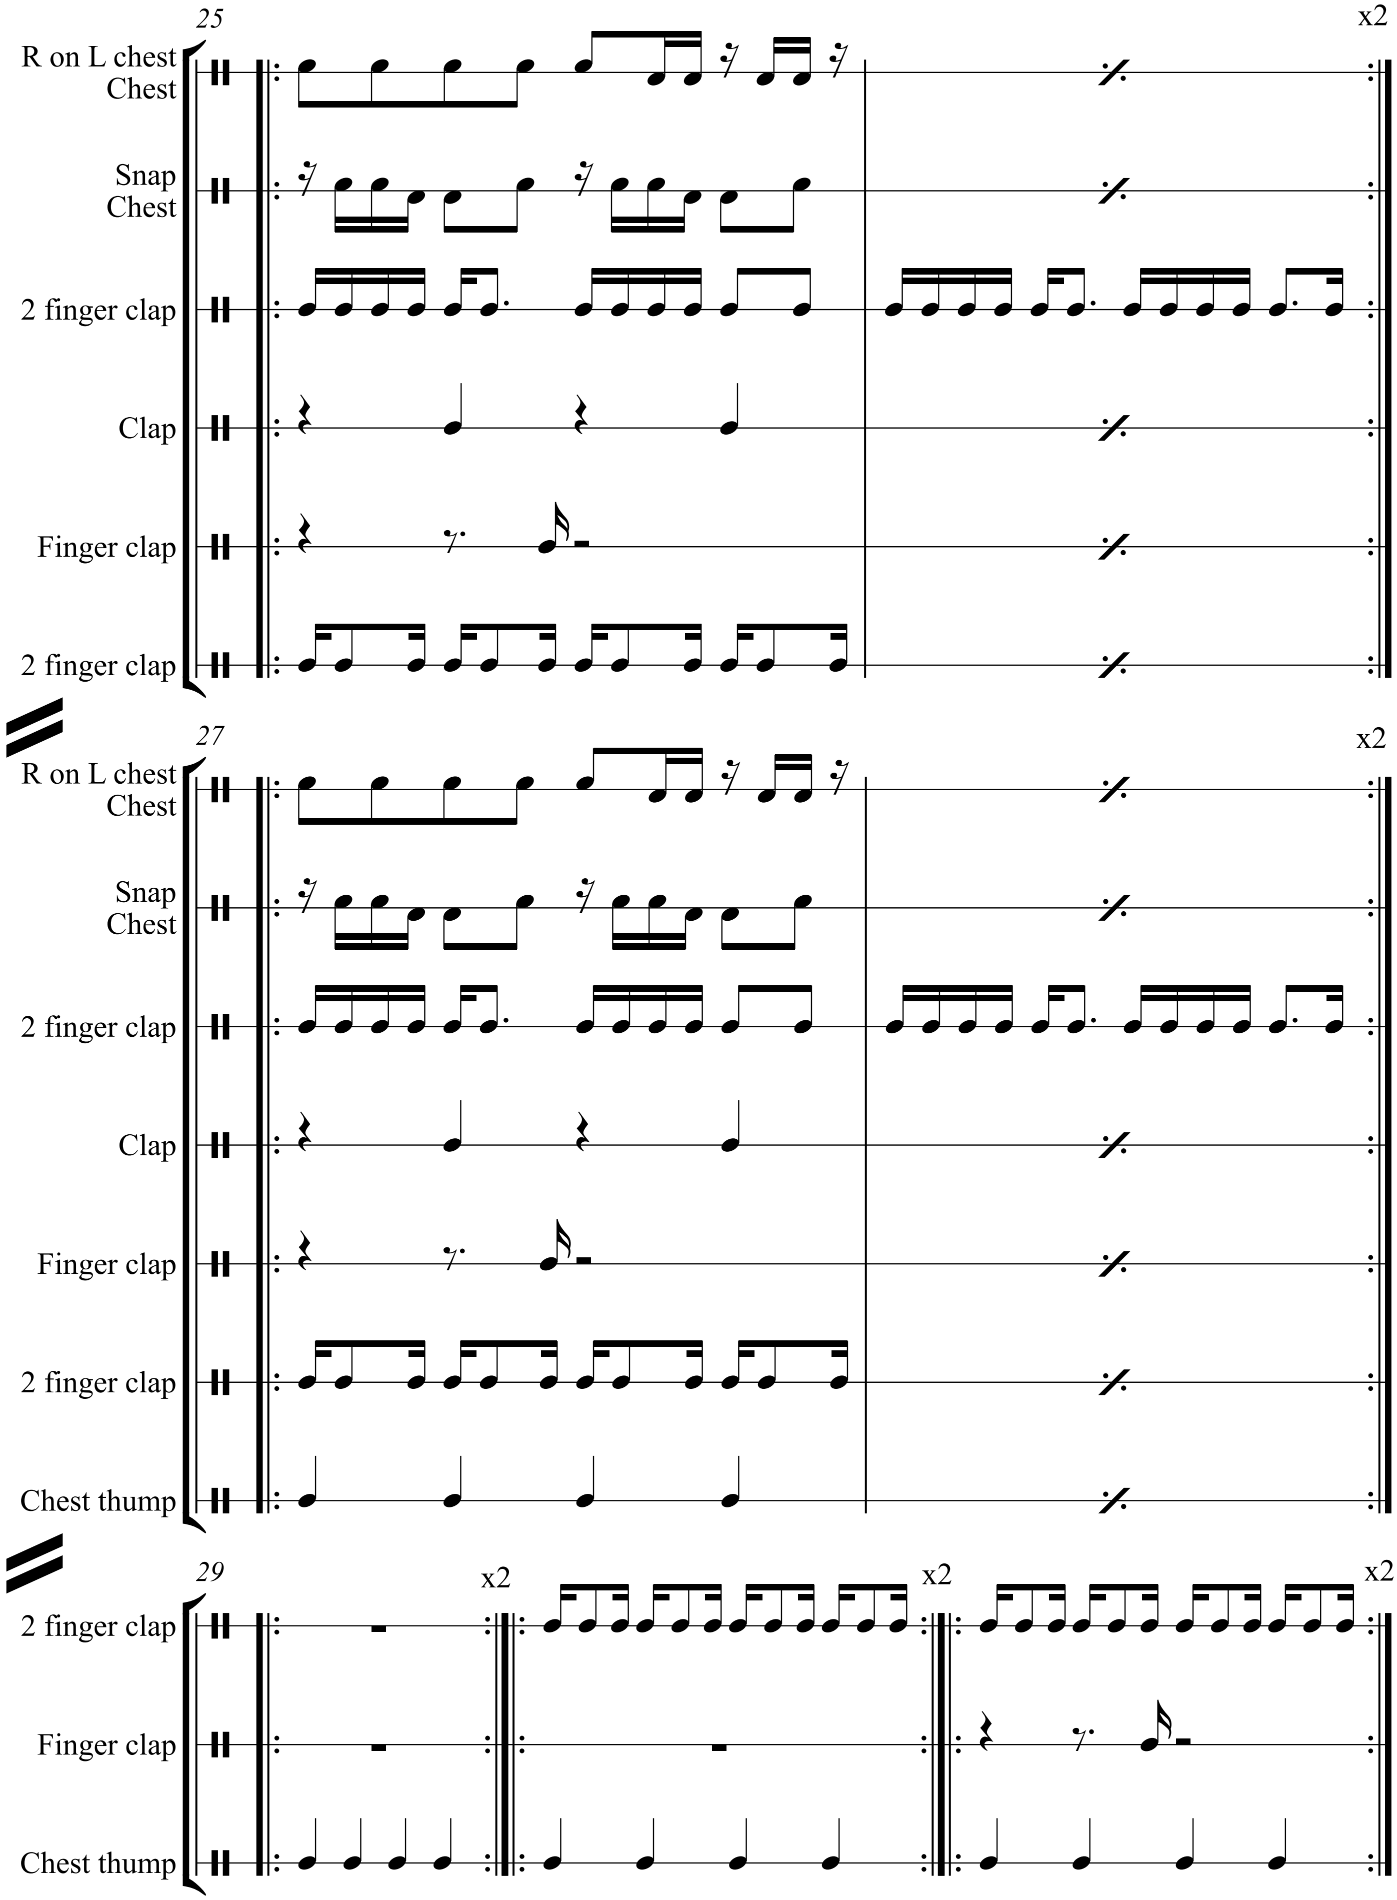
25
x2
R on L chest
Chest
Snap
Chest
2 finger clap
Clap
Finger clap
2 finger clap
27
x2
R on L chest
Chest
Snap
Chest
2 finger clap
Clap
Finger clap
2 finger clap
Chest thump
29
x2
x2
x2
2 finger clap
Finger clap
Chest thump

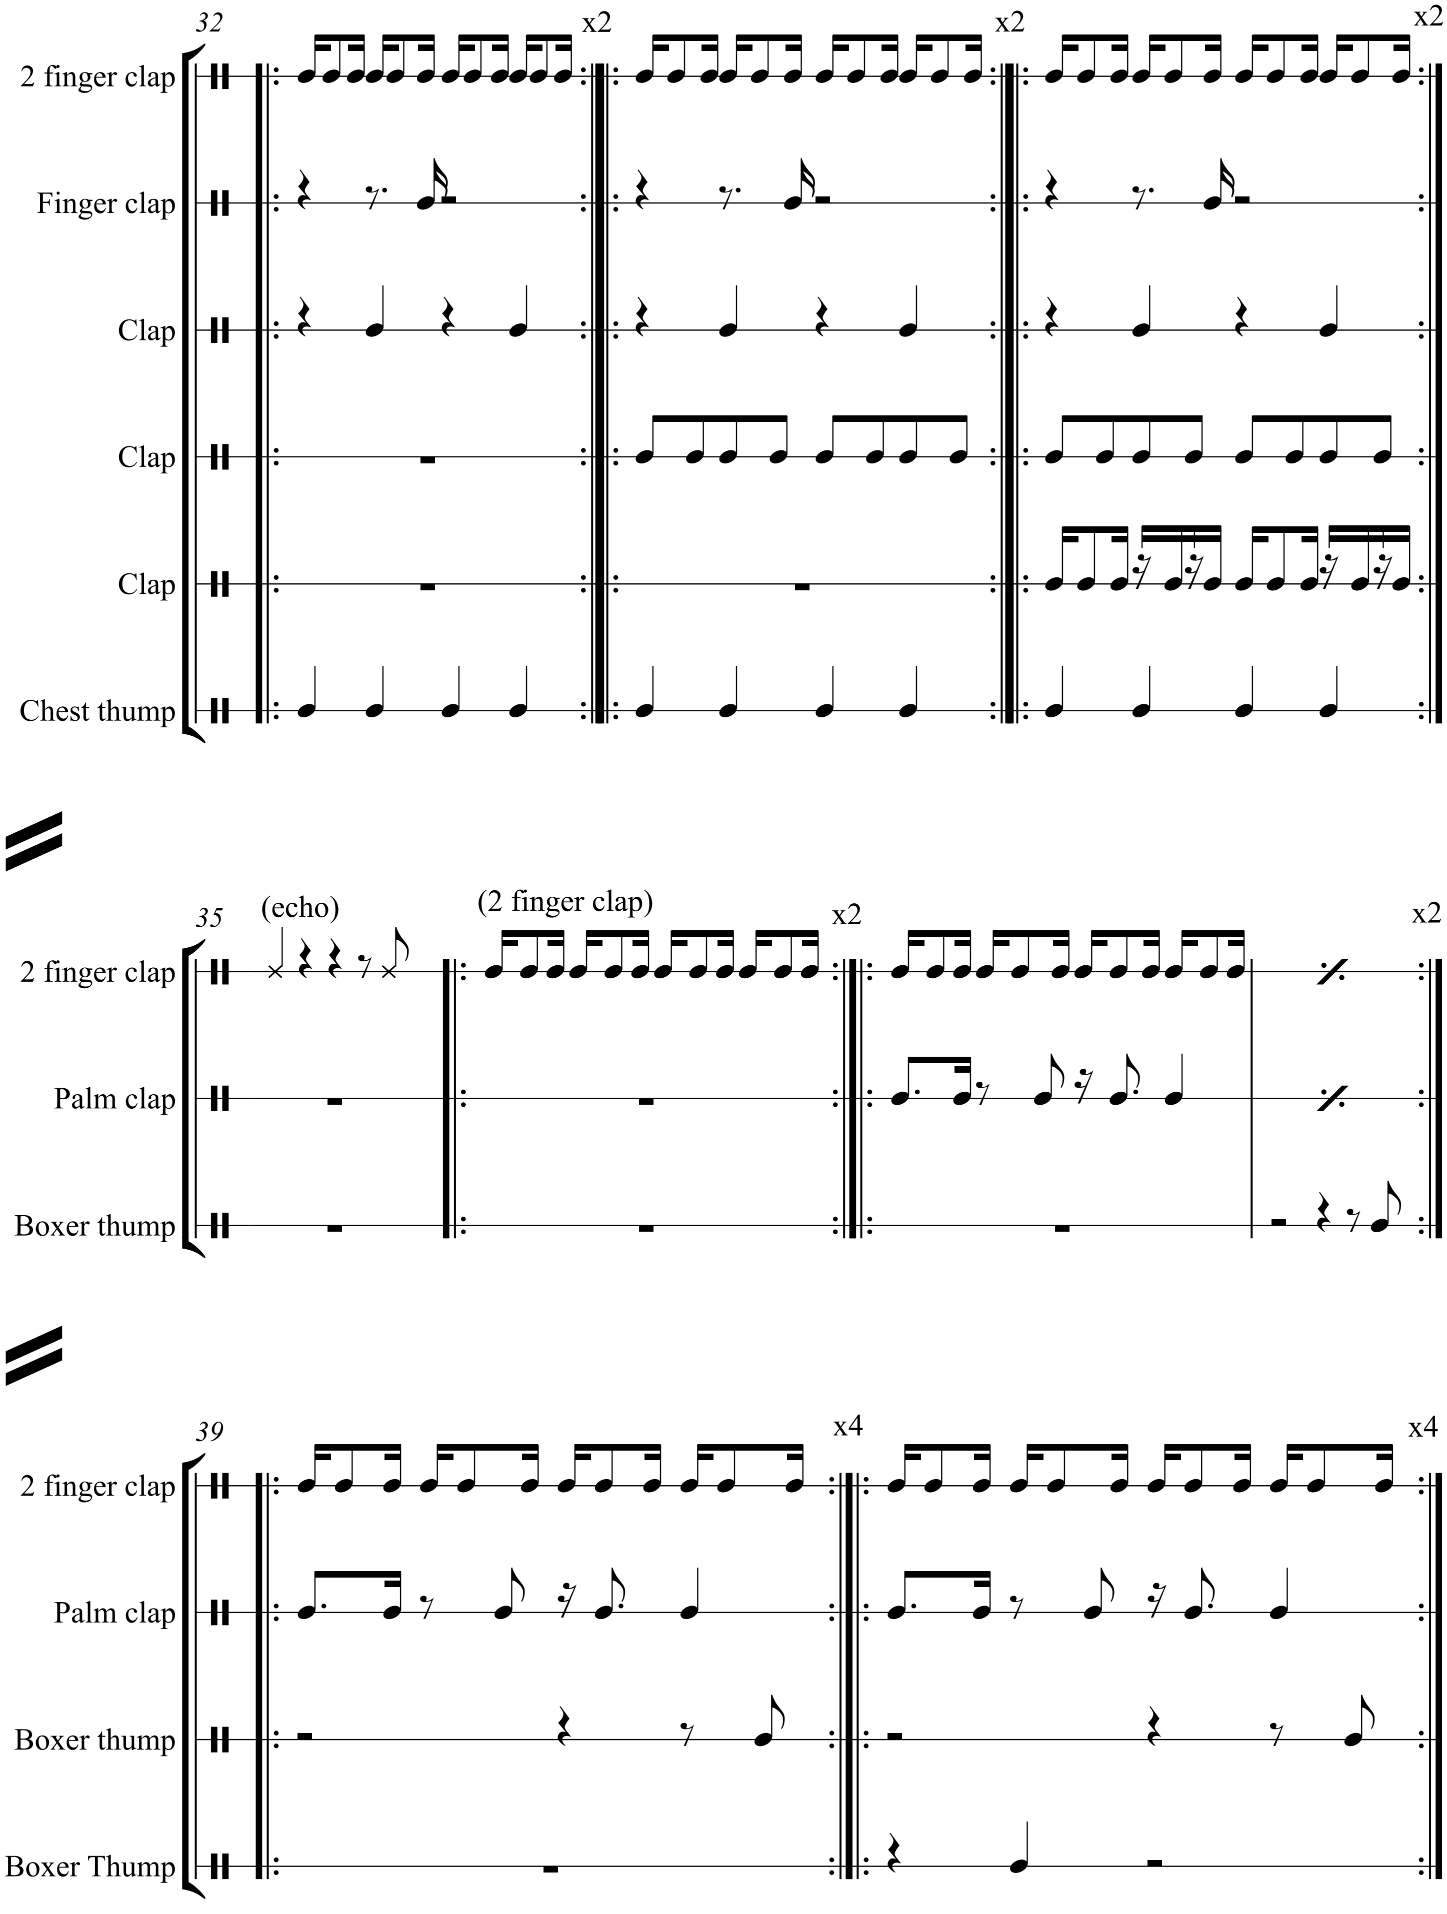
32
x2
x2
x2
2 finger clap
Finger clap
Clap
Clap
Clap
Chest thump
35
(echo)
(2 finger clap)
x2
x2
2 finger clap
Palm clap
Boxer thump
39
x4
x4
2 finger clap
Palm clap
Boxer thump
Boxer Thump

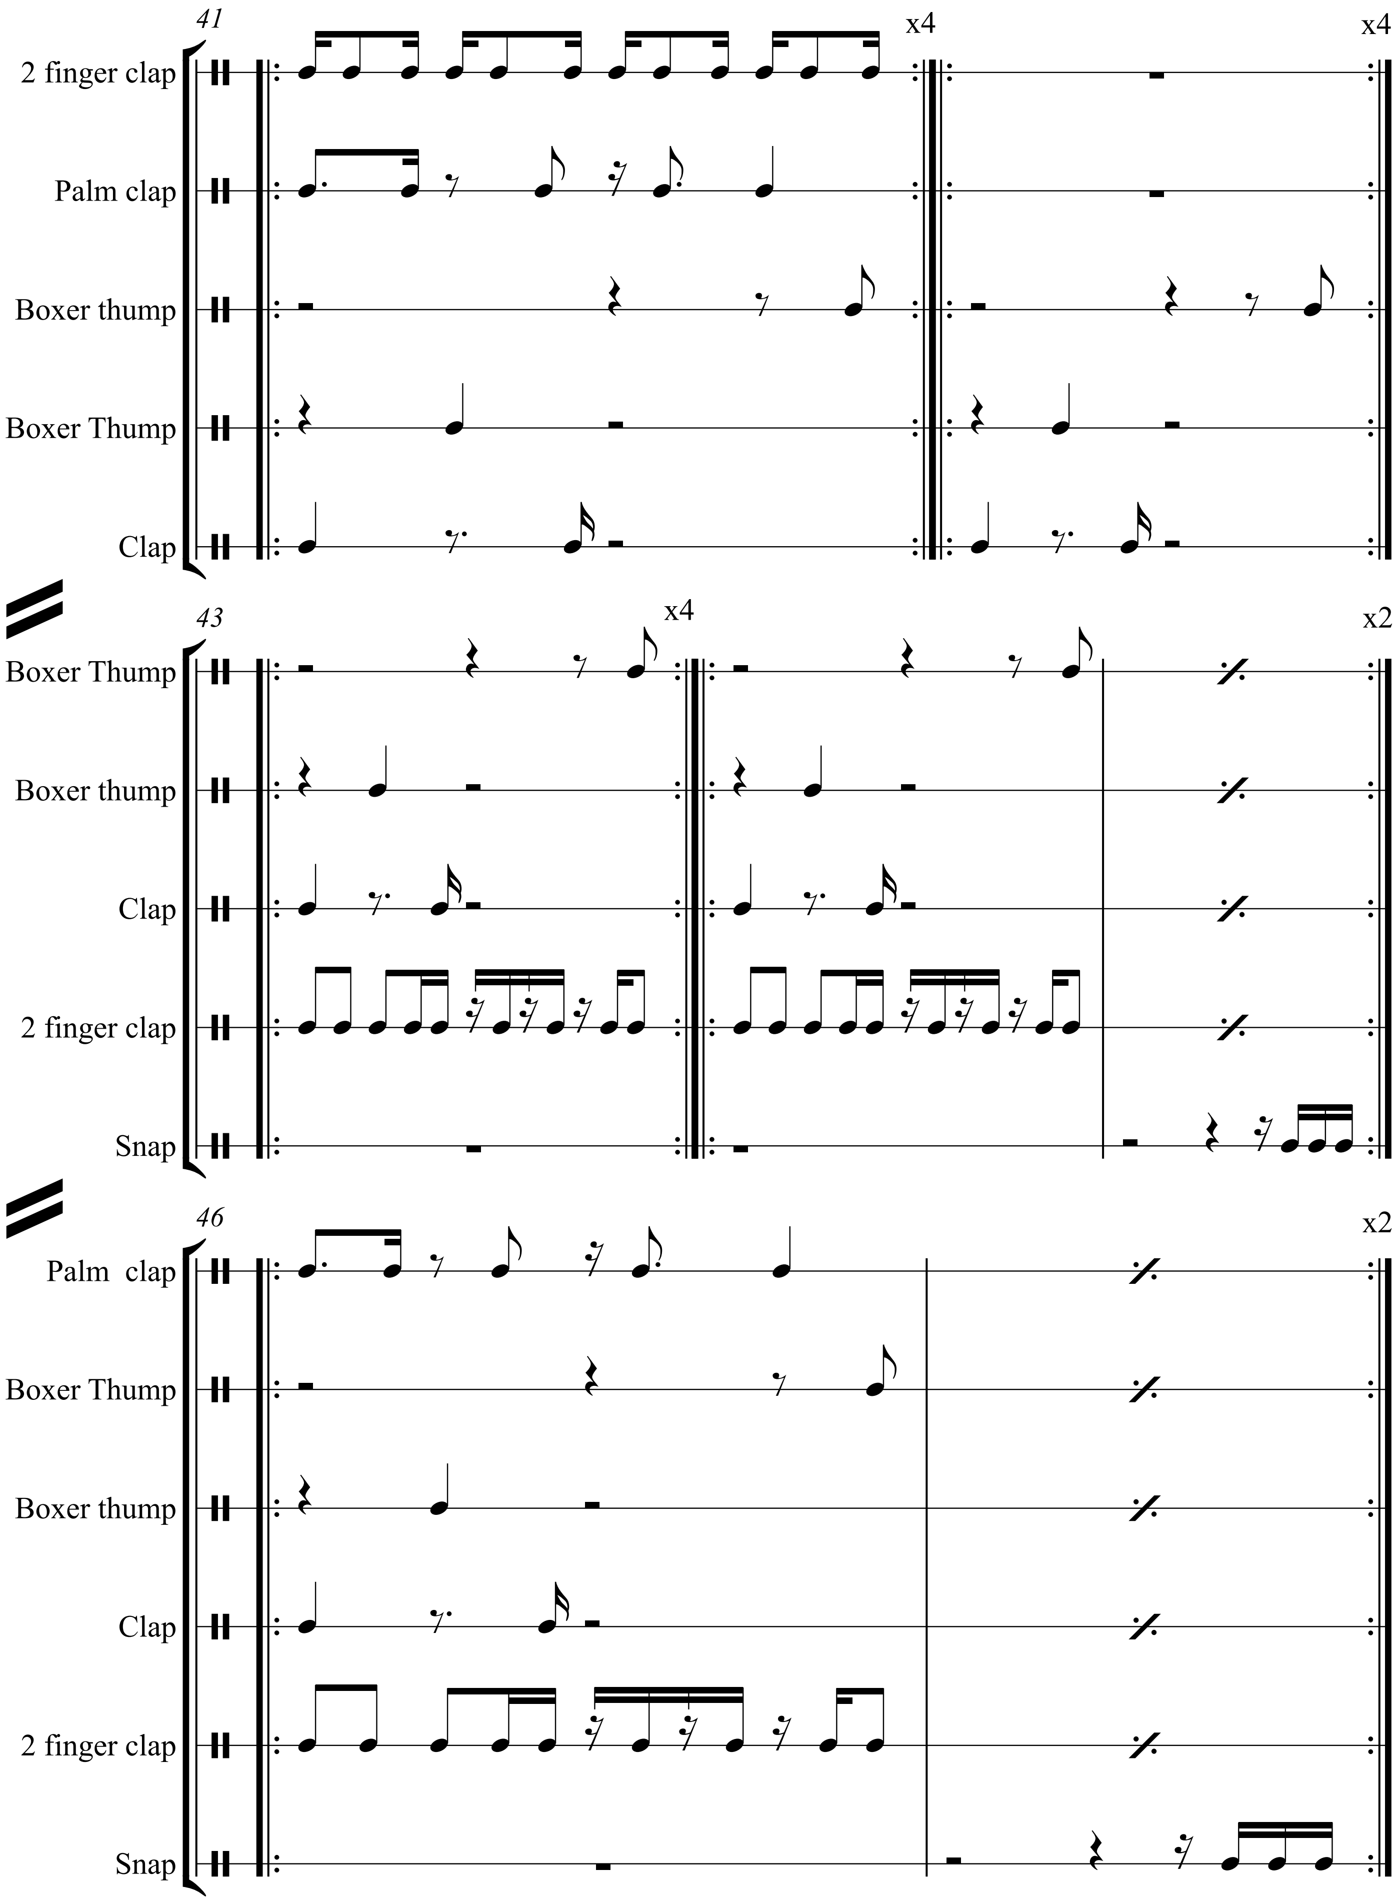
41
2 finger clap
Palm clap
Boxer thump
Boxer Thump
Clap
x4
x4
43
Boxer Thump
Boxer thump
Clap
2 finger clap
Snap
x4
x2
46
Palm clap
Boxer Thump
Boxer thump
Clap
2 finger clap
Snap
x2

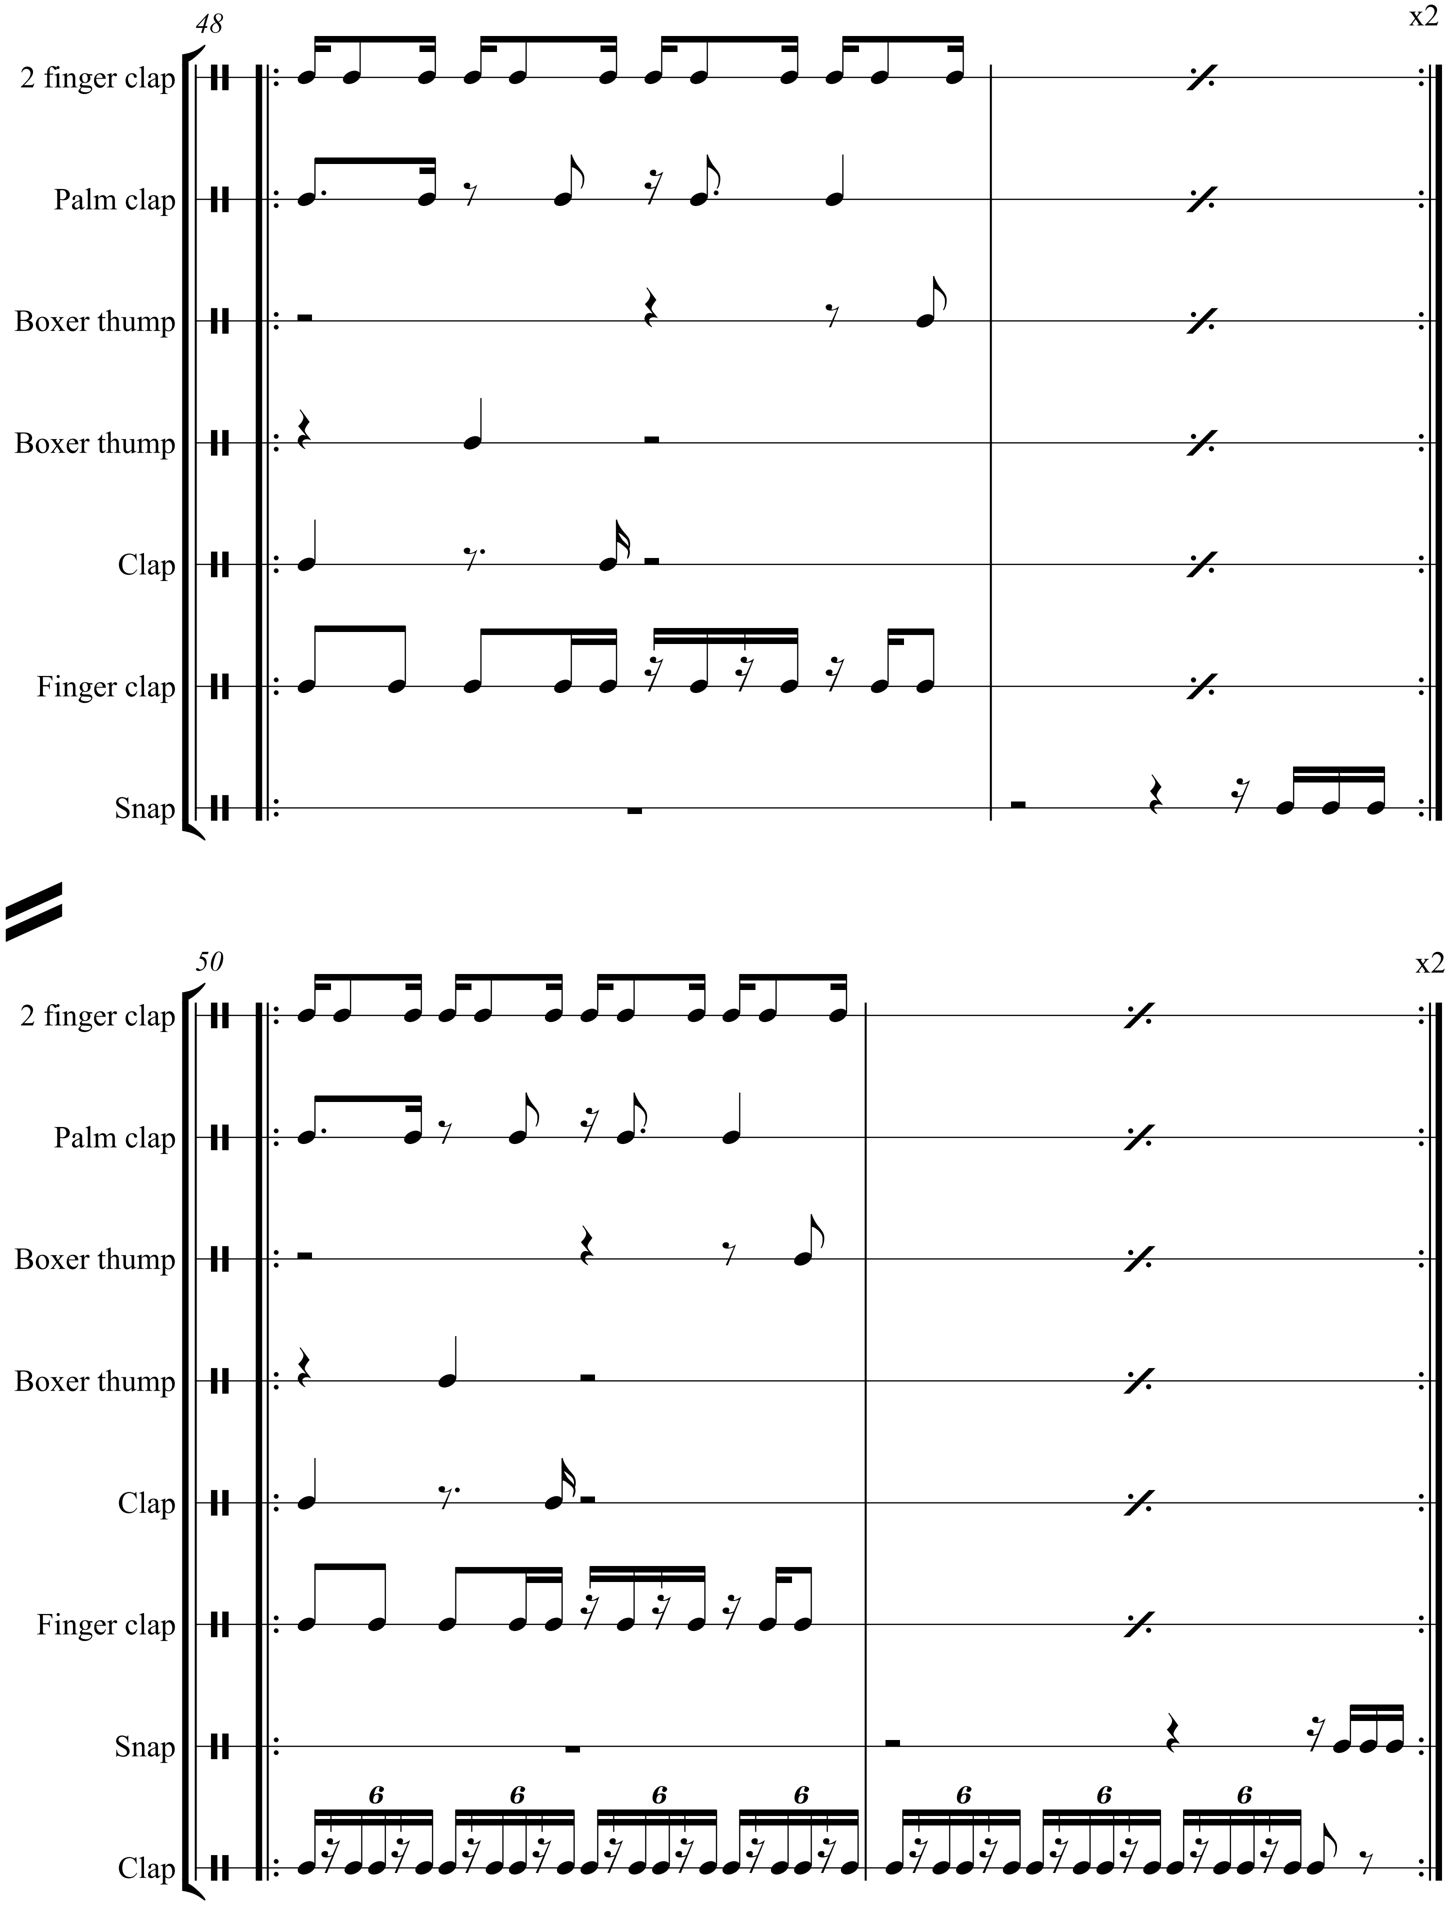
48
x2
2 finger clap
Palm clap
Boxer thump
Boxer thump
Clap
Finger clap
Snap
50
x2
2 finger clap
Palm clap
Boxer thump
Boxer thump
Clap
Finger clap
Snap
6 6 6 6 6 6 6
Clap

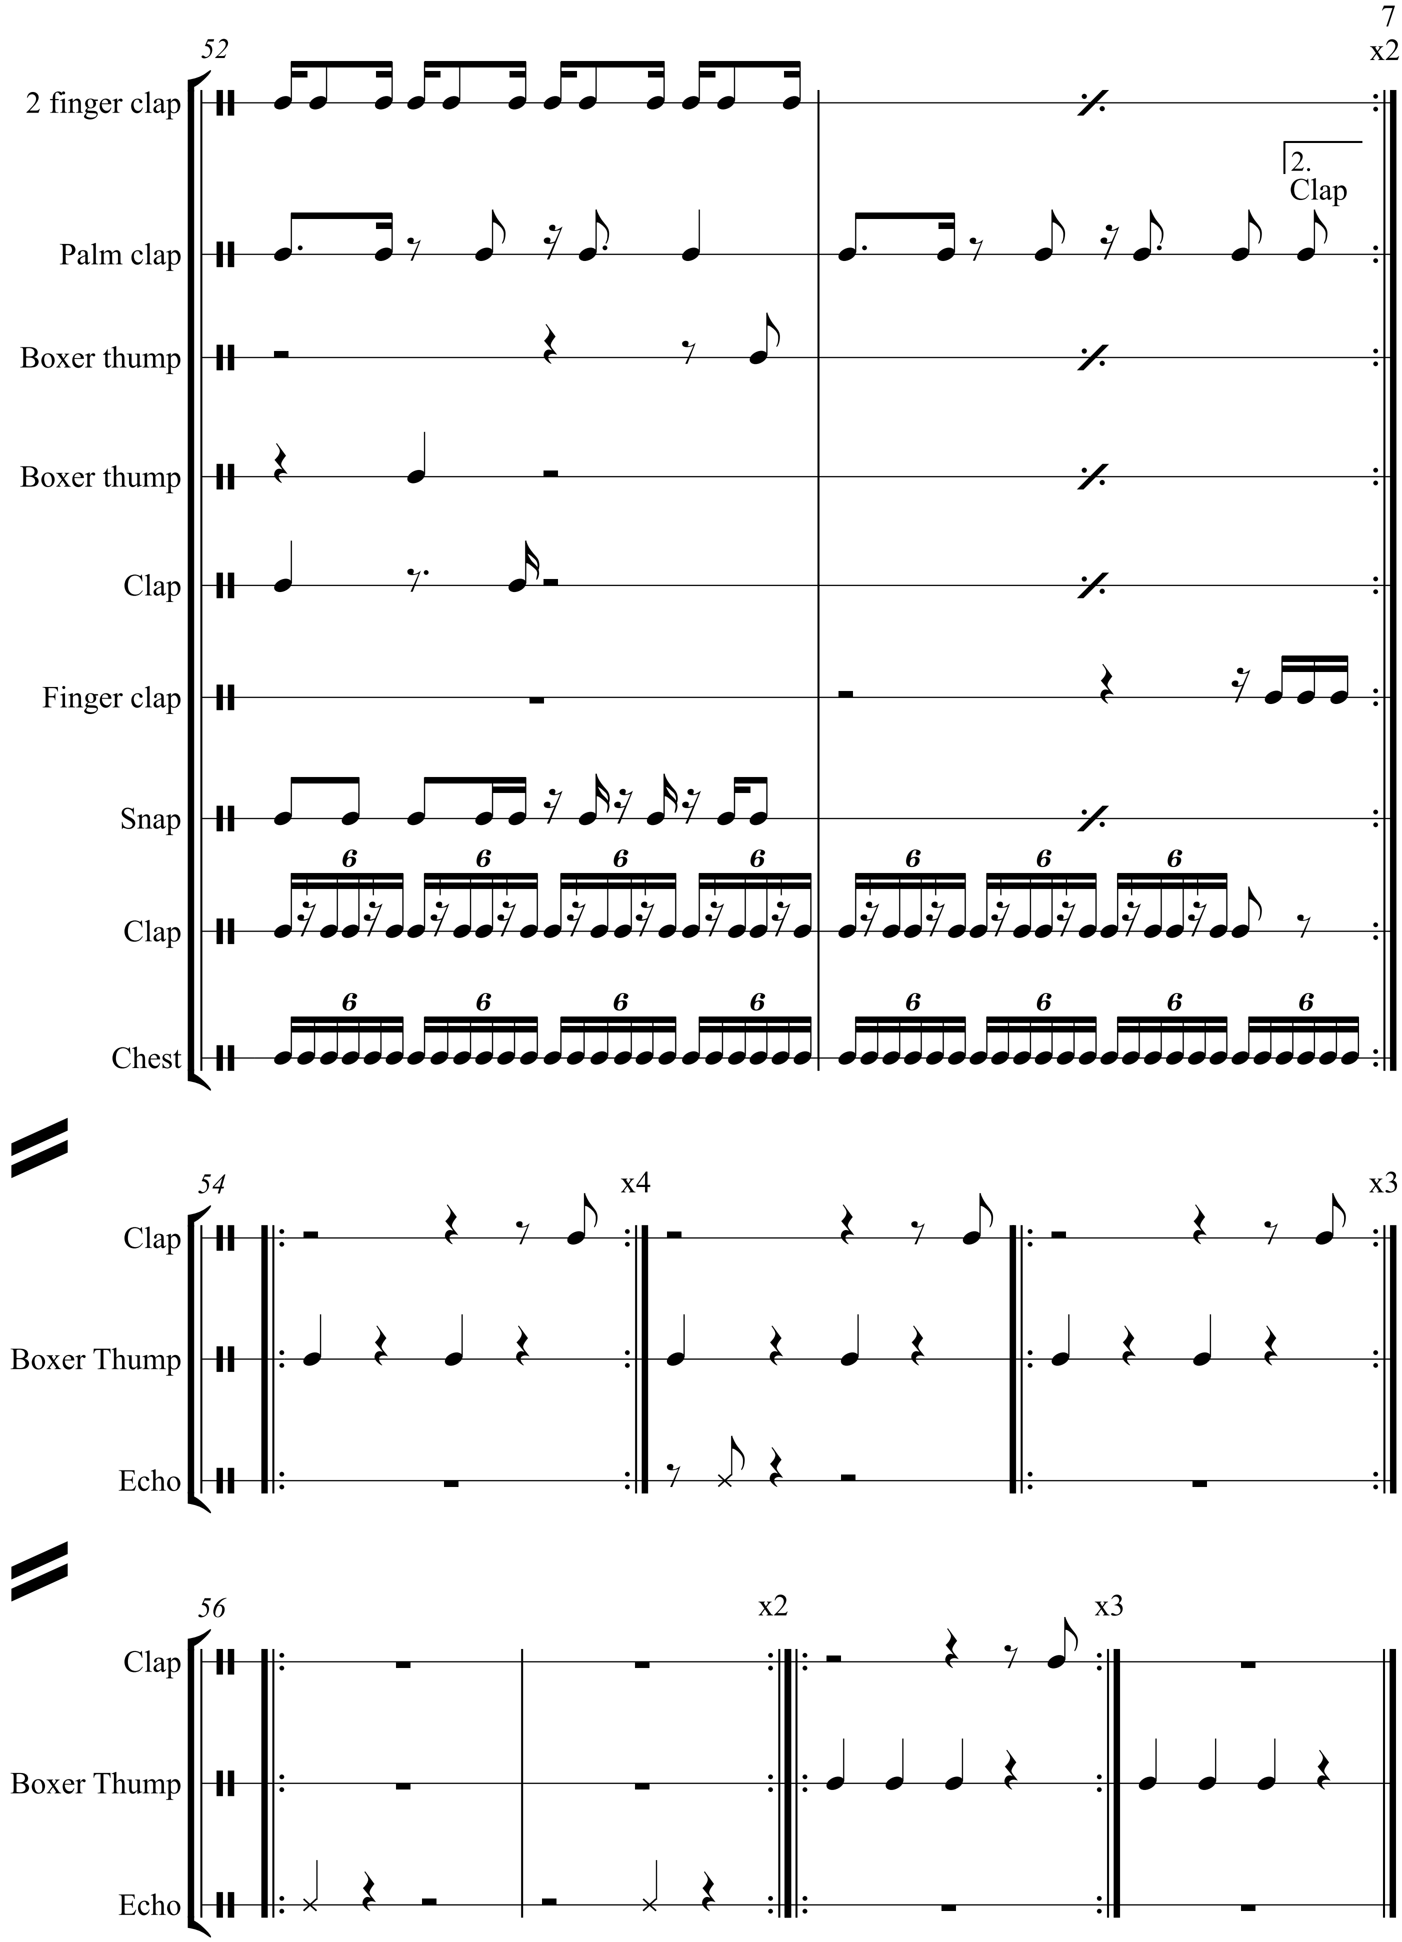
7
52
x2
2 finger clap
2.
Clap
Palm clap
Boxer thump
Boxer thump
Clap
Finger clap
Snap
Clap
Chest
54
x4
x3
Clap
Boxer Thump
Echo
56
x2
x3
Clap
Boxer Thump
Echo