Portfolio of Electroacoustic Compositions with Commentaries

Shane Byrne

Dissertation submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy

Department of Music
Faculty of Arts, Celtic Studies & Philosophy
National University of Ireland, Maynooth

September 2018

Head of Department: Professor Christopher Morris
Supervisor: Dr Gordon Delap

Contents

List of Figures

List of Csound Code Excerpts

Acknowledgements

Firstly I would like to thank my supervisor Dr Gordon Delap for his constant support and guidance throughout my time at Maynooth University. Without his friendship and sage advice this would undoubtedly be an inferior body of work. I would also like to thank Professor Victor Lazzarini for his encouragement over the years, Dr Iain McCurdy for the many inspiring and informative conversations after Hack Club, Paul Keegan, and all the staff within the music department of Maynooth for making these years so enjoyable. There are so many teachers (both academic and otherwise) that I have had the good fortune of coming into contact with over the years, and although there are too many to thank here, I would like to acknowledge their contribution to this body of work.

Dr Patrick McGlynn, Brian Robson, Dr Brendan McCloskey, Simon Kenny, Dr Eoin Smith, Ror Conaty, Kevin MacNamara, Conor Donelan, Rob McComish and Mark James, sincerest thanks to all of you (and others that I have no doubt forgotten to mention) for indulging me and setting aside your time listen to my tunes and for taking part in my often chaotic experiments. I'm greatly indebted to everyone that I ever have played music with, played music for, and to everyone who has shared their music with me. Each and every one of these people have had an influence on this portfolio, some major, some minor, all significant.

Lastly, but by no means least, I would like thank my family, especially my wonderful parents Marie and John, who have been a constant source of support and strength. And of course my eternal gratitude to my beautiful fiancée Dani, you're always there for me.

Abstract

This portfolio consists of electroacoustic compositions which were primarily realised through the use of corporeally informed compositional practices. The manner in which a composer interacts with the compositional tools and musical materials at their disposal is a defining factor in the creation of musical works. Although the use of computers in the practice of electroacoustic composition has extended the range of sonic possibilities afforded to composers, it has also had a negative impact on the level of physical interaction that composers have with these musical materials. This thesis is an investigation into the use of mediation technologies with the aim of circumventing issues relating to the physical performance of electroacoustic music.

This line of inquiry has led me to experiment with embedded computers, wearable technologies, and a range of various sensors. The specific tools that were used in the creation of the pieces within this portfolio are examined in detail within this thesis. I also provide commentaries and analysis of the eleven electroacoustic works which comprise this portfolio, describing the thought processes that led to their inception, the materials used in their creation, and the tools and techniques that I employed throughout the compositional process.

1

Introduction

This portfolio of electroacoustic compositions represents the work I have undertaken over the four years of my doctoral research. It is primarily concerned with investigating ways in which composers of electroacoustic music may use corporeal movement and human presence in general as a tool for both compositional and performative ends.

Ever since Max Mathews developed the first digital recording and digital sound synthesis systems at Bell labs in 1957 (Lazzarini, Yi, et al., 2016), the terms computer music and electronic music have become increasingly linked and today the two terms have become almost synonymous. This is partly due to the ubiquity of personal computing devices and the musical freedom and processing power that they offer to the modern musician. While this has increased the range of musical possibilities available to the average composer of sound, it has also reduced the amount of physical interaction required in order to make music to nothing more than click of a mouse or pressing a button on a QWERTY keyboard, "hence any idea of 'physical' performance is severely constrained" (Smalley, 1997). With the frequency that new technologies for human-computer interaction (HCI) are being developed, I feel that it is time to embrace the physicality of performance in new and exciting ways, celebrating the impact that human presence has in the overall musical experience.

I personally have found it difficult at times to fully empathise with minute performative gestures that take place from behind a computer screen or a mixing desk when attending a performance of electronic music. There are certain meta-communicative cues hidden within the expressions and movements of performers during *traditional* musical performances which help the audience establish a connection with the both musician and the music alike. The power of extra-musical gestures can be observed whenever one has the opportunity to witness a musician performing a piece that truly seems to resonate within them[1]. In cases such as this, the musician seems to become more than an interpreter; they instead begin truly inhabit the music, allowing their corporeal movements and gestures to elucidate the sentiment of the piece. They become a conduit; they embody the music. The correlation between physical gesture and musical gesture is incredibly important in the assimilation of the musical experience, as I will discuss later in this thesis.

The act of embodiment does not only allow for greater communicative potential between the performer and audience, it also grants the composer of electronic music the opportunity to occupy their sonic creations in a completely different way than that of the *traditional* computer musician. The manner in which an artist interacts with the materials they use has a profound impact on the resulting creative output. It has been demonstrated that cognitive states are affected by the body and the environment (Leman, 2008; Godøy, 2003; Risko and Gilbert, 2016; Pfeifer and Bongard, 2007) and I will examine this phenomenon in further detail in the following chapter. Jan C. Schacher highlights the potential for embodiment to enhance technologically mediated creative practises in his essay "The Body in Electronic Music Performance" when he states that,"when choosing workflows that invert the hierarchy between technology and immediate intuitive action, unexpected results tend to emerge. The order in which the elements of a piece are 'composed' can be reversed, letting the desired 'gestural' space inform the musical processes and materials"(2012). If the composer has the ability to embody their work during the early stages of creation, the potential for fluid and intuitive interaction with the musical materials is potentially enhanced, perhaps leading to serendipitous events and unexpected outcomes.

[1] A good example of a musician wearing the emotion of the piece on their face is David Oistrakh's recital of Debussy's *Clair De Lune* which can be found at https://www.youtube.com/watch?v=SKd0VII-l3A

This *physicality of performance* is something that I want to explore within the context of electroacoustic music with the aim of achieving a stronger connection between the performer and the audience, and the composer and the materials they work with. In order to achieve this, I have been experimenting with ways of interacting with computers and computing devices using motion sensors and wearable technologies. The goal of these experiments is to create a synergy between the physical actions of a performer and the musical output of the computing device that is used in the composition of the material.

In the next chapter, I will explore these concepts in more detail and highlight the theoretical frameworks that have helped form the conceptual scaffolding for this thesis. This will include a small introduction to embodied cognition, before I examine the ways in which embodiment has been applied to my own creative practises and how it might be of potential benefit to others. The level of detail concerning embodied cognition in general in this chapter is far from exhaustive, but for the purposes of this thesis it will suffice in order to contextualise my compositional strategies.

The following chapter serves to outline the general methodology that I developed over the course of my research. Here I present a basic usage example of this methodology and also provide a general overview of my compositional practise by separating my method into three distinct stages. I also discuss the merits of an autoschediastic approach toward electroacoustic composition, the merits of celebrating necessary constraints and how I resolved the tension between the roles of composer, performer and technician.

In chapter four, I deal with the ten compositions around which this entire thesis is formed. In this chapter, I go into detail regarding the materials used in these pieces, the compositional strategies employed in their creation, and the conceptual or narrative basis of each piece. Some of the compositions were more heavily influenced by the technologies used during the compositional process than by any external narrative and in others the inverse is true. I have weighted the discourse between technical detail and narrative accordingly in each case.

The main output of this thesis is a portfolio of electroacoustic compositions and accompanying commentaries, however, I feel that in order to fully contextualise my work and the technologies that inform my compositional practise, a further discussion of material not directly related to this portfolio is warranted. For this reason I have included several additional sections in the appendices that accompany this thesis to serve as an addendum to the main body of work.

In Appendix A, I introduce and provide some background to the tools, both hardware and software, which I have used in the creation of the works within this portfolio. Here I give a brief introduction to the operation of Csound, the audio programming environment that I have used extensively throughout this portfolio. I also introduce Processing, a visual programming environment, and discuss the manner in which these programs (and others) communicated with one another throughout my work. I review some of the mediation devices that I used to realise the live performance pieces within this portfolio as well as briefly describing some hardware that I have developed and modified myself.

Appendix B is concerned with the interactive pieces that I created over the duration of my doctoral studies. Although these works do not feature in the portfolio of compositions which accompany this text[2], I feel that their inclusion is important in order to fully contextualise my work in general. These take the form of installations, musical games and custom built instruments, all of which are as much a part of my musical journey as any of the works featured in the portfolio. It was often when working on these interactive pieces that I discovered techniques and devices which would go on to inform my musical compositions.

Appendix C is a brief commentary on *Ouroborus*, the first live electroacoustic performance piece that I composed.[3] Although *Ouroborus* does not feature in the portfolio of compositions, the techniques I utilised during the inception of this piece would go on to have a great influence on my future work.

The final three sections provide a list of the pieces featured in this portfolio, the contents of the accompanying flash drive, and brief summary of the supplementary material provided.

[2] Materials related to these pieces can be found on the accompanying flash-drive

[3] The misspelling of the word to describe the arcane symbol of a serpent eating its own tail is intentional. The initial goal of a seamless feedback system was only partially achieved, hence the partially correct spelling of the the title.

2

Frameworks

> *"All true ritual is sung, danced and played. We moderns have lost the sense for ritual and sacred play. Our civilisation is worn with age and too sophisticated. But nothing helps us regain that sense so much as musical sensibility. In feeling music we feel ritual. In the enjoyment of music, whether it is meant to express religious belief or not, the perception of the beautiful and sensation of holiness merge, and the distinction between play and seriousness is whelmed in that fusion."*
>
> – Johann Huizinga, *Homo Ludens*

During my time composing and performing I have come to understand music as an animate, dynamic entity, constantly in a state of flux, something that cannot be strictly defined or labelled. This is a difficult concept to convey but perhaps it may help to borrow from several observations made by others. In his book *Musicking*, Christopher Small speaks about how problematic it is to try and define music in terms of using the concept as a noun, as a definite article. Small sees this as a futile exercise, due to the fact that music is not one singular entity which can be examined outside of the environment in which it exists. On the contrary, music is a process, or as Small puts it, "the present participle, or gerund, of the verb *to music*[emphasis in original]."(2010, p. 9).

This point of view allows one to consider music not as a static entity, rigid in its definition, but as a process, ever changing and fluid, comprised of a whole host of individual moving parts. Even the musical score, which for some in the tradition of western classical music represents the *purest* form of musical intent, "is best regarded not as an encoded representation of sound, but as a stimulus or provocation for the performer to react to and against"(Croft, 2007).

This philosophy resonated with me and I decided to apply it to my own work. I began to deliberately construct pieces that would invite agents[1] capable of a certain level of plasticity[2] into the compositional process. This became increasingly relevant as I personally felt that there were some shortcomings in the dissemination and live performances of some of my previous works. The manner in which they were presented seemed somewhat sterile; a "performance" usually amounted to no more than someone I had never met pressing play on a laptop in order to begin my piece. The whole experience was lacking in human presence and devoid of any spectacle or context. While techniques such as spatial diffusion do offer the composer of electroacoustic music the opportunity to interact with their musical output in a live performance setting, I still felt that the full potential of human agency to influence the performance of my electroacoustic works was not being realised.

I began investigating possible techniques that I could employ to enhance the overall musical experience created by my work, not only for the sake of the potential audience but for my own interests too. I felt that there was some untapped potential for the human body to be used as an expressive tool in my music and after coming across pieces by pioneers of electroacoustic music performance such as Michel Waisvisz(2004) and Laetitia Sonami(2000), I was sure this was the direction I wanted to take in my own work.

Having performed live music for many years, I knew there was something ineffable about the act of performing and of viewing a live show which set it aside from the act of reduced listening. I began to ask myself why exactly did certain pieces sound better live? It was when I began to delve deeper into some theoretical frameworks supporting

[1]It may be useful to define exactly what I mean by the word agent. For my purposes it is sufficient to use the definition provided by Simon Emmerson which is: "An *agent* is an entity(a configuration of material; human animal or environmental) which may execute an *action* (a change in something, usually involving a transfer of energy)"(Emmerson, 2016, p. 3)(emphasis in original)

[2]I use the term plastic here in the context of psychology, the ability for something to be malleable, dynamic and adaptive

this philosophy that I discovered the work of Marc Leman and his writings on embodied music cognition. The works of Leman not only served to confirm the importance of human presence in the performance of music, but also provided scientific evidence of the effect that viewing a human articulating a musical phrase through physical effort has on the brains of observers. I will now discuss how the work of Leman and others has helped my understanding as to how the body can be used as a musical tool, not only in the context of live performance, but also in the act of composition.

2.1 Embodied Cognition

In his book *Embodied Music Cognition and Mediation Technology*, Leman raises the point that emotive and expressive musical performance is heavily based around empathetic understanding between the performer and the audience. When enacting musical ideas through bodily movement we add expression to the movements themselves. Even though the performer may not be actually experiencing a certain emotion outright, the physical activity of acting out that emotion actually induces some of the effects of the emotion upon the actor(2008, p. 44).

Although the emotion produced by the performer can be perceived as being authentic, it affords the person experiencing it the luxury of not actually undergoing the same stimuli that would normally be associated with that emotion. This is a phenomenon that many will be familiar with in terms of the pursuit of the emotion of fear; people watch horror movies or go on roller coaster rides in order to experience the emotion of fear without actually fearing for their lives. On a certain level this experience is nothing but a simulacrum for the actual emotion. To quote from Leman:

> To have empathetic feelings with the other, the self must identify with it but, at the same time, must detach itself from it (Decety and Jackson 2006). From this point of view pleasure may result from shared representations of action and perception (ibid., p. 44)

Leman then goes on to reference *Gestaltung und Bewegung in der Musik* in which the author Alexander Truslit makes the point that in order to experience music in its totality, it needs to be accepted that the driving force behind music is the expression of inner movement. When a musical idea is presented in a manner that evokes the emotion and intent of the piece, it can cause the observer to physically react in a manner that mirrors that emotion. Not only do the gestures of musicians performing with sounding-body instruments[3] manifest this inner movement into corporeal motion, but the observer too, through the principle of mirroring and empathy, acts out these physical movements, thereby achieving a greater involvement with the piece.

The multimodality inherent in the musical experience is discussed at length by Leman in *Embodied Music Cognition and Mediation Technology.* Going beyond the novel cognitive dissonance demonstrated by experiments such as the McGurk effect[4], where the influence that visual stimulus has on aural perception is explicitly demonstrated, Leman highlights the importance of certain neurons within our brains that encourage understanding and empathy through mimicry.

2.1.1 Motor-Mimesis

These neurons are known as mirror-neurons, and were initially presented in the paper "Understanding motor events: a neurophysiological study" (Pellegrino et al., 1992). The studies in this paper demonstrated that motor neurons in the brain of a monkey became active when watching another monkey performing a manual task. This activity was observed in a particular set of neurons known as audio-visual mirror neurons, which become activated not only when observing a sound-producing action but also when hearing the sound which that action may have produced. This discovery of motor-neurons not only bolsters hypotheses such as the motor-theory of speech perception (Liberman and Mattingly, 1985) but also demonstrates how important the physicality of musical performance actually is when viewing music as a tool for communication.

The sympathetic firing of these neurons in the brain of the perceiver is known as

[3]In contrast to instruments that are electronically or digitally constructed, sounding-body instruments require some level of physical interaction in order to cause them to vibrate, thus producing sound.

[4]The McGurk effect demonstrates that through the act of coupling the audio component of one sound with the visual component related to a second sound, the observer will perceive a third sound. This experiment is commonly performed using the audio component of the phoneme *ba* and the visual component of someone mouthing the phoneme *ga*. The result is that the observer will perceive the phoneme *da*

"resonance behaviour"(Vickhoff, 2008). This behaviour could perhaps offer up an explanation as to why musical intent and conveyance of emotion is often stronger when viewed in a live context and through multiple modes of perception. But how might the mechanical movement of a human body elicit an emotional response or subjective experience in the observer? Leman posits that the transduction of corporeal motion into intention is realised through "the coupling of action and perception"(2008, p. 93). Certain corporeal movements are bound up in social and cultural concepts, either through past experience of the gesture, the manner in which the gesture is articulated, or through mimetic qualities inherent within the gesture. The measured corporeal movement of a performer can imply a range of meta-communicative statements, the semantics of which may not necessarily be pan-cultural. Although there is certainly a lot more to unpack here in terms of discerning what a certain physical gesture may mean to a particular observer and what ontological significance it may have in a given context, these are questions that go beyond the scope of this thesis.

Rolph Inge Godøy discusses this phenomenon further in his paper "Motor-Mimetic Music Cognition", in which he refers to the resonance behaviour previously discussed as "the motor-mimetic elements of music cognition"(2003, p. 317). This theory draws upon the idea that we have contextual experience of how a certain sound may be created, therefore when we hear that particular sound, even though the source of the excitation may be hidden from view, our motor neurons mimic the process that would be required for us to exert force on the imagined excitation source and thus create that sound. This can almost be likened to an air guitar or air drumming that takes place within our minds upon hearing a familiar rock song.

Godøy states that this form of mimesis not only helps us identify the possible source of the sound and comprehend the reality of the action, but it also helps us engage with and literally empathise with the action needed to create said sound. As Godøy puts it so succinctly, "any sound can be included in an action-trajectory"(ibid., pp. 317-318). With this in mind (along with the previously mentioned coupling of action and perception) I will briefly discuss how I deal with this action-trajectory in the performance of my live electronic compositions.

2.2 Source Bonding

In his paper "Spectromorphology: explaining sound-shapes", Denis Smalley echoes the aforementioned perspective of Small when he states "a piece of music is not a closed, autonomous artefact: it does not refer only to itself but relies on relating to a range of experiences outside the context of the work"(1997, p. 110). In this paper he also draws the readers attention to the importance of what he calls "source bonding", and how it is an inherent part of the perceptual practice of an observer; if a sound object is heard, it is natural for the listener to imagine the physical process of how such a sound may be created. This is historically rooted in the fact that before the advent of electronic means of sound production, all sounds had to be created using a certain amount of physical effort on the part of the performer. In an attempt to tap into the motor-mimetic potential previously discussed in this chapter, I use what Smalley terms "bonding play"(ibid., p. 110).

In my live performance pieces *Ouroborus*, *Kinesia*, *Proprioception* and *Conatus*, I create an alternative means for representing the actual source of the sonic material through use of physical gesture. The notion of replacing a concrete musical action with another gesture which is removed from the actual sounding source is referred to by Smalley as "gestural surrogacy"(ibid.). The particular approach of using the human body as a surrogate serves to qualify both human presence and physical gesture within each of these works.

For example, in *Proprioception*, the act of moving both hands downwards in a swift motion creates a loud, harmonically rich sound object that can be described as having an attack-decay trajectory (ibid., p. 113). The actual physical effort required in order to produce this sound object is nothing more than the effort required to press a key on a QWERTY keyboard, however, in order to convincingly perform the creation of this sound object with a view to excite the mirror neurons of the audience, I use the physical gesture of moving both hands in a swift downwards motion to bind the audio output to the physical gesture required to create it. The inherent energy-motion trajectory required to create the sound object is substituted for a more performative and physically engaged one.

2.2.1 Articulations

The decisions I have made when composing the live performance pieces of this portfolio were all made with a view to enhance the perceptual experience of the performer and audience alike and perhaps to add a level of authenticity[5] to the piece. Years of culturally embedded and experiential knowledge concerning the sounds produced as a result of certain action trajectories applied to certain sounding objects cannot be ignored. The act of making sound often requires a sound-making action.

This leads me onto discuss three forms of music related actions, as posited by Godøy in his essay "Chunking in Music by Coarticulation", which can be observed within almost all live music performances. The first is that of the *sound-producing action*; a pianist strikes the key with his finger and a note is produced. The audio output from the instrument is a direct result of physical interaction with the sounding body. This particular kind of action includes bowing, blowing, striking, picking, rubbing and any other means of physical interaction with an instrument that results in the vibration of a sounding body. For the purposes of this thesis I will extend this definition to include interaction with digital and electronic instruments, where the relationship between effort input and audio output is slightly more arbitrary, due to the fact that all musical signals and triggers eventually are reduced to nothing more than a Boolean binary. In this case, a sound-producing action could simply be the pressing of a button, the wave of an arm or essentially any HCI mediation technique that one chooses to use as an interface with a digital instrument.

The second of these music related actions is the *sound-modifying action*. An example of this would be a pianist pressing down on the sustain pedal with her foot, raising the dampers from the strings within the body of the piano, allowing each string to freely

[5] I refer here to the definition of authenticity as used by Walter Benjamin in *Art in the Age of Mechanical Reproduction*. In this context, the word refers to the ability of a work of art to occupy a specific place both spatially and temporally. The compositional systems used within the pieces I refer to are designed to be subject to multiple agents of change in an effort to create works which are unique in some way each time they are performed

vibrate and resonate, thereby modifying the sound emanating from the piano. The act of pressing down on the sustain pedal did not create any sound [6] but acted in conjunction with a sound-producing action to modify the overall sonic output. Once again, in the context of digital and electronic music, this modification does not necessarily have to follow a set of physically determined rules and any HCI strategy can be used when designing a means for modifying audio output from a particular compositional system[7].

The third music related action described by Godøy which I feel is relevant to discuss in this thesis, is the *sound-accompanying action*. This action does not create sound or manipulate any sonic aspect of the music whatsoever. The sound-accompanying action only has relevance in the context of meta-communication, beyond what is written in our hypothetical pianist's sheet music. Sound-accompanying actions include any physical gesture or movement used during the act of sound production that do not directly impact the sonic quality of the output. Moshing, dancing, posturing and the use of exaggerated facial expressions all fall within this category. The non-musical gestures provided by a performer serve to enhance the total musical experience either by framing the act itself as performance (Schechner, 2015), facilitating meta-communication between performers (Carson, 1996; Glassie, 1982) or simply to enhance the effect of a particular performative or musical gesture.

I will now examine three gestures from one of my own compositions, *Proprioception*, and discuss how they can described using the taxonomy posited by Godøy. In the first minute of the piece the act of aligning both of my hands along the y-axis creates a low rumble that heralds the beginning of the performance. This, when allowances are given to the fact that this is a digital instrument and therefore has no sounding body, is an example of a sound-producing action. Although no significant mechanical vibration or excitation took place as a result of aligning my hands, that particular posture was the trigger that caused the low-rumbling instrument to initialise.

---

[6] At least not any sound likely to be intended by the composer, but there will of course be exceptions to this for all sound can be considered valid in the right context.

[7] I use the term compositional system to describe the assemblage of code (usually compiled with Csound), technology and materials that I use in order realise a particular musical work. The compositional system is neither the finished musical piece nor the individual components that were used in its inception. It is a system which facilitates the interaction of multiple agents during the compositional process.

In the third section of the same piece, I cast my hands out in front of me, moving both hands alternatively back and forth. This not only modulates the grain density of the sound object, but this physical gesture also has an impact on how the audio output is situated within the stereo field. Due to the fact that no new instrument or sound object was initialised by this gesture, but it is indeed linked to the modulation of several musical parameters, this can be considered a sound-modifying action.

Shortly after this gesture is performed, I once again bring both hands out in front of me, hold them at arms length away, level with my chest, and proceed to bring them back toward my body, gradually increasing the distance between them. As I am doing this, the conceptual narrative of this piece suggests that I am *summoning* these sounds and in an attempt to control these conjured entities I shake my hands as they move, alluding to the fact that the act of controlling these sonic elements is difficult. This shaking gesture has no impact on the compositional system whatsoever but it does serve to add to the spectacle and the performative quality of the piece, thus providing an example of a sound-accompanying action.

2.2.2 Effort

There is another aspect to this particular sound-accompanying action that I would like to discuss further, and that is the allusion to effort. Each time I perform this particular passage of the piece, I have become increasingly aware of the muscle tension in my arms and legs, and I also observed in video recordings that my face wears the expression of deep focus. The way in which I hold my body throughout this section is an attempt to consciously address the lack of physical effort required to produce audio when using the compositional system of *Proprioception*. Effort and work are inherent in the production of energy and acoustic energy is no different. However, due to the fact that digital audio does not require a large amount of physical effort in order to be initiated, I felt that the energy-motion trajectory of this particular gesture needed to be exaggerated somewhat in order to enhance the spectacle of that particular musical passage. The link between effort and music production has been discussed at length, particularly from the point of view of electronic instruments (Ryan, 1991; d'Escriván, 2006) and although it has been suggested by d'Escriván that "the widespread use of computers and computer interfaces sets the scene for a new way of appreciating performance skills"(2006, p. 190), established paradigms based on high levels of dexterity and physical effort being equated to superior musical output still seem to prevail, for now at least.

2.3 Embodiment as Process

My reasons for investigating the means for technologically augmenting the performance and composition of electroacoustic music can also be understood through the lens presented by Milhaly Csikszentmihaly in his work *Flow: The Psychology of Optimal Experience*. In this book, he specifically addresses how one might engage as many sensory stimuli as possible, in my case this means exercising not only the compositional mind but also the body, in order to pursue a heightened musical experience. Much like the ancient practice of Yoga, my goal here is to use the body as a means for controlling what happens in the mind (Csikszentmihalyi, 2009, p. 105). To elaborate a little on this point, the idea is that if I as a composer or performer can engage with a musical passage through as many modes of interaction as possible, to see the passage from as many view points as possible, then I can fully engage with the task at hand and become completely immersed within the passage in question.

In terms of performance, Csikszentmihaly makes the point that pre-recording era music carried with it some sort of mystical quality that can be linked to a time when its ritualistic uses may have been more common. This ritualistic appreciation of the performance and the spectacle of music can often encourage a Durkheimian "collective effervescence" (Durkheim, Cosman, and Cladis, 2001), a collective experience that thrives in ritualistic settings. These rituals need not necessarily be religious gatherings or strictly spiritual practices, and the framing of the ritual has been extended to include music festivals and concerts (Murphy, 2011).

I believe that these collective experiences are particularly potent during a performance that not only provides the audience with the spectacle of performance but also invites the audience to empathise with the performer. By this I mean that the performance should be somehow relatable to the audience. They should be able to understand (even on a subjective level) the physics of the task that is being performed, the conceptual basis behind the performance or the emotive intent that is driving the performance. Perhaps the ideal performance should include all three?

2.3.1 Cognitive Offloading

Live performance is not the only context in which I use my body as a musical tool. When composing and compiling material for my fixed media pieces, I often use the same compositional systems that I use for live performance pieces such as *Ouroborus*, *Proprioception* and *Kinesia*. When developing the sonic landscape of fixed media pieces, I often find myself using my body to articulate and act out the musical gestures that I had envisioned fitting into the composition at hand. Somehow this corporeal engagement, even though it is merely serving as a *surrogate* to the action required to create a sound object as opposed to having influence on the actual manifestation, helps to solidify the compositional narrative and aids in developing the desired flow of the piece. It helps give form to concepts which initially are formless. One reason that this practise may lend itself to the manifestation of compositional ideas lies in the actual embodiment of these musical gestures.

To provide one general example of how this corporeal engagement can aid in my compositional process, I often use it to combat difficulties I sometimes find when attempting to construct satisfying gestural cadences within my tape pieces. When creating music on a sounding body instrument, the action required to produce any sonic material would inevitably influence the manner in which the resulting musical phrase would be articulated. The physical constraints and affordances associated with the particular sound-producing action required to make the phrase would inherently influence the shape and cadence of the overall gesture. The necessary restraints which accompany music making on sounding body instruments are often not present in the composition of electroacoustic music, meaning that some inherent cues as to when the piece (or the performer) needs to take a breath are not as explicit. The act of corporeally representing the intended musical phrase serves to solidify abstract concepts within my own particular practice and naturally provides indications as to when the piece should breathe. In phenomenology the process of physically manifesting a cognitive process such as the the act of composition is known as *cognitive offloading*.

Cognitive offloading is a term given to the act of physically representing a mental task or idea (Risko and Gilbert, 2016). For example, the act of counting on one's fingers is a simple form of cognitive offloading. This can also be seen when one documents ideas on a piece of paper in order to capture abstract thoughts and ideas with a view to add a degree of coherence to those thoughts. Cognitive offloading has been attributed to extending our mental capacity of conceiving complex tasks or abstract thoughts (Pfeifer and Bongard, 2007). Unbeknownst to me at the time, this was exactly what I was doing when I began to compose electroacoustic music at an undergraduate level. However, the form that cognitive offloading took in this context was not one of notating or scoring musical events as is the norm in the practise of traditional western music composition. My form of cognitive offloading was to try and give physical shape to the musical gestures I was preparing for my fixed media compositions. I would try and perform physical gestures with my limbs that would in someway reflect the character (both spectral and dynamic) of the sonic-object that I was representing.

To provide an example of this style of performative composition, consider a sound object that begins as a low rumble and slowly evolves into a spectrally rich texture. This musical gesture could possibly be physically represented by the composer/performer crouching down with both hands extended out in front of him. As the sound object evolves he may begin to straighten his legs and slowly raise his arms in an upward and outward motion. The result of this process is that the composer has just provided a physical analogy for the sonic development of a musical gesture. This cognitive offloading not only allows the composer to engage with the material across multiple modes, but it also provides an outwardly manifestation of the "inner movement" (Leman, 2008, p. 44) that I have posited as the driving force behind music.

This particular example may result in the slowly evolving texture being punctuated by a loud impact sound to bookend the textural evolution. If this was the case, the composer may have perhaps thrust both hands out to his sides in a short, fast movement and then possibly retain this posture for an extended period of time. The retention of this posture would indicate that the impact which served to bookend the evolving texture should be allowed to decay and give the piece time to breathe before the next musical gesture is made. It is this form of offloading that I used most frequently when first approaching electroacoustic composition and it allowed me to give form to that which at first had seemed abstract and formless.

2.4 Performance as Process

Even when creating material without an audience present, there is still some level of performance taking place during my compositional process. This indicates that any separation between composer and performer becomes blurred when using compositional systems that require at least moderate amounts of corporeal interaction. This sentiment was also expressed by STEIM founder Michel Waisvisz in an interview published in the *Computer Music Journal*; "The term 'electronic music composer' implies being a performer as well [...] I think that a composer has to be able to make immediate compositional decisions based on actual perception of sound rather than making decisions derived from a formal structure that - as happened in serialism - tends to drift away from our pure musical needs" (Krefeld and Waisvisz, 1990, p. 28).

In my own experience, the process of composing a piece of electroacoustic music is rarely straight forward. There are no clear start and end points where concept becomes process and process becomes art. The initial idea that will become the catalyst for a piece can come from a whole range of sources such an instrument played in an unusual way, an interesting narrative, using a piece of equipment in a manner that it was not intended, or combining a strange assemblage of sounds captured on a mountain hike.

I generally spend quite a bit of time assessing the manner in which I will interact with the materials I am using. For example, if part of the impetus to compose a piece is to test the functionality of a certain piece of custom-built equipment, (as was the case with the joystick instrument in *Disintegrate*) some of my compositional choices have been made in so much as some of the sonic materials and the means for interacting with them have already been established. This sets the compositional process off in a particular direction, even before any sounds have been produced. In other cases, the direction of the work is uncertain and I will try out a range of different tools until I am satisfied with how the piece is evolving through the use of those tools. This was the case when composing *TenterHooks*, where the progression of the piece was directly linked to the tools being used during its inception (see chapter 5.10).

The manner in which I work revolves around two aspects of the compositional process which Waisvisz refers to as the *formal structure* and *the physicality of performance*(Krefeld and Waisvisz, 1990). The formal structure represents the analytical mind, concerned with the implementation of rules, patterns and the precise delineation of stylistic choices. The physicality of performance refers to a more intuitive manner of thinking, where ineffable elements of composition, such as feel and flow, come to the fore. It is in this phase of the compositional process that unexpected and serendipitous events tend to occur, mistakes result in the illumination of new creative paths that would have never been discovered if one was to only use the analytical part of their mind.

2.4.1 Formal Structure

In terms of engaging with formal structure[8], what I tend to do is create generative engines, the function of which is to create a large body of sonic material. These generative engines often comprise of multiple file players constructed in a very intentional way; the level of control exerted on musical parameters such as dynamic and timbral evolution tends to be quite measured. In order to remove a certain level of sterility to this process, I tend to incorporate pseudo-random[9] operations within these generative engines. The importance of adding some level of capriciousness to digitally created sounds can not be

[8] I refer here to my general practice of organising musical passages and the means for generating musical gestures. The use of the term is informed by Michel Waisvisz's definition; "as a synonym for the beauty of patterns, cold reasoning, a law abiding mind and dogmatic thinking"(Krefeld and Waisvisz, 1990)

[9] I use this term to describe the algorithims used by computers to create the illusion of randomness as opposed to using the term in reference to weighted chance operations.

stressed enough, as the ear tends to notice when something sounds too perfect, "for form unvitalized by spontaneity brings about the death of all the other elements of the work" (Cage, 1961, p. 35). Unlike some of the aleatoric work of John Cage and the *Music for Magnetic Tape Project*, I am not content to use chance operations to organise macro elements of my compositions, such as the structural evolution or the ordinal progression of sound objects[10].

The function of these pseudo-random processes in my work is merely to add a level of coarseness and unpredictability to the overall audio output in a manner that seeks to mimic the nuances present when a human articulates instances of sound on a sounding body instrument. Untreated, digitally processed audio, can often sound too refined for my taste. The indeterminate elements of the generative engines do not have a significant affect on the overall shape of the resultant sound object, they instead cast the object with a slightly variant hue. A good analogy for this process would be to compare it to a photographic mosaic. It is possible to replace one singular frame within the collage for another frame without distorting the overall master image. The result may be a slightly different shading of that frame or pixel, but the master image is still recognisable[11]. The pseudo-random processes which I use are strictly bound within specific ranges, therefore I still retain a large amount of structural and timbral control.

When speaking about formal structure, I do not necessarily refer to the application of serial techniques or similar deterministic processes[12]. The rules of which I speak may be as simple as: the overarching structure should make reference to a certain narrative, or the source materials can only come from objects that display particular physical attributes for example. This is the first step into a process that very quickly seems to take on a life of its own. Once even the smallest of rules has been conceived and acted upon, the compositional journey has begun. Often, if not always, the final form of the compositional process will remain nebulous until the work nears completion. Once source material has been gathered, as per the rules of the compositional game, and the first examinations

[10] I refer specifically here to works such as Cage's *William's Mix* and *Imaginary Landscape No.5*

[11] This analogy can also be used when describing the relationship between improvised and composed elements of *Proprioception* and *Conatus*. The overall structure is retained but there is scope for variance in the individual parts that comprise that macro structure.

[12] However, I did play with some serial techniques in *Sinmara*.

and arrangements take place, the compositional process tends to evolve in an organic fashion. The musical material itself begins to dictate the direction of the piece.

2.4.2 The Physicality of Performance

The second of these aspects Waisvisz spoke of, the physicality of performance, is a strategy that invites play into the compositional process, ensuring that there is scope for unexpected things to happen when working with sonic material. For example, I often assign a musical parameter or modulation parameter to a sensor of some sort. This immediately changes my interaction with the material from the arbitrary movement of a graphical user interface (GUI) slider to a method that feels more intuitive. This can take the form of mapping my hand gestures to a certain parameter via the LEAP motion sensor (see Appendix A.5) or my whole body using the Xbox Kinect (see Appendix A.2). It may simply be the act of mapping parameters to a series of potentiometers. The advantage of mapping musical parameters to physical components is that while there may only be one level of interaction when using a mouse (one can only move one GUI dial at a time), several parameters can be accessed simultaneously when there is a means for physical interaction and the interface has been thoughtfully designed. It is through the act of mapping physical agents to musical parameters that less calculated, serendipitous events tend to occur.

2.5 Inspiration

In this section I will provide a brief introduction to the work of four composers that have developed their own means of facilitating corporeal engagement within their musical compositions. As this is only a brief overview I will not discuss or critique their work in any great detail, instead I will use this opportunity to situate my own work in the context of their apparent live electronic performance strategies and briefly discuss any similarities or differences that exist between the output of these composers and my own work.

2.5.1 Laetitia Sonami

Sound artist Laetitia Sonami has been using an instrument of her own design which she calls *The Lady's Glove* for performances of her music since 1991 (Sonami, 2010). Sonami makes clear her motivations when developing The Lady's Glove stating that "the intention in building such a glove was to allow movement without spatial reference, and most importantly to allow for multiple, simultaneous controls" (ibid.). This sentiment resonates with my own work in that it explores the potential for coarticulation to circumvent the discrete procedural process of making computer music. Musical gestures are not restricted to a *one-after-another* approach of implementing control over compositional choices (I have spoken about this when discussing the idea of coarticulation in chapter 2.2.1) which allows for interaction with the musical materials in the same intuitive manner that takes place when one interacts with a traditional instrument. To quote Sonami directly; "the sounds are now "embodied", the controls intuitive, and the performance fluid. It has become a fine instrument"(ibid.).

Although the glove I use for *Conatus* is quite different in both design and application, I feel that it too provides a means for intuitive and fluid performance. The notion that Sonami is actively involved in the construction of her own tools, therefore increasing her sense of agency when using these tools, was a revelation for me and was certainly a catalyst for seeking out ways in which I could create my own tools to enhance my compositional practise.

2.5.2 Rajmil Fischman

Rajmil Fischman made use of the P5 glove to realise his live electronics piece *Ruraq Maki* (which translates to hand-made)(Fischman, 2012). The P5 glove is a wearable peripheral device that was designed to facilitate gestural control for applications such as games and educational software(Mindflux, 2002). This glove gave Fischman the ability to interact with his Manual Actions Expressions System (MAES), through which he could create and control sound objects via the use of predetermined hand gestures.

This piece was influential in that it provided an insight into how one might move away from live performance pieces which rely heavily on improvisatory passages and begin to formulate systems which followed very measured and precise direction. It was also one of the first live electroacoustic performance pieces that I came across which placed a strong emphasis on source bonding(Smalley, 1997), or Fischman refers to it, "the intuitive metaphor"(Fischman, 2014).

2.5.3 Michel Waisvisz

The Hands were developed by Waisvisz at STEIM in Amsterdam in 1984. Initially they were used to control a DX7 synthesiser but soon were improved upon to become a powerful generic MIDI controller(Waisvisz, 2006). While they do pursue the intent of the composer to create a means of fluid interaction between the performer and digital instrument, the emphasis on performance seems to take a back seat in favour of the development of sound objects. In contrast to artists such as Sonami and Donnarumma, the actual spectacle of performance does not seem to be as prevalent as the musical impact of the instrument, however, one aspect of Waisvisz's work that appealed to me was the level of detail and control in his pieces.

Waisvisz inspired me to take it upon myself to explore and design methods of musical interaction that tended towards moving away from buttons and sliders as the sole means of generating music materials. His views on the performativity of composition (as discussed earlier in this chapter) also gave me the tools and language to formulate a working compositional methodology.

2.5.4 Marco Donnarumma

In contrast to the examples of the *Lady's Glove* and *The Hands*, Marco Donnarumma's implementation of his *Xth Sense System* seems to set its focus squarely on performance art. Although the audio content has obviously been considered, it seems to be secondary to the spectacle of performance. The audio content is generated through the use of sensitive microphones placed on the body of the performer. The sound of interactions between muscles, bone and sinew are picked up using this system and then the audio signal are further processed (Donnarumma, 2016b).

The pieces that Donnarumma creates using this technology cannot necessarily be evaluated on their audio content alone. In contrast to the three composers mentioned previously, Donnarumma identifies as a performance artist and as such the inclusion of the physicality and spectacle within his performances cannot be separated from the totality of his pieces. This can be observed in his stunning (and at times disturbing) performances *Corpus Nil*(Donnarumma, 2016a), *Eingeweide*(Donnarumma, 2018) and *Hypo Chrysos*(Donnarumma, 2012).

Although his works do not carry the same level of focus on sonic composition as Waisvisz, Fischman and Sonami, the metric used to judge them must be slightly altered to include the spectacle of the piece itself. Where it may fail to be compared on a purely musical level with the other pieces I have mentioned, it more than makes up for this in that the actual spectacle of his performances is meticulously designed and expertly presented, especially in his more recent works. Observing how works such as *Corpus Nil* are presented has led to me to seriously consider the role of spectacle in the performance of live electroacoustic music and as a result I have attempted to consider the overall aesthetic when presenting pieces such as *Kinesia*, *Proprioception* and *Conatus*.

3

Methodology

The previous chapter has served to situate my compositional practice within a theoretical framework and indicates the inspirations and influences which have informed my compositional practices. This chapter will investigate the implementation of these frameworks and map out the methodological approach I have developed over the course of my research.

3.1 Outline of Approach

My process of composition evolved over the course of several years until I finally arrived at a methodology that I began to consistently revisit. This methodology became fully established toward the end of my PhD research, particularly during the composition of pieces such as *Conatus*, *Djinn* and *Tenterhooks*.

This approach consists of three stages of composition, none of which are clearly delineated from the others and often spill over into one another. The first process that I undergo is to construct some means of generating a large amount of sonic material from which I can extract musical phrases that will eventually become the building blocks for my compositions. This will entail the use of either a generative process to create a large body of material or involve creating a compositional system which affords some level of corporeal interaction with sonic materials via a HCI technique of some sort. In the first instance, I would construct a framework for the generation of sound objects within the Csound environment. Software instruments would be designed to exhibit the specific characteristics (if any) that I would like to explore within the piece in question. I often employ a number of pseudo-random processes to the parameters of these software instruments to add variation and a level of uncertainty to the instruments output.

Although the output of the instrument is in some part informed through the use of pseudo-random processes, these processes are always bound by sets of clearly defined rules. For example, when generating material for *Sinmara*, I used a generative software instrument to build the graduated continuant signals heard throughout the piece (particularly at 01:30). There were a number of pseudo-random processes taking place within these instruments on a micro level, such as adding subtle variations to amplitude envelopes and slight modulations to frequency content. However, on a macro level the instrument was clearly defined and the output of the instrument could, to a certain degree, be predicted.

The use of weighted random values and other such processes does not create a totally unpredictable outcome, instead they serve to add some humanising elements to the software instrument in much the same way that randomness is used when attempting to model the output of a human performance in the design of a software instrument. It is common to add some randomness to the pitch values of the software instrument in order to correctly emulate the Gaussain distribution of output pitches produced by an instrument that does not exhibit discreet harmonic output such as a violin or the human voice (Lazzarini, Yi, et al., 2016, p. 82). The goal here is not to create an entirely unpredictable musical output but instead to facilitate the potential for subtle variation in that output.

3.1.1 Stage One

In terms of creating source material through HCI strategies, the compositional system used for *Conatus* will serve as a good descriptive example as I have used this particular compositional system to generate the source material for *Tenterhooks*. I discuss the functionality of the *Conatus* compositional system in-depth in chapter 4.10 but for now it will suffice to say that this system uses a number of time-based and granular techniques to manipulate source audio. This source audio can either be fixed media or from a live input. In the case of *Tenterhooks*, I prepared a short passage of audio that would serve as an input for the compositional system. This passage was composed of largely unprocessed (except for some high-pass filtering) audio recordings, arranged in a manner which allowed for a lot of space in between sound objects. In total this source material lasted roughly ten minutes. The entire passage was fed into the *Conatus* compositional system as a loop which allowed me to generate a large body of material through corporeal interactions facilitated by the Kinect sensor and the Glove that I had built for the performance of Conatus.[1] As this loop iterated through the sound objects present in the source material passage, I created new sound objects through the act of performance within the studio space. Altogether I created roughly forty five minutes of new material using this technique.

3.1.2 Stage Two

Whether I was using a generative process or a HCI strategy to create the initial body of source material, the next stage of my compositional practice was to choose passages or phrases which I determined to be interesting enough to build upon. I find myself constantly referring to the analogy of a sculptor, slowly hewing away at a large piece of marble until a form is defined whenever I reflect upon this stage of the compositional process. If the first part of the compositional process could be considered somewhat synthetic in terms of the thought process it occupies, this can be thought of as a far more analytic process. If possible, I will begin this process after a substantial break from the initial creation of the source material. I will then methodically listen to any musical phrases which appear, either through design or serendipity, and separate them from the larger body of material. These phrases will then be used to construct a more defined musical form. During this stage I will make some refinements to the extrapolated phrases if necessary, but these can be for the most part considered to be broad strokes, only implemented in order for reconcile opposing or clashing musical ideas into a more

[1] Full descriptions of these technologies can be found in the appendix

coherent structure.

3.1.3 Stage Three

If the second stage of this process can be considered to be the broad strokes of the entire operation, the third stage can be thought of as the fine detail phase. It is here that I will further edit material and focus on the interplay between the sound objects. It is at this point that aspects such as the careful treatment of amplitude envelopes and the vertical evolution of this piece takes place. To expand upon this a little more, the vertical evolution of the material can be thought of the interaction of multiple sound objects at a singular point in time. This leads to the creation of rich textures and can be considered analogous to the use of harmony within tonal music. The previous stage of the compositional process is far more focused on the horizontal progression of the piece, or how one sound object or phrase flows into the next.

This third stage of the process is often an exercise in distillation. To once again use the sculptor analogy, it is at this stage that any remaining marble not integral to the final form is removed. During this final stage of the process, I meditate on the cadences and breathing patterns of the piece (see chapter 2.3.1). Although I refine and mix the spectral components of the piece from the second stage of this process onward, it is really at this third and final stage that I consider the piece as a whole and apply final changes to the overall mix, altering dynamic and spectral relationships.

The three stage methodology that I have described here is in truth an oversimplified account of how I approach my compositional practice. In reality, these processes tend to bleed into one another more than is implied in the above account, and the point at which one stage ends and the other begins is at times quite nebulous. It should also be noted that although I have employed this methodological approach for many years, it was only towards the end of my PhD research that I truly became fully aware of my compositional methodology, such that I could consciously take this route when beginning to compose a new piece.

3.2 Art as Process

For much of my time as a composer of electroacoustic music I have battled with, what I initially saw as, the tension surrounding the exact definition of my role in the creative process. First and foremost I like to consider myself an artist that uses technological tools in order to achieve my creative goals. However, in order to fully harness technological tools, a modern composer must also be adept as a technician. To quote Pierre Boulez, "a virtual understanding of contemporary technology ought to perform part of the musician's invention; otherwise, scientists, technicians and musicians will rub shoulders and even help one another, but their activities will only be marginal one to the other."(1986, p. 12). This does not necessarily mean that one must be able to invent or develop new technologies in order to create their art, but one *should* have the ability to use pre-existing technologies to an advanced level.

The frustration felt when one does not have sufficient control over the materials they use is something that I have experienced before when working with software designed by another party, the functionality of which is limited to the intentions of the original designer. The potential for interaction with a tool or a system is bound by what James Gibson called "affordances" (2015). With greater knowledge of the tools and materials a practitioner works with, the greater the potential for skillful and innovative interaction with those tools and materials.

Gibson's notion of affordances was further developed by Don Norman in his book, *The Design of Everyday Things*, when he stated that, "to be effective, affordances and anti-affordances have to be discoverable-perceivable"(Norman, 2013, p. 11). This implies that in order for an affordance to be present, the practitioner must also be aware of the presence of such affordances. When working with third party software for example, a user may not be fully aware of every possible action afforded by the software due to that fact they may only interact with the software on a high-level; the software can only mostly be used as intended. However, if the practitioner has designed the software themselves, they will have worked with the software at a lower level and will likely have a more complete idea of the totality of actions available to them. If a means for interacting skillfully with a tool or a material is not perceived it cannot be exploited, therefore the more knowledge the practitioner has of a particular tool or material, the greater their ability to exploit this affordance. The potential for the skillful sonic manipulation of a sound object is greatly increased if the tools and technology that are available are familiar to the practitioner.

Initially I became frustrated with the amount of time that I would spend designing and soldering a circuit, or writing some code for digital signal processing purposes. I felt that the amount of audio material I was generating was quite minuscule when compared with the amount of time I was spending working on the technical aspects of a compositional system. I began to question whether I was indeed a composer or if I was simply a technician who developed and tested out custom built hardware and software. It was only after two years of working on this portfolio that I began to realise the work I had put in to learning a coding language or becoming familiar with a certain piece of hardware meant that I could skillfully interact with the materials in a way that granted me far greater musical and creative control.

Soon after this, it became apparent that when trying to define my role as practitioner, there was in fact, no tension between the roles of technician and artist in my work. To once more make reference to Christopher Small's work *Musicking: The Meanings of Performing and Listening*, "to music is to take part, in any capacity, in a musical performance, whether by performing, by listening, by rehearsing or practicing, by providing material for performance (what is called composing), or by dancing" (2010, p. 9).

3.2.1 Autoschediasm

The composer of electroacoustic music which attempts to push boundaries of whatever kind should be constantly using tools and techniques that are on the fringe of the established norms of the discipline. To quote the immutable wisdom of Frank Zappa, "without deviation (from the norm), progress is not possible" (Zappa and Occhiogrosso, 1999). The message here is clear; if one does not intend to be bound by pre-existing paradigms and conventions, it is necessary to have the ability to alter the functionality of pre-exiting tools to perform innovative tasks. Is that not what experimentation is all about? Progression from one existing mode of thought or paradigm to the next? I began to realise that rather than being two separate roles, the technician and the artist both fell under the umbrella of composition. This is not to say that every composer of electroacoustic music must be able to write his or her own feedback delay networks or build their own hardware controllers, but in my case these technical concerns are as much a part of my practice as arranging sound objects in a digital audio workstation (DAW).

To have this kind of relationship with ones tools and materials is not uncommon. Many painters will create their own pallets, unique to their work. In one case, the English painter Stuart Semple has created the world's "pinkest pink" which enables him and users of his paints to create images and textures that set their work apart from the works of others, at least in terms of specific range of colour(O'Connor, 2016). Many musicians will build their own instruments out of found materials, even vegetables have been used as material for creating musical instruments[2]. The idea that there is a certain power in being the complete author of ones own work is not new, as it was once said that blacksmiths had special powers and the ability to cast spells and curses due to the fact that they were the only craftsmen that could craft their own tools (Dillon and Doolan, 2017).

The ability to not only create with a set of tools but to be able to create the tools themselves grants the composer a huge amount of artistic freedom. Rather than having to use a compositional system of someone else's design (Live9 or Reaper for example), the technically competent composer can design their own compositional systems which compliment the work-flow of that individual. Not only does this approach increase the creative potential of a musical concept, it also carries with it a character that is unique to that particular composer. The more bespoke elements within a compositional system, the greater the potential for truly unique outcomes. If the software instruments, the hardware interfaces and the means of arrangement are all custom built, the chances of creating a piece of art that is *sui generis* are vastly increased.

[2]The Vegetable Orchestra are an Australian group formed in 1998 who perform music exclusively on their own custom built instruments made from various vegetables.

There is also, in my case at least, a certain egotistical boost whenever I use tools that are of my own creation. The level of control that I have over a system or a particular aspect of that system is both reassuring and empowering. If any element of the system fails during a rehearsal or performance, I like to know that I have enough understanding of how everything works that I can either resolve or circumvent the issue. This is a double edged sword however, because although it may often be more educational and cheaper to take the DIY approach, the quality of the finished product rarely matches that of one designed and built by a specialist. This consideration does not bother me too much and the roughness of some elements of my compositional systems actually tend to inform my work more than hinder it.

3.2.2 Kintsukuroi

An interesting phenomenon occurs when one designs their own interfaces, especially when the interface is somewhat experimental in its nature; there will be inevitably some bugs within the system with which the performer will have to contend with. These bugs do not necessarily result in break down of the efficiency of the compositional system[3] but can instead shape the performance in a unexpected manner. The workarounds and the allowances which need to be made for the idiosyncrasies of a novel interface, grant further agency to the tools within a compositional system, and can be viewed through the same lens as the Japanese concept of *kintsukuroi*, which means something is more beautiful for having been repaired (Kopplin et al., 2008), or to put it in the words of the late Leonard Cohen, "forget your perfect offering, there's a crack in everything, that's how the light gets in" (Cohen, 1992).

The limitations within any compositional system are integral to the artistic choices the composer makes. Not only do the constraints of a system force the artist to behave in a specific way, they also direct the future choices of the artist. I personally like the nuances and idiosyncrasies that exist within my compositional systems. They invite serendipitous events to occur and ensure that musical outcomes, as is the case when

[3]Although this can indeed be the case. Certain bugs can cause the system to cease functioning entirely, as was the case during the first performance of *Kinesia*. These issues are best described as fundamental flaws within the compositional system as opposed to bugs and must be addressed before further developments can be made.

playing with human counterparts, are uncertain and in a constant state of flux. To quote Tom O'Sullivan and Dan Igoe, "like a good magician, you should make your necessary constraints look like a perfectly natural part of your system" (2004).

Another point to be made explicit is that the more stability within a system, the more predictable the outcome. While it may be kinder on the nerves to have a system that will behave in precisely the manner it is expected to behave, it can be more exciting to be unaware of the exact output from a system and it is often the case that the audio output from a system that is operating just below the limits of stability can yield the most interesting results (Delap, 2016). This *playing with fire* ethos is echoed in the following quote by American author Kurt Vonnegut;"I want to stay as close on the edge as I can without going over. Out on the edge you see all kinds of things you can't see from the center. ... Big, undreamed-of things — the people on the edge see them first" (Vonnegut, 2006).

3.2.3 Locating the source

This realisation also had me asking myself the question, when does composition actually begin? Is it when I start to arrange sound files in the edit window of a DAW, is it when I begin to capture audio recordings, or perhaps it begins when I design a software instrument in Csound. In order to answer this question, I first had to ask myself what I understood the definition of composition to be and the answer that I was most comfortable came from a interview with Pierre Henry. In this quote he is speaking specifically about Musique Concréte but I think that the sentiment is applicable to all forms of composition; "musique concréte is about the art of decision. It's the art of choice"(Darmon, 2007).

The sentiment in the above quote, when extended to the total gamut of my artistic practice, succinctly conveys how I view each stage of my artistic process. Firstly, I may have a concept that I may wish to explore. I will then choose the manner in which I would like to realise that concept and what tools and materials I will use. The tools that I choose will then dictate the character and the manner in which the materials are generated and processed. If these tools are custom built or unique in some way, the potential for an unexpected or novel outcome is increased. The materials generated are then subjected to aural analysis and they will either be disregarded or developed depending on how they sound in the context of the other materials, or how well they lend themselves to the initial concept behind the work. Once I have chosen the sound objects that I wish to use, I will then choose the manner in which I will arrange them and knit them into each other in order to create sonically interesting textures and gestures. Using this process as a guide, it would appear that the act of composition begins when the first choice is made as to how the initial concept becomes manifest.

However, it would be equally valid to posit that since the act of composition is as much reliant on how one interacts with the tools available to them as the materials that those tools are used to create, the act of composition could possibly begin with the construction of a tool. In the case of my work, pieces such as *Conatus*, *Ouroborus* and *TenterHooks* all make use of custom built interfaces and mediation technologies. The manner in which the material is generated is a direct result of how the particular interface was designed. So perhaps in cases such as these, the act of composition does not begin with the microphone or keyboard but with the soldering iron.

It seems that the more one attempts to find the source of the river that is composition, the more elusive that source becomes. The technical processes that I undertake are as much a part of the art as the creation and arrangement of sound objects. In my opinion, the line separating the technician from the composer does not exist, and any attempt to separate process from art is futile. One simply cannot exist without the other. The process is the art.

4

Commentaries on Compositions

> *"My theory is this: I have a basic mechanical knowledge of the operation of the instrument and I've got an imagination, and when the time comes up in the song to play a solo, it's me against the laws of nature. I don't know what I'm going to play; I don't know what I'm going to do. I know roughly how long I have to do it, and it's a game where you have a piece of time and you get to decorate it."*
>
> – Frank Zappa

4.1 Sinmara

Sinmara was my first attempt at writing a piece of program music within the context of electroacoustic composition. This piece attempts to convey the fear and anxiety associated with sleep paralysis and nightmares. It was written at a time when sleep, and rest in general, were becoming less common than I would have liked. This was having a negative effect on my general state of mind. During this period, I was experiencing bouts of anxiety and night terrors; my nocturnal experiences began to feed into my daily thoughts. I felt that one way I could exorcise my head of *l'appel du vide*, would be to try and convey how I was feeling through the arrangement and composition of sound objects, in the hope that it would offer some catharsis.

The title, *Sinmara*, has its roots in the etymology of the word *nightmare*. Initially, the term nightmare referred specifically to what we now call sleep paralysis and was thought to be as a result of a demon visiting the place in which you slept. This demon, who was known by many names across the world but commonly in Europe referred to as the Night Hag, would sit on the chest of the sleeping person and ensure that they found it difficult to breathe. The term nightmare has also been used to refer the spirits and demons that were thought to possess people as they slept (Davies, 2003).

4.1.1 Process

In an effort to clear my head, I began taking long walks through the Wicklow mountains which are located close to where I live. It was January, and the mountains were covered in a blanket of snow and ice. The patches of ice which formed upon the mountain paths provided me with much of the source material for this piece.

In addition to the recordings of ice, I used several other techniques to generate source material for this piece. One of these techniques was to use a six-string electric bass guitar which was processed using a synthesis technique known as Adaptive FM (ADFM) synthesis which was developed in order to perform frequency modulation synthesis on audio signals(Lazzarini, Timoney, and Lysaght, 2007). This particular DSP technique affords a wide range of transformations when processing live audio, and the results of this can be heard in the section beginning at roughly 04:38. This particular section features the ADFM bass weaving in and out of the revisited ice-based material until about 05:27, where this conversation begins to make way for the introduction of additional sound objects.

Another source of material for this piece was a simple homemade pulse wave oscillator. This small device is basically a version of an Atari Punk Console[1] which grants the user control over the the pulse width and trigger frequency of two square wave oscillators.

[1] The APC is a simple lo-fi sound synthesiser built using the commonly found 555 timer IC. Variations on this circuit use the 556 timer chip

This can be clearly heard unprocessed at roughly six minutes into the piece.

Chance

This was the first of many compositions within this portfolio where I employed some chance operations in order to generate source material. At the time I had become very interested in the notion of indeterminacy and how this had been used by Earle Browne and John Cage (Cage, 1961) in the early 1950s. Both used randomness to inform not only the formation of structure in their compositions but also the creation of the sound objects themselves (ibid.). Keen to experiment with chance operations in my own work[2], I wrote a script in Csound which would create a constantly evolving landscape of synthesised sounds using pseudo-random processes.

```
instr RandOsc

        ia random 100, 500          ;value a
        ib random 500, 1000         ;value b
        idur random 1, 5            ;duration
        irand random 0, 24          ;random number
        iphs random 0, 1            ;phase position
        iat1 random 0.001, 1        ;attack 1
        idec1 random 0.001, 1       ;decay 1
        irel1 random 0.001, 1       ;release 1
        islev1 random 0.2, 0.9      ;sustain 1
        iat2 random 0.001, 1        ;attack 2
        idec2 random 0.001, 1       ;decay 2
        irel2 random 0.001, 1       ;release 2
        islev2 random 0.2, 0.9      ;sustain 2
        icps random 60, 600         ;cycles per second
        iamp1 random 0.1, 0.5       ;amplitude 1
        iamp2 random 0.1, 0.5       ;amplitude 2
        iamp3 random 0.1, 0.5       ;amplitude 3
        iamp4 random 0.1, 0.5       ;amplitude 4

...
```

Code Excerpt 4.1: Variables used for generating psuedo-random material

[2]Although as I have discussed in chapter 2, I was not entirely subscribing to the same model as Cage et al.

This script consisted of four instruments. The first was an additive synthesis design that mixed four sine waves of varying pitch and amplitude envelopes. The second of these instruments was a complex FM synth which again made use of pseudo-random processes to control the amplitude envelope as well as the carrier and modulator frequencies. The remaining two instruments were trigger instruments, designed to trigger each of the aforementioned synthesisers at indeterminate intervals for as long as the script was running. The material generated by this script provided me with a scaffolding on which I could begin to build up various other sound objects.

```
        ;Spectral Envelopes
        kfreq1 linseg ia, idur, ib
        kfreq2 linseg ia, idur*0.5, ib, idur*0.5, ia
        kfreq3 linseg ia, idur*0.25, ib, idur*0.25, ia,\\
        idur*0.25, ib, idur*0.25, ia
        iband random 10, 200

        ;Amplitude Envelopes
        kenvL madsr  iat1, idec1, islev1, irel1
        kenvR madsr  iat2, idec2, islev2, irel2

        ;Virtual ICHING
        if (irand <= 10) then
                kfreq = kfreq1
        elseif (irand <= 20) && (irand >= 11) then
                kfreq = kfreq2
        else
                kfreq = kfreq3
        endif

        ;Glissandi
        islide random 0.3, 4
        iline random 1, 100
        icps2 = icps*islide

        kcpsline line icps, idur, 200
        icpsline = i(kcpsline)

        if (iline >=40) && (iline <= 50) then
            icps = icpsline
        endif

        a1 oscils iamp1, icps*(2^(irand/12)), iphs
        a2 oscils iamp2, (icps*2)*(2^(irand/12)), iphs
        a3 oscils iamp3, (icps*0.5)*(2^(irand/12)), iphs
        a4 oscils iamp4, (icps*4)*(2^(irand/12)), iphs
        a5 oscils 1, icps*0.25, iphs

...
```

Code Excerpt 4.2: Implementation of pseudo-random processes

4.1.2 Structure

In terms of its arrangement, the goal of this piece was to emulate a situation in which the listener is perpetually being jolted awake, unaware of their new surroundings and uneasy with the change in environment. As soon as the listener grows accustomed to the new reverie they are once again awakened with a start and plunged into a new strange landscape. The act of being jolted awake is articulated using a sound object with an attack-decay trajectory (Smalley, 1997). Figure 4.1 shows a spectrogram plot of once such gesture which occurs at approximately two minutes and twenty-three seconds into the piece.

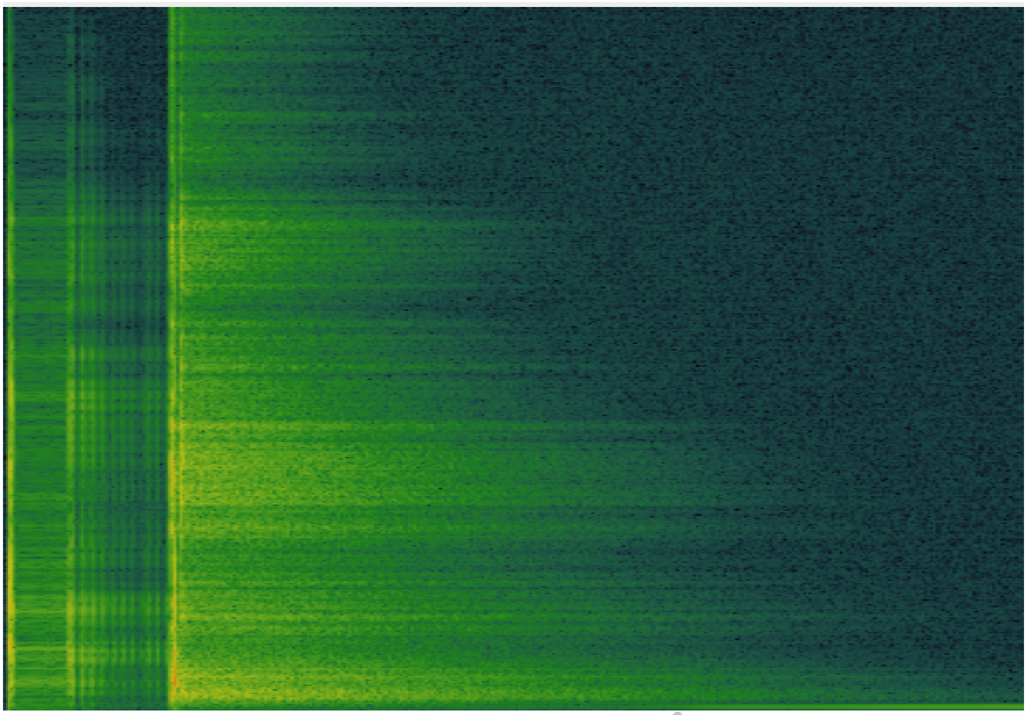

Figure 4.1: A gesture with an attack-decay trajectory preceded by several short percussive gestures

4.1.3 Conclusion

Some compositions seem to come from somewhere within the ether, appearing to almost write themselves. In such cases the composer seems to be no more than a conduit for some divine artistic message. This work was certainly not one of these pieces. *Sinmara*, rather than flowing from my mind into a passage of musical phrases, crawled slowly and begrudgingly out of me, eventually becoming manifest in the collection of sounds presented in this work. This work is the only example within this portfolio that relied on no augmented gestural input during its creation. Although the compositional process was difficult, it did indeed serve as a means of catharsis. I put so much of myself into this piece that when I revisited it for the purpose of analysis for this portfolio, I found myself once again occupying the head space that defined this period of my research. Needless to say, that was not a particularly pleasant experience; however, upon reflection, I was surprised at the amount of visceral emotion that could be contained within a such a piece.

4.2 Kinesia: The Third Law

Kinesia:The Third Law is the first composition in this portfolio where I make use of the Kinect as a means for enabling physical gesture and total corporeal movement as a compositional tool. In early 2015 began experimenting with gesture based control over musical parameters using the LEAP motion sensor.[3] While this did afford greater expressive potential in the performance of my music than a mouse and a QWERTY keyboard (Emmerson, 2016), I found that the limitations of this sensor were at times too great to achieve the level of control that I desired. I was becoming increasingly interested in using the whole body as a musical interface and had invested quite some time in exploring the possibility of wearable devices to achieve this. I was also investigating the potential for a collaboration with someone that was trained in the art of movement. I felt that the potential for skillful interaction with a corporeally informed interface would be greater if it was used by a dancer. I made the acquaintance of Conor Donelan, a dancer who was eager to experiment with a corporeally informed musical interface and we began experimenting with methods for translating motion into music.

4.2.1 Inspiration

After several meetings with Conor, and following a number of experiments involving attaching microcontrollers to the body of a performer, I decided that using wearable devices to collect and transmit data related to movement was not a viable solution. This was mainly due to the encumberment caused by the multitude of physical connections required to connect the necessary sensors to a microcontroller such as an Arduino. Instead I decided to use the Microsoft Kinect for the task of tracking Conor's movements. I had seen the Kinect being used in conjunction with Processing, a graphical programming environment, to allow users to virtually paint on a screen using physical gestures.

After some further investigations into accessing the raw Kinect sensor data via Processing, I came across a library from the OpenNI project (see Appendix A.2.2) which granted me access to the skeletal data collected by the Kinect. This library grants the

[3]These early experiments into gestural control are discussed in detail in Appendix B and Appendix C in relation to the installation *EAREYEMOUTH* and the compositional system *Ouroborus*

user access to the Cartesian coordinates of fifteen points of the tracked body (see Figure A.1). I then normalised the coordinates of each of these fifteen points and sent the normalised data to Csound using the OSC protocol. In addition to sending the skeletal data to Csound, I collected and sent data pertaining to the speed of movement of both arms, the distance between the hands of the users and a variable to check whether or not the user is being tracked by the software. These variables would serve as mechanisms for manipulating musical parameters and triggering sonic events within Csound.

The inspiration for some of the instruments and the overall narrative of the piece came from conversations that took place between Conor and myself regarding aspects of eastern philosophy and mythology. I began researching the use of sound in Buddhist and Taoist religious ceremonies and decided to make use of Tibetan singing bowls as a musical device. American composer John Bower did some work in documenting the modal frequencies of Tibetan singing bowls of varying sizes and these frequencies were contributed to the Csound manual by fellow composer Scott Linderoth (Linderoth, n.d.). I used these modal frequencies as the basis for the design of several instruments within this piece.

4.2.2 Plenary Instruments

Singing Bowl

The first instrument I designed for this piece is inspired by the sound achieved while employing a singing bowl technique known as rimming. With this technique a graduated continuant of sound is produced by rubbing a specially designed mallet in a circular motion along the rim of a Tibetan singing bowl. This action causes the body of the bowl to resonate and produce a harmonically rich tone in much the same manner as one might use a finger to cause a wine glass to resonate.

The modal frequency ratios provided me with the virtual resonant body of the bowl but, just as in the real world, an excitation source was required in order to produce sound. In this particular instrument, the excitation source I used was white noise, which emulated the constant circular rubbing motion used to cause the bowl to resonate. The speed at which this circular motion takes place is determined by the position of the dancer's head on the y-axis. There is the choice of using three different bowls when this instrument is triggered, this choice is made using a pseudo-random process in Csound which functions similarly to a coin toss each time the instrument is triggered. The result of this coin toss determines whether the bowl used is a 180mm bowl, a 152mm bowl or a 140mm bowl. The reason I included this element of pseudo-randomness into this instrument was to add certain nuances to the sound and to allow the software itself to hold a two-way conversation with the dancer throughout the piece.

The resulting sound is then filtered using a Butterworth band-pass filter, the bandwidth of which is determined by the distance between the dancer's left and right hand and the center frequency of the filter is influenced by the position of the dancer's left hand along the y-axis. This creates a connection between physical and musical gesture in the sense that the sound of the instrument will appear to open up (due to the increased spectral content) as the dancer's hands move further apart. The position of the dancer's body across the x-axis determines the fundamental frequency of the bowl and the position along the z-axis determines the amount of resonance present in the band-pass filter.

A pseudo-random process, similar to the process that determines which bowl is to be played at a given time, is used to position the audio output in the stereo field. The instrument itself is triggered when both hands of the dancer are in close proximity to the torso. This particular gesture is also inspired from research into Eastern philosophy, specifically the Samadhi or Dhyana mundra posture. This allows the sound to begin quite subtly as the hands are required to be close to one another in order to trigger the instrument, which in turn means the bandwidth of the filter is quite narrow. This allows harmonic content of the instrument to slowly grow from the moment of the instruments instigation. The first instance of this instrument can be heard at 01:06.[4]

[4] These times refer to the audio version of *Kinesia* featured in the main portfolio of compositions which can be found in the flash drive that accompanies this text

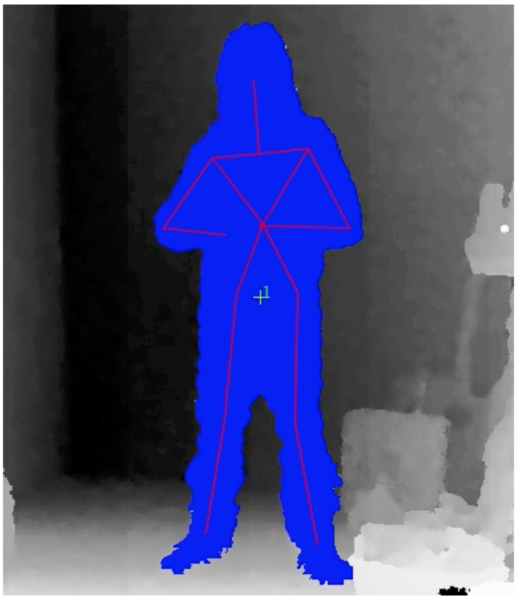

Figure 4.2: Posture required to trigger Singing Bowl Instrument[5]

Transient Bowl

The second instrument uses the same method of physically modeling Tibetan singing bowls as the first. It also makes use of the same pseudo-random process to select between different virtual bowls. The main difference between this and the previous instrument is that it employs a different excitation technique. Instead of using white noise to emulate the sound of resonating bowls, this instrument uses short impulses to create transient bursts of sound that have a sonic character similar to that of dropping a marble on a metallic surface. The rate at which these impulses occur is determined by the position of the left hand along the y-axis. The higher the hand is raised, the greater the interval between the impulses. This allows the dancer to mimic the natural physical behaviour of a bouncing ball; as the distance from the two surfaces between which the ball is bouncing decreases, the amount of times per second the bouncing object makes contact with either of the surfaces increases. This is another example of an attempt at trying to make an intuitive connection between physical and musical gestures.

The right hand also affects the instrument as it moves up and down along the y-axis. When it is at its zenith, the resulting sound is allowed more time resonate due to

[5] All postural images are stills from skeletal tracking windows provided by the Kinect depth tracking OpenNi Processing sketch

higher Q factor for each modal frequency. As the hand is lowered, the Q is decreased and the resulting sound becomes a lot less resonant, displaying a deadened or muted sonic characteristic. The left hand's x-axis has a subtle influence on the fundamental frequency of the modeled bowl. Pseudo-random processes also have an influence on the fundamental frequency of the bowl and the position of the sound in the stereo field.

This instrument is triggered when both the dancer's hands are above a certain threshold on the y-axis, which usually results in the dancer having both arms raised above his head. This posture follows nicely from the dancer having to spread his arms open to fully activate the first instrument creating a pleasing fluid physical motion between activation of instruments. An example of this instrument can be heard at 00:17.

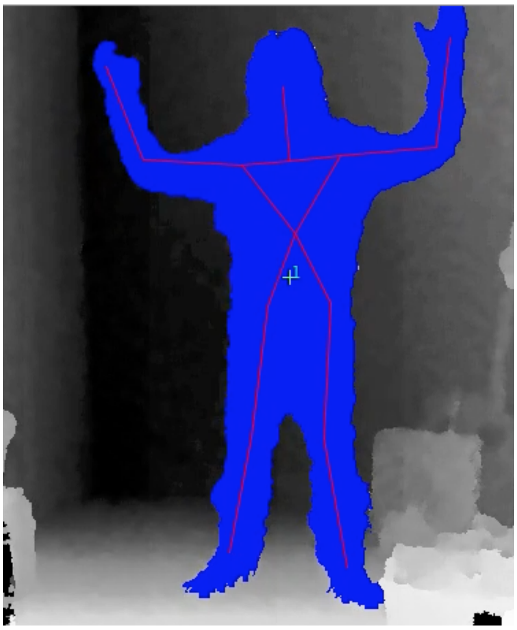

Figure 4.3: Posture required to trigger Transient Bowl Instrument

Pad

Instrument number three, or as I refer to it, the pad instrument, originally resembled a soft sweeping pad sound. I eventually carved it into a more sinister, colder sounding manifestation in order to differentiate it from another instrument I had used in a different project. It is essentially an additive synthesiser consisting of seven oscillators. The waveform of each of the oscillators is either a pulse wave or a sawtooth wave in various manifestations. These oscillators are mixed and sent through five different Butterworth filters arranged in series. The cutoff frequencies of these filters are determined by position of the dancer's feet along the y-axis. The resulting signal is passed through an amplitude envelope and filtered once more by a bandpass filter.

In this instrument the bandwidth of the bandpass filter is dictated by the distance between the dancer's hands and the center frequency of the filter is determined by the dancer's overall position on the z-axis. This allows for both a sonic and physical connection between this instrument and the first instrument if they are played in unison. The resulting musical output is similar to two instruments playing a harmony with one another with each instrument following the others dynamics very closely.

The triggering of this instrument occurs when the dancer's head is below a certain threshold but the hands are above a certain threshold. The resulting form resembles that of a bird preparing its wings for the downward flapping motion during flight.

An example of the pad instrument can be heard subtlety entering at 05:34. This instance had in fact been triggered earlier alongside the gong instrument but due to the dancer's posture it was filtered such that it remained unheard until roughly twenty seconds after it had been triggered.

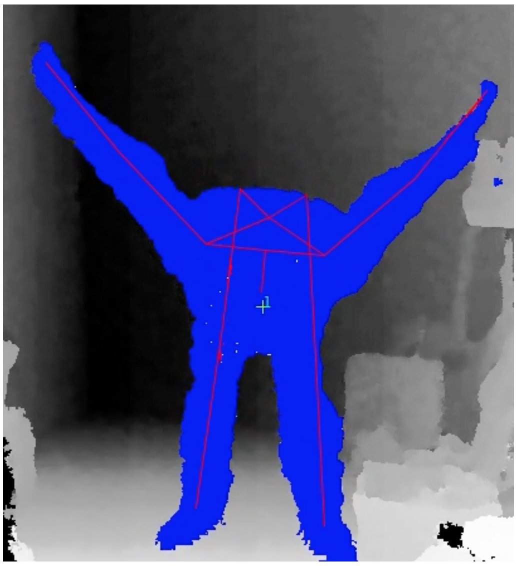

Figure 4.4: Posture required to trigger the Pad Instrument

4.2.3 Apportioned Instruments

The previous three instruments have parameters that can be adapted and changed by the dancer after they have been triggered. Three of these instruments can only be triggered by the dancer and no subsequent movements of the dancer are used inform their spectral evolution. Another forth instrument behaves conversely to the other three instruments as I will discuss later in this section.

Pseudo-random processes in Csound, all of which are restricted by a set of predetermined boundaries, dictate much of the performance information concerning these instruments. While there is still a small amount of parametric information received from the dancer's movements, the dancer's influence on the real-time properties of these instruments is drastically reduced in comparison to the previous instruments.

Bell

The first of these particular instruments is designed to resemble a large bell or a gong. The triggering of this instrument is reliant on the speed of movement of both of the dancer's arms. The speed of movement is calculated in Processing and that information is then sent to Csound via OSC. Once the speed of motion crosses a certain threshold the gong is then triggered. The speed across both the x-axis and the y-axis is measured in order to not restrict the triggering of this instrument to one particular vector. The initial position of the dancer's right hand along the y-axis has a small modulating influence on the frequencies of the each of the nine oscillators used to synthesize this bell-like sound.

This instrument can be heard throughout the piece, a particularly satisfying instance with a nice amount of low frequency content can be heard at 00:38.

Pūrerehua

The next instrument of this category originally resembled the sound of an object such as a Pūrerehua, also known as a bullroarer, a Maori instrument which creates sound as it is whirled in cyclical motion through the air. The final manifestation of this sound was softened using reverb which makes the character of the instrument a bit more like a swoosh than the aggressive sound of the Pūrerehua. The basic synthesis of the instrument is similar to that of the singing bowls, but the manner in which it is triggered is much akin to the gong instrument.

The trigger for this instrument is the speed of motion of the dancer's right hand. This instrument serves to provide an obvious musical analogy for the physical movement used to instigate it. The resulting sonic gesture is described by Adrian and David Moore as a swipe or a swish (A. Moore and D. Moore, 2011). As I have already mentioned, several pseudo-random processes dictate many of the parameters of this instrument but certain input from the dancer will still influence its behavior. For example, the position of the dancer's body along the z-axis will change the cutoff frequency of a low-pass filter and the position of the sound in the stereo field is influenced by the position of the dancer's left hand along the x-axis.

One of the most consistently triggered instruments in the piece, the Pūrerehua is used to begin the performance.

Gen

This instrument is triggered upon the first occurrence of the Gong instrument. Unlike the other three apportioned instruments, the performer can influence the behaviour of this instrument once it has been triggered. The Gen instrument uses the same physical modelling mechanisms used for the Transient Bowl and Singing bowl instruments but the manner in which these models are articulated is determined through a mix of pseudo-random processes and data pertaining to the location of the performer in relation to the Kinect sensor. As long as the Gen instrument is engaged it will continuously trigger instances of itself

This instrument is disengaged upon a subsequent triggering of the Gong instrument but it does not cease to trigger instances of itself immediately, instead its presence tapers away over a short amount of time.

Gong

The final instrument that I created for this composition plays a very important role in the arrangement of the piece and also acts as a structural flag to ensure that not only is the dancer aware of how far the piece has advanced but he or she can encourage its continued evolution. Spectrally, this instrument is quite similar to the singing bowls in terms of frequency ratios, but the amount of harmonics present here is much greater, therefore creating a far richer sound. The amplitude envelope is also drastically different to that of the bowls, which in-turn means that the entrance of this instrument (which I refer to as the gong) into the sonic environment is quite sudden and dramatic. The performer has even less control of the spectral morphology of this instrument than any of the other instruments. None of its parameters are randomized, meaning that the sound of this particular instrument is predetermined and not susceptible to any environmental or performative influences. It is in fact, the only sonic constant throughout the entire piece.

The gong is triggered when the dancer's head is below a specified threshold. This threshold is so low that the dancer has to be almost completely lying down in order to trigger this instrument. The main function of this instrument is to mark the transition between the four sections of this composition. I included a counter within Csound which monitors how many times the gong has been triggered. Different sections of the piece are activated according to how many times the gong has been triggered. There are certain restrictions imposed on this instrument, such as the fact that each section of the piece has a minimal amount of time that it has to be active for in order for this instrument to be triggered. For example, the first section of the piece has to run for a minimum of 90 seconds before the transition into the next section occurs.

The first instance of the gong can be heard at 02:47. The marks the transition from section 1 to section 2. The second instance can be heard at 05:05, this gong marks the beginning of section 3. The final gong occurs at 06:25 and heralds the end of the piece.

I will now speak briefly about the characteristics of each of the sections within this particular piece.

Instrument	Trigger	Modulation Effect	Modulator
Singing Bowl	(R.Hand.x-Torso.x)&(L.Hand.x-Torso.x)<40	Q,Band-Pass,Initial Frequency	Torso.z, Hand Distance, Torso.x
Transient Bowl	(R.Hand.y)&(L.Hand.y)<40	Interval,Q, Frequency	R.Hand.y, L.Hand.y, R.Hand.x
Pad	(R.Foot.y>0.5)or(L.Foot.y>0.5)	Hi-Pass, Low-Pass, Band-Pass	R.Foot.y, L.Foot.y, Hand Distance, Head.y
Bell	Hand.Speed>20	Frequency	Hand.x, Random
Pūrerehua	(R.Hand.x)&(L.Hand.x)<40	Frequency, Cut-off, Panning	R.Hand.y, Torso.z, R.Hand.x
Gong	(Head.y<0.3)&(time>90)	None	None

Table 4.1: Functionality of plenary instruments used in *Kinesia*

4.2.4 Structure

The overarching shape of this piece was conceived when reflecting on Joseph Campbell's concept of the *monomyth*(2008). Without drawing too directly from the precise structure of the archetypal themes of the *heroes journey*(ibid.), I simply wanted the piece to be informed by the general shape of its form; moving from exposition and refinement, into chaos and surrender of control, and back into an evolved manifestation of the initial ontological state. This concept is reflected in the levels of control and potential for chaos afforded by the compositional system at any given moment of the piece. In short this composition follows a ternary form. The following time stamps presented within this commentary relate to the performance of *Kinesia* provided in the main portfolio which accompanies this thesis

Section 1

This movement functions as an exposition for the sonic material within the piece. There are several restrictions put in place to allow the exposition to occur gradually and for each of the instruments to be demonstrated as an autonomous entity. The restrictions that are utilized here include a minimum amount of time to be passed before the gong instrument is triggered (meaning that the first section can not be finished too early) and the singing bowl cannot be triggered before a certain amount of time has passed.

The reason for the latter restriction is that the singing bowl tends to have quite a rich spectrum when the bandwidth of the filter is increased therefore it takes up a lot of spectral space. The limit imposed on its behavior ensures that some of the more subtle instruments such as the Pūrerehua, have their time in the limelight so to speak. After these minimum time conditions have been satisfied the performer then has complete control over what instruments are triggered and how they will subsequently behave. Conceptually this section represents the illusion of control that one may experience in any given real life scenario.

Section 2

Once the performer has sounded the first gong at 02:23 the second section commences. In this section, the performer no longer has control over the triggering of sonic events, meaning that any of the instruments, with the exception of the gong, can occur at any given time. What the performer does have control over is how these events behave once they have been triggered. The construction and design of the instruments in this section is essentially the same as the previous section bar one or two subtle alterations, but the manner in which they are manipulated is slightly different. For instance, the absolute position of the dancer's body on the stage has more of an effect on the sonic events than in section 1, and due to the subtle differences in their characteristics, some of the sounds generated are unique to this section. There is more focus on indeterminate processes in this section and as a result, the compositional system itself is the driving force behind of this movement. This section represents the chaos that can occur when factors beyond our control influence a given situation.

Section 3

This section occurs at 04:41 and sees full musical control returned to the performer and reconciliation between order and chaos. The ghost of chaos still haunts this section but is overpowered by the influence of the performer. Musically, this section is a recapitulation of the first movement except none of the restrictions of the first movement are in place. As time passes, the chaos subsides completely until the final gong is sounded. The piece then breathes its last as the sonic objects gently fade into the distance.

Score

The graphic score shown in Figure 4.5 is presented in three stave-like sections. The top stave corresponds with Section 1 of the piece. Each of instruments with the exception of the Gong is represented on the left side of the stave. The probability of a particular instrument occurring at a particular point in time throughout this section is indicated by the thickness of the horizontal lines. For instance, the piece always begins with the Pūrerehua, therefore the corresponding line is very thick at the beginning of the piece due to the fact that it is highly probable that the Pūrerehua will be featured.

There are three time markers for sixty, ninety and two hundred and fifty seconds into piece. Each of these relates to a constraint on whether or not an instrument can be activated. The first time marker at sixty seconds relates to the fact the Singing Bowl cannot be triggered until at least this amount of time has passed in the performance. The other two are used to represent the time necessary to have passed before the first and last Gong can be triggered respectively. The Gong instrument itself is represented by the character delta (Δ) and can be seen above the stave at the end of Section 1, towards the end of Section 2 and midway through Section 3.

The double barlines indicate transitional points within the piece. These can indicate a change in the availability of certain instruments, as can be observed at the sixty second mark with Singing Bowl, or a complete change in agency within the compositional system as shown at subsequent occurrences of the Gong instrument.

Section 2 moves away the singular representation of instruments due to the fact that the performer's control over that part of the compositional system is relinquished at that point in the piece. Instead, what is shown are rough guidelines as to where the performer should be positioned on the stage which will in-turn affect the behaviour of the Gen instrument. The letter L and R are used to indicate stage left and stage right and the upstage and downstage indications are shown on the left of the stave. T white circles suggest that the performer should linger at given positions for extended periods of time. After the next set of double barlines the position of the performer onstage is left to their own discretion.

Section 3 sees the ability to trigger instruments regained by the performer and this reflected in the return of the use the horizontal bars. After the final transition indicated by the double barlines, the performer is directed to slowly reduce the occurrence of additional instruments by reducing their overall movement as alluded to in the narrow horizontal bars at the end of the score.

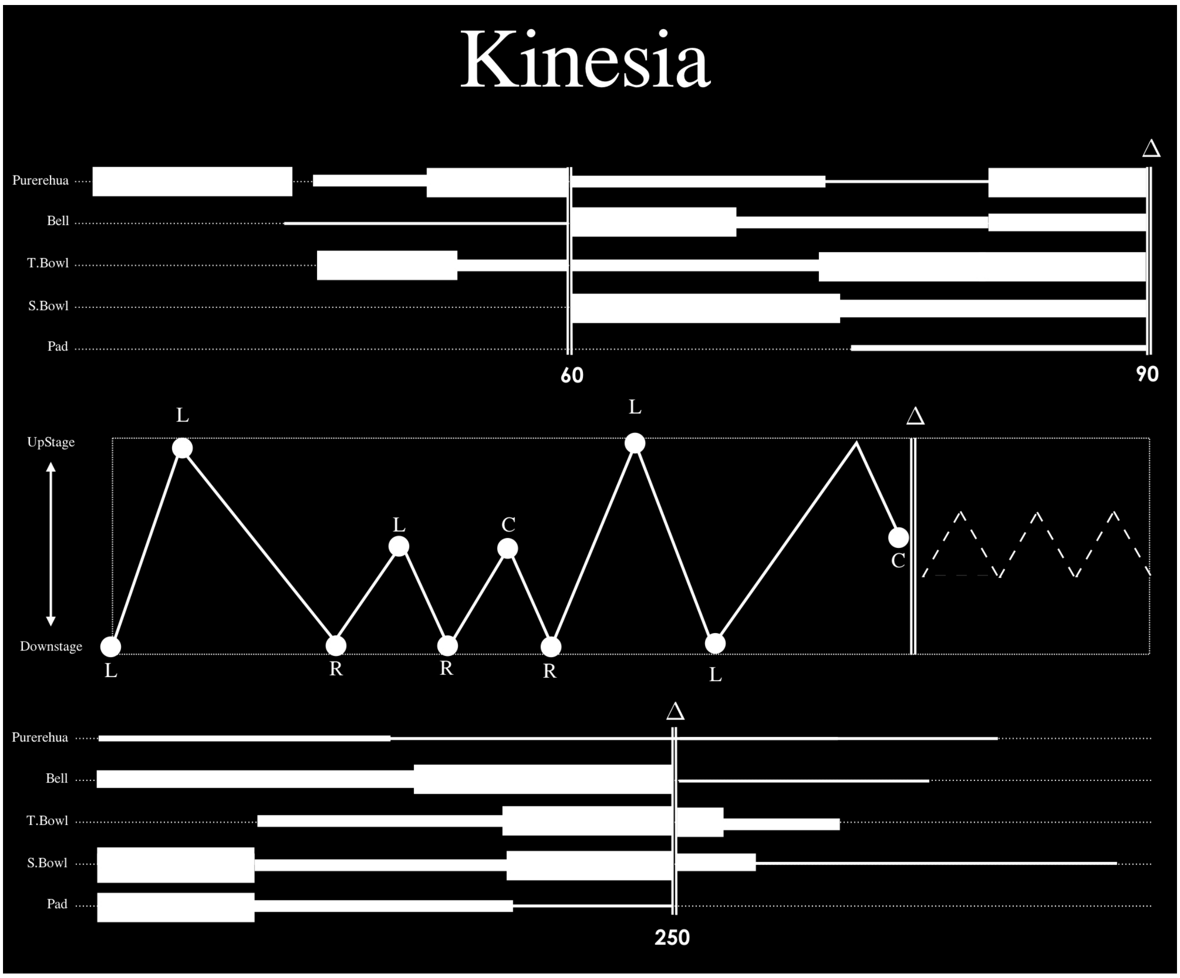

Figure 4.5: Graphic score for *Kinesia*

4.2.5 Conclusion

This piece was effective in the sense that it overcame some of the performative issues I had with the LEAP motion sensor. It enabled the performer to fully embody the piece and as a result of this it circumvented the issue I had with the lack of performative gestures afforded to the practitioner of live electronic music. However, there seemed to be several problems with the particular system that I used in this composition.

Firstly some of the triggering was at times unreliable. Curiously, the system seemed to behave unreliably if it was being used by a performer other than myself, perhaps due to differences in relative distances between points of the body but this is only a hypothesis, as I was unable to faithfully recreate the errors myself.

When rehearsing, I would demonstrate a posture and subsequent musical gesture to the dancer but when he attempted to replicate the same gesture, either the audio output was not exactly comparable or the sound object would not be triggered whatsoever. This inconsistency made designing the compositional system a regular source of frustration for both myself and the dancer due to the fact that I could not faithfully replicate any systematic faults without the dancer present. This meant that much of our rehearsal time was spent trying to reproduce errors rather than address the artistic concerns of the piece.

Secondly, one major failing of the system was how susceptible it was to variations in environmental lighting. This problem was largely out of my control due to the fact that the issue lay within the hardware of the Kinect and the lighting of the performance space. On occasion, I could influence the lighting of the performance space to a certain degree but mostly this factor was out of my control. The only way in which I could see myself preventing this was to devise a more robust triggering system that incorporated a number of fail safes to prevent the triggering of false positives. However, this did not fully address the issue caused by certain lighting restrictions, meaning that sometimes a physical gesture would not trigger a musical gesture at all but it certainly was a step in the right direction to refine the overall reaction quality of the system.

In one particular instance, I was put in a situation where I had to perform the piece myself. I was initially reluctant to do so but the circumstances demanded it, and in hindsight it was an incredibly rewarding experience. Not only did I get opportunity to perform my own work but I also found that I was less nervous performing than I would have been when watching someone else use my compositional system. This, initially stressful event, went on to inform the direction I took with my subsequent corporeally informed work.

4.3 Cascando

Cascando was written as an accompaniment to a live performance of the Samuel Beckett poem of the same name by Dublin based theatre company, *Mouth On Fire Productions*. This work is based on the 1936-37 poem *Cascando*, not to be confused with the 1961 Beckett radio play of the same name (Beckett, Lawlor, and Pilling, 2014).

The piece was constructed around a spoken word performance of the Beckett text from the company's artistic director Cathal Quinn. The final production of the piece was to be set to an improvised dance and I was asked to provide a sonic backdrop, based on the meter and inflection of the spoken word recording, to which the dancer could move.

4.3.1 Process

I was given very little direction in relation to the exact sounds to be used or the general aesthetic that the director had in mind for this piece, which was simultaneously liberating and stifling creatively speaking. A creative paradox is commonly experienced by practitioners of any artistic expression when given little or no limitations on possible materials, this paradox is directly related to the phenomenon know as the paradox of choice, which psychologist Barry Schwartz investigates in his book of the same name (Schwartz, 2004). In an effort to counteract this paradox, which refers to the fact that the amount possible choices are often inversely proportional to overall satisfaction and sense of freedom when making decisions, I decided to limit my source materials to cymbals alone. This was not for any conceptual reason, it was simply because I personally enjoy the presence of harmonic and inharmonic qualities that reside within the spectral makeup of a resonating cymbal. The layers of tonal colour, especially when bowing cymbals, seem to become more complex the more they are examined and this has led me use cymbals time and time again for source material in my work. Complex, inharmonic sounds in general tend to attract my attention due to their pareidolic qualities.

4.3.2 Mimesis

Linguistic

The process of selecting and composing material for this piece was mainly informed by the mimetic qualities of the sonic material in terms of its relationship with the linguistic properties of the spoken word material as well as the sonic properties (Emmerson, 1986). For example, the onomatopoeic qualities provided by the line “thud of the old plunger”, provides the inspiration for the choice of sonic material which immediately precedes it. Another example of this linguistic mimesis can be heard accompanying the lines “the churn of stale words in the heart again”. In this case the repetitive amplitude envelope which periodically rises and falls seeks to emulate the sickening, cyclical motion implied in these lines.

Phonetic

To provide an example of mimesis in terms of the phonetic properties of the spoken word material, there is a overt example which follows the phrase “pestling the unalterable whey of words”. A small splash cymbal was bowed in such a way as to serve as an antiphonal comment to the word “whey”.

Physical

I realised that there was potential for the material to be composed in such a way that the musical gestures could be used to inform the subsequent physical gestures of the dancer in what Rajmil Fischman refers to as an “intuitive metaphor” (2014). Denis Smalley discusses this action/consequence relationship in his paper *Spectromorphology: explaining sound-shapes*, in which he defines the term *source-bonding* as “the *natural* tendency to relate sounds to supposed sources and causes, and to relate sounds to each other because they appear to have shared or associated origins[emphasis in original]”(1997, p. 110). An effort to facilitate intuitive metaphors and source bonding with the movements of the dancer became a consideration when composing the sonic material.

Although the choreography was not completed when I began composing this piece, I did get a chance to view the dancer performing an improvised dance to the spoken word recording and this allowed me to get an idea of what tendencies the dancer had in terms of gestural vocabulary. It was these early observations that informed the spectro-morphological evolution of many of the sound objects in this piece.

The dancer was given quite a bit of direction when improvising the piece and this indirectly gave me the opportunity to get a feel for the correct flow and meter of the sound objects. For example, many of the dancer's movements were smooth and quite measured; to complement this I elected to use sound objects with graduated-continuant spectro-morphologies as opposed to sound objects with attack-decay morphologies.

It is in the context of gestural mimesis that *Cascando* is corporeally informed. The gestural vocabulary of the dancer, in conjunction with the delivery and meaning behind the words of the poem, were used to guide the movement and shape of each musical phrase.

4.3.3 Conclusion

The entire piece was an interesting, non-linear way of composing, in the sense that I was initially taking cues from the phonetic and linguistic features within the spoken word recording and to a certain extent I was also inspired by the brief moments of dance rehearsals that I witnessed. The musical composition then went on to influence the manner in which the dance was performed in its final iteration. Being aware that I was providing musical gestures which would be interpreted by a dancer had the effect of making me compose in a more horizontal manner than usual. I would generally tend toward layering many sound objects into dense vertical structures, but the fact that these musical gestures were to be manifest as physical gestures made me consider the sound objects in a more sequential manner. This perhaps stemmed from the sequential manner that one may think in when choreographing a series of corporeal movements.

As an exclusively audio work in its own right I do not feel that this piece is particular strong, however, it was written to disseminated along with the spectacle of a choreographed danced. In this regard the piece is satisfactory. Upon completing this piece and contemplating its validity as a musical work, I decided to attempt to compose live performance pieces that could stand under scrutiny of purely aural assessment when taken out of their performative context.

This piece was written for Mouth on Fire productions as part of their show *Time Passes*. It premiered at the *Theatre X International Performing Arts Festival* in Tokyo, June 2016.

4.4 Galilean Moons:Callisto

Galilean Moons: Callisto was written as part of a collaborative project directed by Dr Michele Biasutti and Dr Eleonora Concina of the University of Padua in Italy. This process began after I answered a call from the university looking for composers to collaborate on a fixed media electroacoustic composition.

After I was selected to take part in this project I was set up with an account for the University's Moodle page. Here, the details of the project were fleshed out for the participants. The project team consisted of fifteen composers which were then placed into five teams with three composers in each team. The teams all had their own message forum, live chat facility and project diary where they could collectively note the progress of their piece. We were then encouraged by the project facilitators to converse and exchange ideas on the forum provided which in hindsight proved to be an excellent resource for collecting and reassessing compositional ideas.

4.4.1 Process

My team consisted of myself, a London based composer called Daniel James Ross and Kyle Stewart, a composer based in Glasgow. All three of us expressed an interest in how we might interact with the sound material that would eventually be used within our composition.

Daniel suggested that we approach this piece like what he termed "an exquisite corpse drawing". An exquisite corpse is a drawing which consists of multiple sections of the body of a character being drawn by several different people. As the paper is passed around to the different participants it is folded in such a way so that none of the participants know what the other has drawn until the picture is completed, at which point the collaborative drawing of the *exquisite corpse* is revealed. To elaborate on how we translated this concept into an electroacoustic composition, the idea was that each of us would at first compose our own short section, based around a unifying theme of some sort, then we would begin to reassess the work as a whole whilst making reference to the compositional choices we each made within the piece thus far.

The theme that we chose to direct our compositional choices was Galileo Galilei, who served as a professor of mathematics at the University of Padua. During this time Galileo published the work *Le Meccaniche*, which introduced the law of falling bodies. This work contradicted Aristotle's earlier theories on the nature of falling objects and paved the way for Newtonian physics. With this in mind, I generated much of source material by collecting objects of various shapes and sizes and dropping them onto surfaces made from materials such as fabric, wood and metal. The initial sound created from the falling objects, as well as the resonance created from the interaction and any subsequent bouncing or ricochet-like sounds, were what provided me with the source material for this piece. I further processed this source material within Csound, controlling processing parameters through the use of the LEAP motion sensor.

Kyle took a similar approach to creating his source material and used sounds created from spinning spherical objects. Both Kyle and I stated in the discussion forum that although we were basing our initial compositional choices on a mimetic discourse (Emmerson, 1986), our final judge would be our ears. Once the material was gathered from the heavily conceptualised Galilean ideal, aural discourse would become the main source of compositional motivation. Daniel on the other hand, who stated that his main compositional interest at the time was in the area of algorithmic composition, took a slightly more deterministic approach to the compilation of his source material. He analysed the orbits of four of Jupiter's moons, all of which had been first observed by Galileo, and used the ratios of these orbits to inform the length of audio loops used within the section of the piece. These moons were Io, Europa, Ganymede and Callisto and the ratios of their orbits (approximated by Daniel) were 9.4: 4: 2: 1 respectively.

4.4.2 Arrangement

The actual assemblage of the piece took place within an online DAW called Ohm Studio. Ohm Studio is a cloud based DAW which allows multiple users to collaborate on compositions and mixes in real-time with all changes being saved to the Ohm Studio cloud server. As far as the DAWs functionality was concerned, there were a few issues with access and several occasions where workflow was quite slow due to connectivity issues, but because of the fact that all three of us were only using Ohm as a canvas on which we could arrange our material, as opposed to actually processing material, the software served its purpose well.

I was the first to complete my section of our *exquisite corpse* and I arranged and uploaded my material for my teammates to hear and provide feedback. They were both very positive about the material I had provided and they both expressed a desire to base their work upon some of the gestures that I had created within that material. In fact, Daniel ended up using the sound objects I created almost exclusively for the source material of his audio loops.

Once all three sections were complete and uploaded to the Ohm Studio project, we each discussed what changes might need to be made to the overall structure of the piece in order to make it coherent. Due to the fact that we had worked closely with each other over the months in which the piece was composed, coupled with the fact that Daniel had used my sound objects as a recapitulation of sorts, the piece was surprisingly coherent as it was and we all agreed only minor adjustments were necessary.

4.4.3 Conclusion

Upon further reflection, while this collaboration was a breath of fresh air after composing alone for so long, I couldn't help but feel that at times some people were more committed to contributing to the project than others. I think that this is apparent in the end result and while at the time I was happy with the coherence of the piece, listening to it one year later has presented me with some doubts. This piece as a whole may not be a structurally unified composition partly due to social pressures, politeness and ego, but it certainly highlighted the power of collaboration once again. The composer's life can be a lonely one, especially when working with fixed media, and any opportunity to share ideas with others who are like-minded and enthusiastic should be welcomed with open arms. The end result may not be exactly what I had intended, but the possibility for the work to grow beyond the initial expectations of a single mind is far greater.

4.5 Proprioception

Proprioception is the second piece in this portfolio which uses the Kinect as an interface for triggering and manipulating sound objects. It was composed after reflecting upon what I felt were the shortcomings of *Kinesia:The Third Law*. One major issue I had with that composition was that the original performer was unfamiliar with the technology being used which meant that while the physical movement within the performance of the piece was strong, the level of control that the dancer commanded over the compositional system was somewhat lacking. As I have mentioned in chapter 4.2.5, I began performing *Kinesia* myself with a good degree of success. After several performances of *Kinesia*, I decided to write a piece that was less improvised and more focused on subtle movement rather than motions which are spread across the entirety of a stage. *Proprioception* is the result of these experiments.

4.5.1 Concept

The performance aspect was not as controlled in terms of movement, as I am not trained as a dancer, but the overall spectacle of the piece, physical movement and audio material combined, was much more potent. Whereas *Kinesia* was a composed improvisation, *Proprioception* was constructed in such a way that the musical gestures within the piece could be replicated with a fair level of accuracy. The mapping scheme is almost identical to that of the one used in *Kinesia*.

With this in mind I composed *Proprioception* to be performed by myself and I constructed a compositional system that relied less on large movements within the performance space and had more of an emphasis on finer control of software instruments through minimal movement. One major inspiration for the physical aspect of this piece was *Butoh*, the dance form which was was developed in post world war 2 Japan by Tatsumi Hijikata (Toshiharu and Parsons, 2003). This dance form is based around the paradigms of Noguchi Taiso which, according to Paola Esposito and Toshiharu Kasai, "appears not to rely on the practitioner's intention to move, but that he or she is moved by invisible forces" (2017). Another main concern of Butoh is to focus on the smaller, more controlled movements as opposed to grand physical gestures (ibid.). Although I was inspired by these elements of Butoh, I did not fully engage with all of the conceptual underpinnings of this form such as the use of body paint and the celebration of the grotesque.

Conceptually speaking, I was inspired also by the shamanistic practises and the evocation of spirits that I had read about in the anthropological work *Soul Hunters: Hunting, Animism, and Personhood among the Siberian Yukaghirs* (Willerslev, 2007). This work explores the fluidity of personhood among the shamans of the Yukaghir people of Siberia, and how they believe that they literally become the animals in their environment when evoking certain spirits. My goal was to give the impression of summoning sound objects from an unknown, ethereal environment and have those sound objects affect the subsequent behaviour of the performance so this conceptual framework seemed quite apt for this piece.

4.5.2 Instruments

I created a total of thirteen sound objects which may be summoned throughout the piece. These sound objects are triggered in a manner similar to that of *Kinesia* but the triggering system and conditional requirements in this context are slightly more sophisticated. I will now describe these characteristics of these objects and review the conditionals specific to each individual object in detail.

Ball

The first instrument, Ball, is triggered when the left hand of the performer is directly above the right hand and both hands are in a position roughly in front of the navel. This is the opening musical gesture and it dictates the first position of the performer. The sound object itself is not synthesised but a short sample of a low frequency rumble created in a separate Csound script. The frequency content of this instrument is linked to the distance between the the left and right hands; as the distance between the hands grows, the center frequency of a resonant filter is increased. The resonance of this filter is also linked to the position of the left and right hands. The spectral centroid of the sound object is determined by the initial position of the left hand along the y-axis. This rises over a period of twenty seconds (if the instrument is active for that long) to have a spectral centroid four times that of the original value. Once the distance between the two hands along the x-axis crosses a predetermined threshold the instrument is disengaged. Due to the fact that this instrument can be easily triggered each time one hand passes over the other, there is a possibility of triggering this instrument unintentionally. To counteract this problem, this instrument can only be triggered within in the first 60 seconds of the piece.

Whisper1

The second instrument to be introduced is called Whisper1. This instrument takes a very short sample of a vocal recording and sends it through several delay lines. The feedback level for these delay lines is controlled by the position of the left hand along the z-axis and the maximum delay time is determined by the position of the same hand along the y-axis. This instrument is triggered when the speed of either the left hand or the right hand crosses a specified threshold.

Whisper2

Whisper2 uses the same instrument design as Whisper1 but the overall behaviour is quite different. Firstly, the manner in which this instrument is triggered is unlike Whisper1 as it does not rely on speed as an event cue, instead the movement into a specific posture is what instigates this instrument. When the left hand is positioned in front of the naval and the right hand reaches out, crossing a predetermined threshold along the z-axis, Whisper2 is engaged. As in Whisper1, the feedback level and maximum delay time of the instrument is controlled using the position of the left hand along the z-axis and y-axis respectively.

Vox

The instrument Vox uses the addition of multiple banks of modal frequencies to emulate a human voice. My initial experiments failed when trying to achieve the exact effect

(a) Ball (b) Whisper2 (c) Vox

Figure 4.6: Postures required to trigger the Ball, Whisper2 and Vox instruments[6]

I desired but after using a technique demonstrated by Dr Iain McCurdy in one of his many Csound examples on his website (2012). In this example, he stores the amplitude, center frequency and bandwidth information of each formant in separate function tables which can then be accessed remotely by another instrument. This facilitated a much more streamlined design for the Vox instrument in Csound. The actual triggering of the Vox instrument is done through the inverse posture described for triggering the Whisper2 instrument; the right hand is positioned in front of the naval and the left hand reaches out, crossing a predetermined threshold along the z-axis.

Glass

This instrument was created using, as the name suggests, recordings of stepping on broken glass. The pitch of this sample is varied through the use of random spline curves. The total user input for this input is confined to controlling parameters of two filters which the sample is fed through. The filters are a emulation of a Moog diode filter and a Butterworth band-pass filter which are arranged in series. Each of the stereo channels is given its own dedicated series of filters. This instrument is triggered by the performer raking the left hand across their head. The instrument is disengaged by performing the exact same motion with the right hand.

[6]These images are all taken from a recording of Proprioception at the *Sound Thought Festival*, Glasgow, 2017

Keys

The Keys instrument is also named after the source of the sound object. The design of the instrument is quite similar to that of the Glass instrument too but with the addition of dynamic start and end loop points. In this instrument, the start point of the loop is defined by the position of the left hand along the x-axis and the end point of the loop is defined by the position of the right hand along the x-axis. The natural proximity of the left and right hand serves as a pragmatic way of ensuring no unexpected or out-of-bounds values are received into the system as it is difficult to have ones hands in a position whereby the left hand and right hand are simultaneously at their opposite extremes along the x-axis. This sound object is created by raising the left hand above a predetermined threshold along the y-axis. The motion required to trigger this instrument is purposefully designed to evoke a beckoning or summoning of a entity from above.

Bang

In order to employ harmonically dense attack-decay gestures within this piece, akin to those discussed in chapter 4.1, I created the Bang instrument. The title of the instrument is indicative of the audio output it provides, and it serves to bookend several other musical gestures within the piece. It is triggered when the speed of the right hand and its position on the y-axis cross predetermined thresholds. Bang also functions as a trigger for another instrument in this piece, meaning that its purpose is not only tied into the musical and physical progression of the piece but also the control of flow within the compositional system. The conditions for the triggering of this instrument call for a physical gesture which befits the dramatic nature of the musical gesture it produces. This is a direct effort to address the issue of source bonding (Smalley, 1997) and considerations regarding action trajectories (Godøy, 2003).

Input4

Input4 is designed around the `fog` opcode within Csound. FOG synthesis is similar to FOF (fonction d'onde formantique) synthesis in that it creates streams of grains which

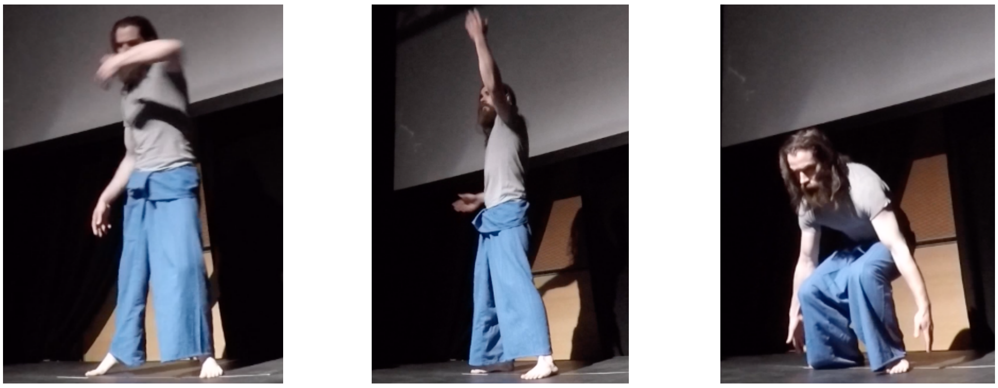

(a) Glass (b) Keys (c) Bang

Figure 4.7: Postures required to trigger the Glass, Keys and Bang instruments

are then shaped using local envelopes to model formant region found in the human voice and other instruments (Clarke, 2000). In the case of FOF synthesis, stored sine-waves are used for synthesising formant regions. FOG synthesis enables the use of sound files for synchronous granular synthesis. I was experimenting with the `fog` opcode in Csound, trying to create the sound that would eventually be provided by the Vox instrument, when I stumbled upon a strange audio output that did not behave as I had expected, but nevertheless provided me with an interesting voice to include within this composition. The sonic characteristics of Input4 were created through manipulating the transposition factor and the density of two separate FOG instruments, one for the left channel and one for the right. There are a number of parameters which are controlled at performance time with the `fog` opcode but in the interest of reproducibility I settled upon specific parameter settings and assigned static values to them. The main manipulation of the sonic output of the instrument is provided by the motion of both hands along the z and y axes.

Dis

The three instruments with the Dis prefix, Dis1, Dis2 and Dis3, are all derived from a sound object featured in the tape piece *Disintegrate* (chapter 4.3). The pitch of the first two of these instruments is randomised between a range of 50% of the original and 150% of the original sound object in each case. The first Dis instrument is triggered when the the right hand crosses a threshold along the x-axis. The motion required to cross this threshold forces the performer to reach outwardly from their base position, alluding to the idea that this sound object is being pulled into being from afar. The manner in which the second Dis instrument is triggered mirrors that of the first, using the left hand instead of the right. The spatial positioning of these objects in the stereo field is controlled by the right hand and left hand for Dis1 and Dis2 respectively. The third Dis instrument is triggered by raising both hands above a defined threshold along the y-axis. The pitch of this particular instrument is fixed at a static value and now spatial control is afforded to the performer. Each of these three instruments act like a set of switches which, when activated sequentially, prime the compositional system for the final sound object.

(a) Dis1 (b) Dis2 (c) Dis3

Figure 4.8: Postures required to trigger the three Dis instruments

Mod

This instrument acts as a tonic to many of the other instruments in this piece in the sense that it is far less abrasive or percussively driven than say the Dis instruments or the Bang instrument. Like Input4, Mod is the result of a failed experiment when trying to synthesize a vocally inspired formant instrument. I made use of the `ModFM` (Lazzarini and Timoney, 2011) opcode within Csound to in conjunction with some spectral filtering to achieve the sonic characteristics of this instrument. Mod is engaged by the performer lowering their head below a specified threshold whilst ensure that the speed that either hand is moving does not exceed the threshold defined in the Bang triggering instrument. Mod also serves to disengage the Vox instrument and careful consideration was given to crossfading both of these sound objects to ensure a smooth transition between them.

Colossus

The final sound object that I designed for this composition was created to serve as a coda for the entire piece. I spent longer designing this sound object than any other within *Proprioception*. There is no real-time performance control afforded to the performer within this instrument design. It is simply to be triggered and left to evolve in a preordained manner. As I have mentioned in the previous section, this instrument must first be primed before it can be engaged. The reason for these three fail safes to be put in place is that after the Colossus has been triggered, all user control over the triggering systems are relinquished. The only system control which remains is to subtly manipulate the bandwidth of a filter by changing the distance between the left and right hand. It also forces all other instances of instruments to decay leaving just the Colossus itself to herald the end of the piece.

The actual design of the sound itself is based around a gesture created using Trevor Wisharts's pichstak technique (Wishart, 2000). An attack-decay trajectory is used but the decay is considerably extended beyond that of the Bang instrument. I wanted to create a sound that evoked an idea of a great amount of power being released from the system and my mind was brought back to when I first saw *The Fellowship of the Ring* in the cinema. There is a scene in the opening few minutes of the film where the antagonist, Sauran, brings his mace down hard onto the ground causing a huge shock wave to propagate through the battlefield. The sound design which accompanies this action has stuck in my head since, a very low frequency, percussive burst with a spectral centroid that descends over the space of a couple of seconds. Colossus was created by layering multiple percussive sound sources over one another and adding a harmonically rich texture underneath to simulate the disturbance of particles that would be caused by such an explosion of energy.

Figure 4.9: Posture required to trigger the Colossus instrument

4.5.3 Structure

The structure of this piece is defined by how each of the aforementioned instruments flow into one another. The measures put into place to establish a strict set of controls over how the instruments behave, in a sense act like a score. For example, the Ball instrument can only serve as an introduction to the piece based on the fact that one of its conditional triggers is time-based, only allowing the instrument to be used within the first minute of the piece. Instruments such as Input4 rely on other instruments (in this particular case, Bang) to act as their triggers. It is the coarticulation of one sound object flowing into the next that creates musical gestures from groups of sound objects which in turn create the overall form of the entire piece.

Godøy states in his article *Chunking in Music by Coarticulation*, "in general, coarticulation is about continuity, about movement as continuous and about the human body as made up of interconnected effectors" (2010, p. 7). It is this sentiment that informed the manner in which I approached the arrangement of this material. When considering how one sound object may transition to another I also had to be aware of the physical positions that I was moving in and out of. As a result of this, the construction of the triggering system, the choreography of the movements and the composition of the sound objects all developed concurrently.

The awareness of coarticulation in the context of musical performance not only helped me to create gestures which were intended to flow into one another with a certain sense of logic and intuition, but has also served to help overcome some of the performative issues that exist when presenting electronic music in a live context. It has been noted that the lack of coarticulation in digital instruments may be one of the reasons that they can often seem to be “unnatural” in how they behave in live performances (ibid.). The use of coarticulation, coupled with a degree of source bounding (Smalley, 1997) is my attempt to circumvent these issues in this composition.

Score

The score presented in Figure 4.11 provides a rough outline of the overall progression of the piece. There are five sections presented, each containing a sequence of symbols. These symbols are to be read from left to right starting with the topmost section and working downwards. Each of these symbols is described in the lexicon shown in Figure 4.10. The score acts as a performative frame, within which a certain level of play is permitted to take place. The exact duration of events (with the exception of the Ball instrument), and to a certain extent the manner in which they are articulated, are open to the interpretation of the performer, but ultimately the sequence of these events should be adhered to. The prominence of instruments such as Vox and Bang are shown through the symbols size, shape and opacity. The capacity to create additional instances of instruments is afforded to the performer, particularly where shading is present.

Proprioception Lexicon

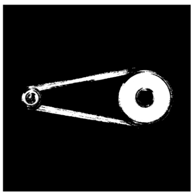

Ball
This symbol represents an instance of the ball instrument. The size of the dark circle within the white circle indicates the presence of high frequency content. The size of the ball indicates amplitude envelope.

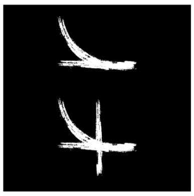

Whisper 2
The topmost symbol is a instance that is played for any significant duration. A vertical line through any symbol (as seen in the lower symbol) indicates that it should be immediately silenced

Whisper 1
Symbol used to represent an instance of the Whisper 1 instrument. Can be triggered by motion with either the left or right hand at a speed that exceeds a specified threshold

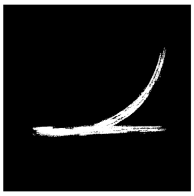

Vox
The Vox instrument is triggered in a posture that mirrors that of the Whisper 2 posture, therefore the symbols used to represent them are also mirrored. The size and opacity of the symbol in the score roughly indicate the prominence of the instrument

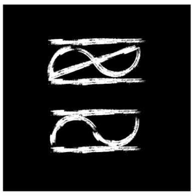

Glass
Both of these symbols are used to indicate an instance of the Glass instrument. The topmost figure is a stifled instance, meaning that the instrument is silenced immediately after its instantiation

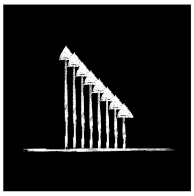

Keys
Much like the Dis instruments, the symbol for the Keys instrument somewhat conveys the nature of the gesture required in order for it to triggered. In this case the performer raises an arm above a specified threshold and slowly draws it back downward.

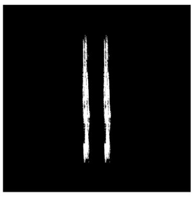

Silence
Two vertical lines indicate that the performer should allow for any instruments to fade and allow for a moment of silence between gestures

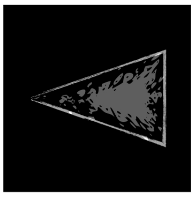

Mod
The mod instrument is to be played for an undefined of time. This acts as a precursor to the Input4 instrument

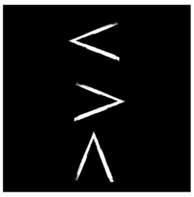

Dis
Each of these symbols represent Dis1, Dis2 and Dis3 as read from top to bottom. The direction that symbol to indicates the direction of the threshold crossing needed to trigger each instrument

Repeat
Traditional music notation symbol indicating that a section should be repeated. It should be assumed that the repeated section will not be an exact replica of the first instance of that section.

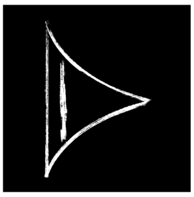

Bang
The symbol used is a comment on the amplitude envelope produced when the instrument is triggered.

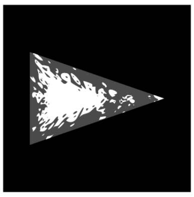

Colossus
The final instrument to be triggered in the piece. The presence of this instrument indicates that the performer does not have the ability to trigger any more instruments

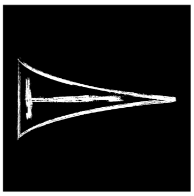

Input4
Triggered in much the same way as the Bang instrument. If the Mod instrument has engaged the Input4 instrument will be triggered instead of the Bang instrument.

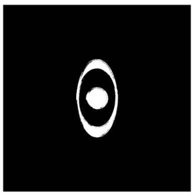

Coda
Based roughly on traditional notation, this symbol indicates the end of the performance. The shape is a reference to the Ball instrument that begins the piece. The gesture used once Colossus has been triggered mirrors that of the Ball.

Figure 4.10: Proprioception Lexicon

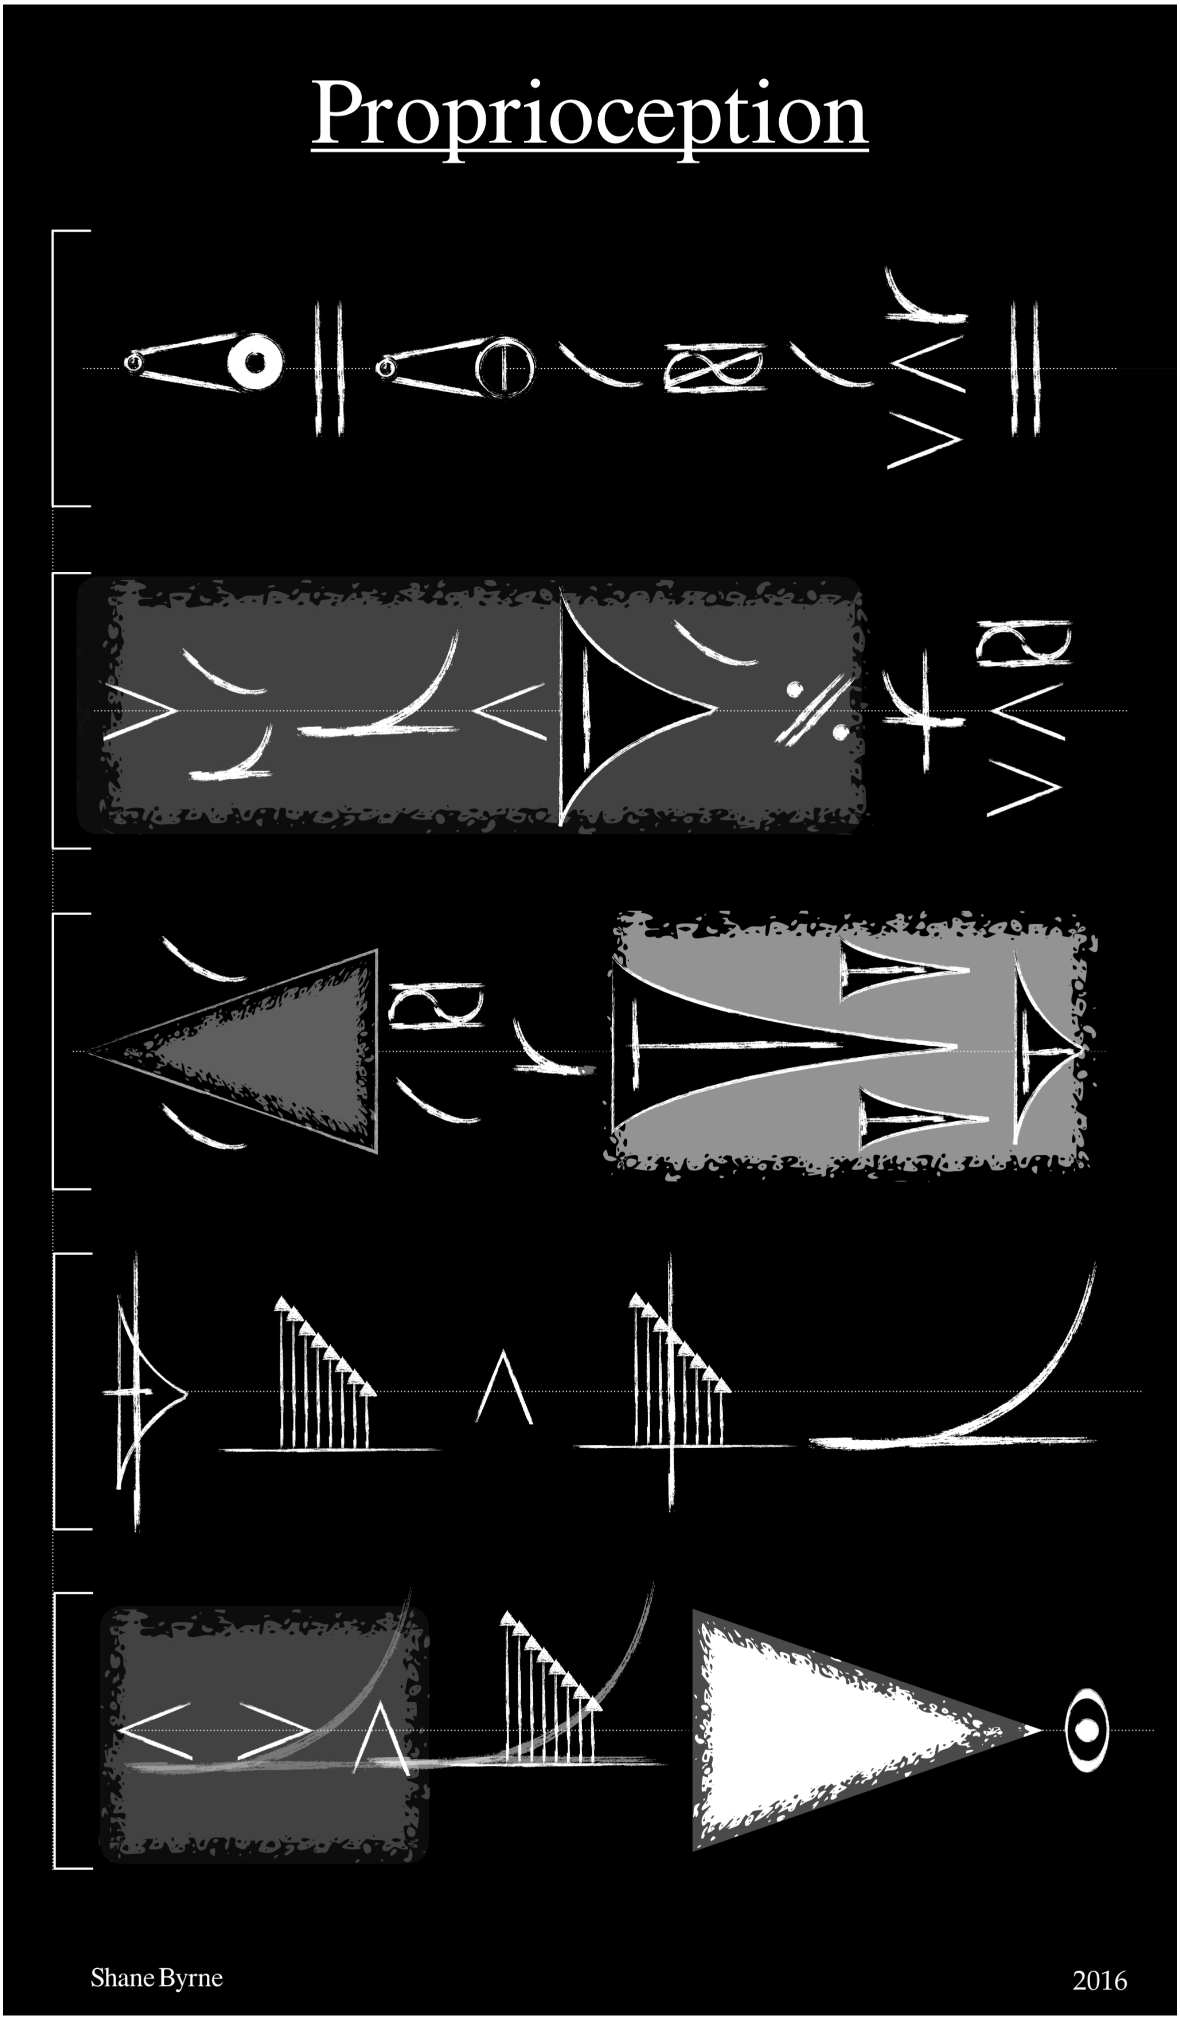

Figure 4.11: Proprioception Score

4.5.4 Conclusion

Proprioception was written specifically to be performed by me and with this established early on in the compositional process I feel I was able to investigate the potential physical and musical output from the system in finer detail than when working with *Kinesia: The Third Law*. There are certainly a lot more affordances present in a system that the operator knows intimately and this was certainly true of the compositional system employed in *Proprioception*. In the many performances of this piece, the system never fully failed (as was the case with *Kinesia* on one occasion) partly due to the inclusion of various fail safes within the system but also due to the fact I had an intimate knowledge of all the nuances and idiosyncrasies within the system. When presented with something that I may anticipate as a potential issue (such as certain lighting conditions) I had enough knowledge of the system to counteract that issue either through subtle adjustments of the movements within the performance or through the recalibration of thresholds within the code. This could not have been done with the same amount of fluidity when another actor was involved.

One issue which was present in this piece and in *Kinesia:The Third Law* was the matter of triggering false positives or even worse, of physical gestures not triggering any event at all. This is a shortcoming of the technology used within the Kinect hardware and although I have put measures in place within the compositional system of *Proprioception* to reduce the negative impact on the quality of the performance the problem remains. After several performances of this piece under various different lighting conditions, I came to the conclusion that my reliance on the Kinect as the sole source of performative data collection needed to be addressed. For the system to be truly robust I required another means for collecting performative data. This line of inquiry led to the conception of *Conatus*, the final live performance piece in this portfolio.

4.6 Disintegrate

Disintegrate was originally written for the purpose of being featured on an album curated by Wicklow multi-instrumentalist, Kevin MacNamara. A number of Wicklow based composers, all working within a varied range of styles, were approached to contribute a piece of original music. No further thematic guidance was given beyond the title of the compilation album, *Disintegrating Consciousness*. The composers were free to interpret this theme in any way they saw fit. This piece soon began to take on a life of its own and the final duration of the piece exceeded the limits of the proposed compilation.

At the time, my work had been focused almost entirely on the composition and performance of live electronic pieces, so I saw this as an opportunity to revisit fixed media composition.

4.6.1 Process

I had recently finished building the Joystick instrument (see Appendix B.4) and I was aware that I would soon have to deliver it to the musician for whom it was built. In order to provide aural documentation of the Joystick, I began to perform several short studies for the instrument, free from any additional processing. The sonic material contained within these studies provided me with the initial source audio for *Disintegrate* and it is this material from which the majority of sound objects within the piece are derived. The unprocessed audio from the Joystick can clearly be heard in the first 30 seconds of the piece.

4.6.2 Structure

This piece follows a through-composed structure consiting of three sections. After the brief exposition of the source audio, the first instance of disintegration occurs. In this case the source audio from the Joystick is used to modulate a broadband, percussive, noise source. The transition from note to noise can be visualised when examining the peak frequency spectrogram shown in Figure 4.12. In this spectrogram, the harmonic sonic content generated by the Joystick at the beginning of the piece is clearly represented by the consistent lines of peak amplitude present within the first minute. After this it can be seen that consistent harmonic patterns cease to exist at around one minute and twenty seconds into the piece.

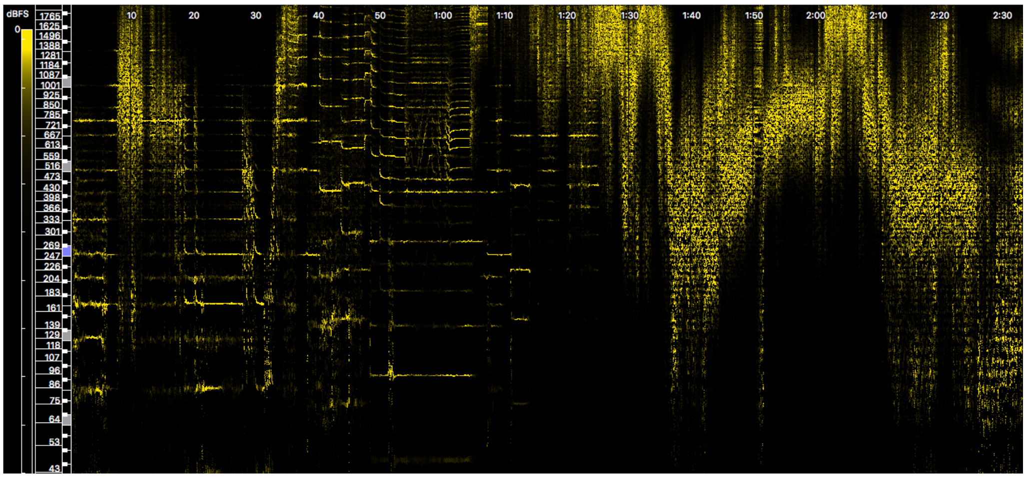

Figure 4.12: A peak frequency spectrogram demonstrating the evolution of the initial sound object

By manipulating the density of the noise source and downsampling the signal on the fly, a granular texture begins to emerge from the swelling gestures created by the Joystick. This evolution continues for roughly one minute until the connection between the processed and unprocessed signal begins to fully disintegrate. Once any semblance of the original source audio has almost completely disappeared, the granular texture ebbs and flows until finally it begins to swell into a wave of sound that passes over the listener and washes away the established grain based texture. This swell heralds the introduction of new source material.

The next section begins at 02:43 and the materials are processed in a manner that aims to facilitate a smooth transition from the disintegrated, granulated audio generated by the Joystick to audio obtained from completely different sources. The sources of this new audio material included short, discrete sounds obtained from manipulating small metallic objects that displayed a short attack and release amplitude envelope. These sound objects were specifically chosen to provide a counterpoint to the source material generated by the Joystick. In the case of the Joystick, a sound object with a graduated continuant trajectory is processed in such a way that it transforms into a stream of impulses, in the case of the material which follows, a stream of impulses are processed in a manner that elongates the sustained portion of the sound object. Here, the spectral typologies and the morphological archetypes (Smalley, 1986) of both audio sources are manipulated in such a way as to provide the antithesis to their respective natural states. This creates a contraflow within these first two sections of the piece, moving from note to node to noise and back again. The gestures created by the sound objects move from an initial state of graduated continuant to impulse and from an initial state of impulse to graduated continuant.

An entirely new landscape is introduced after the stream of impulses developed in the second section begin to ascend, paving the way for a *pichstak* percussive gesture which occurs at 04:30(Wishart, 2000). This gesture serves to provide a entrance for a new palette of sounds, all derived from the previously exposed audio material. The granulated texture derived from the Joystick is reintroduced but this time accompanied by a harmonically consonant sustained pad created using spectral filtering tools available within the Soundhack phase vocoder software.

4.6.3 Convolution Reverbs

From this point on in the piece I begin to use some convolution reverbs on some of the sound objects. The convolution reverbs I use were obtained from impulse responses of locations that were thematically relevant to the narrative of the composition. Since the theme of this piece was disintegrating consciousness, I decided to record impulse responses of locations in which the people may have experienced some sort of destruction or re imagining of self, or somewhere where their sense of self was challenged. After some research into suitable locations, I concluded that churches and psychiatric hospitals (preferably disused) would satisfy both the conceptual and sonic criteria that I was searching for.

Two locations specifically were used to capture impulse responses for the convolution. The first was a small chapel within a psychiatric hospital in county Wicklow. This hospital is still in use so the freedom to record here was limited. In spite of this, Dr Iain McCurdy, Dr Gordon Delap and I managed to capture some interesting impulse responses using multiple hand held recorders and a pair of binaural microphones. In order to capture these impulses, we burst balloons to create the broadband noise source necessary for the creation of impulse responses. In hindsight this may not have been the best choice due to the fact that some extraneous noise was created when parts of the burst balloon fell to the ground. Regardless of this oversight, some nice, natural sounding reverbs were created using these impulse responses.

The second location that was used for the capture of impulse responses was a disused psychiatric hospital in north county Dublin. I had scouted this location on a number of occasions but despite the fact that the hospital had been closed for some time, there was still a significant security presence meaning any guerrilla style recording was not feasible. I tentatively contacted the government body that was acting as custodian over the property to ask for permission to record impulse responses and, to my surprise, I was granted access without major issue. On a warm April morning, myself and Dr Gordon Delap drove to the peninsula where the hospital is located to spend the day capturing impulse responses.

The building itself was constructed toward the end of the nineteenth century and designed in the gothic-revival style of architecture. The building consisted of a central hub connected to several large rooms via long interconnected corridors with two wings extending either side of the main compound. The variance in room size and the presence of the long expansive corridors provided a rich variation of reverberant spaces.

On this occasion of capturing the impulse responses, I decided to use a short burst of white noise from a high fidelity loudspeaker as the impulse as opposed to popping a balloon. Many of the impulse responses captured at this location were a slightly darker than those captured in the previous location, and the spectral decay was a little less natural sounding, perhaps due to the large expansive spaces and the lack of absorbent materials. While it is hard to pick out the individual characteristics of these impulse responses within the piece due to interference and masking from other sound objects, a convolution reverb created using an impulse response of the X-ray room in the hospital (see Figure 4.13b) can be heard in isolation at roughly 06:18.

(a) One of the expansive corridors

(b) The X-Ray room

(c) A 'seclusion' room

Figure 4.13: Some of the spaces in the hospital where the impulse responses were recorded

4.6.4 Conclusion

The composition of this piece was mainly informed through aural discourse (Emmerson, 1986) and soon after early conceptual considerations had been investigated, my ear became the sole tool of divination within the piece. While allowing the Schaefferian practice of primacy of the ear to inform my compositional choices, I neglected to consider the fact that this piece had become too long in duration to fit neatly on a cross-genre compilation album.

I reassessed the piece on multiple occasions but resigned myself to the fact that if I was to abridge the work in any major way it would end up misrepresenting the work. After much deliberation, I decided that rather than reduce the duration of this piece to fit within the scheme of the album I would instead compose an entirely new piece under the same theme, whilst strictly adhering to a maximum duration of five minutes. I will discuss this piece in detail in the following section.

4.7 iAmAtEase

iAmAtEase is the second electroacoustic piece I have composed under theme of disintegrating consciousness. As discussed in the commentary for *Disintegrate*, my original contribution to the concept album was unsuitable. Rather than try to edit my original piece down to a shorter duration I decided to compose an entirely new piece. While the previous composition, *Disintegrate* was indeed inspired by the theme of disintegrating consciousness, I decided to take a slightly different approach to the composition of *iAmAtEase*.

4.7.1 Source Material

During the time when I was researching psychiatric hospitals in preparation for *Disintegrate*, I came across a story in an Adam Curtis documentary which dated back to middle of the cold war. It was a concerning a program of experiments conducted by Dr Donald Ewen Cameron in Montréal(Curtis, 2015). Cameron's theory, known as psychic driving, posited that conditions such as schizophrenia, anxiety and depression were caused as a result of suppressed memories.

Using a mixture of psychoactive and paralytic drugs, as well as electroconvulsive therapy, he would induce subjects into comas that often lasted for weeks, in an attempt to remove any suppressed memories which he felt may have been causing behavioral abnormalities in his patients. When the patients were in a comatose state, Cameron would play loops of noise and repetitive phrases from tape machines placed under their pillows, in an effort to 'reprogram' them (ibid.). These experiments did not yield the results that Cameron had expected, and in fact they were all complete failures. Instead the patients, who initially may have been suffering from only mild ailments, began experiencing severe memory loss, chronic pain and bouts of depression (Klein, 2008).

Although the majority of people subjected to these CIA funded experiments experienced severe memory loss, one thing that many of the patients could remember from the horrific experience was one of the phrases that Cameron played under their pillows on repeat; "I am at ease with myself" (Curtis, 2015). This phrase, recited once, is the sole source of sonic material in this composition. These words are broken down into individual particles through a mixture of spectral processing and granular techniques, creating an entropic representation of the original phrase.

4.7.2 Process

As I have already disclosed, the only source material used for this piece was a short recording of me reciting the phrase "I am at ease with myself". I recorded this phrase using an Electro-Voice RE20 dynamic microphone particularly because I found the proximity effect of this microphone to effectively increase the bass response whilst still retaining vocal clarity which is no great surprise considering that it is marketed as a broadcaster microphone(Wilkins, N.D).

I purposefully articulated this phrase using a technique known as "vocal fry" which has become prevalent among podcasters. Vocal fry, also reffered to glottal fry is described as an "aperiodic, staccato sound that's formed by compression of the arytenoid cartilages [...] which are at the back of the vocal folds [...] the voice becomes sort of a popping or creaking sound"(Purcell Verdun, 2016). The provided me with a large range of sonic characteristics within that single phrase, which I could then dissect and re-purpose throughout the composition.

The processes that I employed to further shape the source material included the use of a software vocoder, granular synthesis (performed using Csound and the LEAP motion sensor) and modal synthesis. Other processes such as re-pitching of material, reversing sound objects and micromontage were also used to achieve the end result presented in this portfolio.

4.7.3 Structure

This is a through-composed piece which focuses on the exposition, or more accurately the restriction of a full exposition, of the source material. Much like the narrative presented by the victims of Dr Cameron's experiments when discussing the difficulty they have remebering exactly the processes they were subject to, the formation of a concrete example of the source material is hazy, nebulous and constantly just out of reach of the listener.

The listener is initially presented with a series of water-like grains of sound which at first are formless. These grains give way to an impression of a more complete, coherent body of sound, first in the form of formants at 00:18, to more complete ululations at 00:30.

The phrase almost becomes recognisable at 00:58 but soon is lost among a sea of wash of granular textures. Particular words are draw out such they again become unrecognisable as can be heard when the word "with" is presented at 01:28. This section is closed with a faint shadow of the word "myself" at 01:47 before a new, harmonic texture is introduced heralding the next movement.

This movement begins to take advantage of the staccato nature of the "vocal fry" delivery and evokes a feeling of being underwater through the use of low-pass filtering of the signal. Having briefly introduced the staccato effect at 02:00, it reenters the foreground of the piece until the exposition of the word "ease" at 02:30.

A granular texture, reminiscent of the sound of ripping fabric, is developed until 03:13 when it gives way for a Xenakis inspired, high-pass filtered granular texture which calls to mind the charcoals used as source material in his work *Concret PH* (Xenakis, 1958).

This texture slowly evolves into a comment on the introductory texture at the start of the piece. When this passage dissolves at 03:38 expectancy of an upcoming gesture is created through a slight crescendo and brief hints at the staccato nature of the impending phrase. From 04:00 onward the piece gradually builds in both amplitude and harmonic content until the abrupt ending at 04:48.

4.7.4 Conclusion

It is not often that I allow a preconceived narrative to dictate the direction of one of my compositions in such a manner. My intent with this piece was to convey the frustration exhibited by victims of Dr Cameron's experiments whilst using that one poignant phrase as source material. The phrase, much like the memories of the aforementioned victims, is clouded by noise in some form or another. The pure, unsullied source material is lost forever behind a cloud of inharmonicity and complex signal processing.

This particular piece was composed under a specific time restraint; the piece had to have a duration of under five minutes in order for it to be featured on the concept album discussed in the commentary for *Disintegrate* and at the beginning of the this particular analysis. This restriction had a knock-on effect in that it not only dictated the total duration of the piece, but it also determined the duration of each movement, each phrase and each gesture. I wanted to avoid a sense of rushing through passages just to meet to duration criteria but I also wanted to give each element of the source material time to breathe and allow for a sufficient exposition of each part of the original source material. My ability to do this was greatly enhanced by restricting the amount of material I had to work with initially. Self-imposed constraints can often yield the most unexpected and intriguing results due to the fact that deep exploration of the source material is imperative to building a rich palette of sound. This notion of self-imposed constraint is a predominant feature in the subsequent compositions featured in this portfolio.

4.8 Djinn

Djinn was created as the result of attempting to compose a piece guided by a set of a self imposed limitations. These rules do not have the same function as the serialistic rules employed by Schoenberg, Stockhausen, Eimert and Goyevarts. In fact these rules do not seek to address the structural nature of the piece whatsoever. They were merely created to counter the paradox of choice that I have previously mentioned. By the time I got to the stage of writing this piece I had accumulated quite an large range of tools and techniques with which I could generate material so these rules were simply put in place to restrict the amount of these tools that I could employ. I also chose my source material from the two objects nearest to me at the time of inception, a plastic wrapper and a harmonica. Another reason for composing Djinn was to step away from the corporeally driven live performances that I had been working with at the time.

4.8.1 Process

Although I was very much enjoying the materials I was producing using these gesturally informed pieces, I wanted to take a break from the heavily technical side of electroacoustic composition and create something that was informed as much through aural discourse as possible (Emmerson, 1986). I wanted to essentially create a small bank of sound objects and allow them to inform my arrangement. With this small bank of sounds I would create sound objects using some simple techniques within Abelton Live.

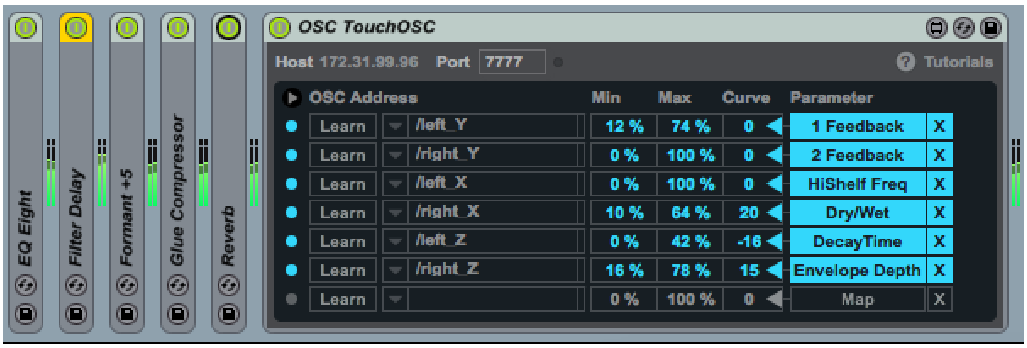

Figure 4.14: An example of controlling audio unit parameters in Live9 using the LEAP

While I sometimes find working within a DAW can sometimes be restricting in terms of the transformative processes afforded to the composer (as opposed to the seemingly inexhaustible processes available within audio programming environments such as Csound or Supercollider), I do find the workflow to be a lot more efficient when processing and arranging audio on the fly. As a result of this strategy, much of the material used within this piece was generated and arranged in a relatively short time period. I mapped the LEAP motion sensor to several parameters of the audio units within the signal chain and performed musical gestures on the fly (see Figure 4.14). This was achieved through the use of Processing for the initial collection of Cartesian data from the sensor and a Max For Live OSC receiver.

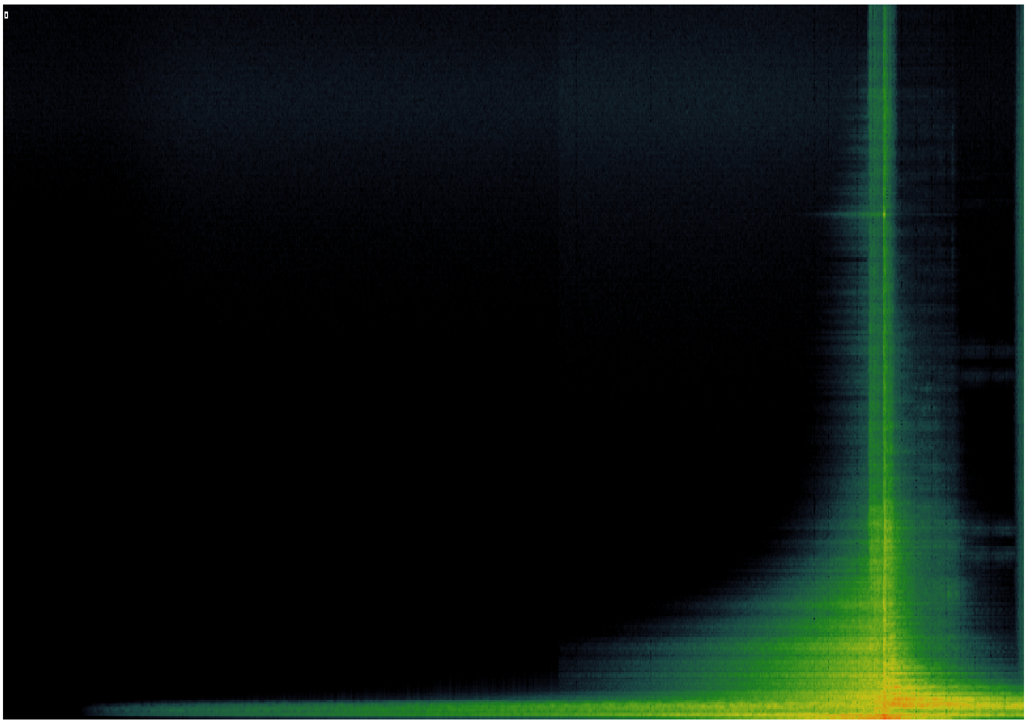

Figure 4.15: Djinn: Expectancy preceding attack-decay gesture

4.8.2 Structure

Djinn opens with a slow build of expectancy (Smalley, 1997) which is released through an attack-decay gesture at thirty-one seconds into the piece (see Figure 4.15). This, conceptually at least, performs the role of the 'big bang' moment in this piece, all subsequent musical gestures are informed and contexualised through this singularity. The sounds that follow this gesture are to be conceived of as sonic motes, created by the preceding attack-decay gesture. It is from these motes that more complex 'life forms' emerge.

Attack-decay style gestures that are common throughout my work are not as predominant in this piece. With *Djinn* I tried to create smooth and subtle transitions between objects, gestures and phrases in a way that can be described analogously as the *sfumato* painting technique.[7] Stark contrasts and sharp edges make way for subtle gradients between the featured sound objects, placing the emphasis on gradual morphologies as opposed to distinct gestures.

4.8.3 Conclusion

I often tend to use densely packed, percussive gestures, especially when composing fixed media material, and in composing this piece I attempted to create a sonic landscape where the sound objects evolved at a pace much slower than usual. I elected to consciously avoid impact-decay gestures in this work and instead focused on slow transformations and graduated continuant morphologies.

The self imposed restrictions regarding source material and the tendency toward a select few specific gestural shapes within this piece, lend themselves to an artistic freedom of sorts. Time and time again I have found myself overwhelmed by the sheer volume of options (both technical and sonic) available to me when beginning a composition. Setting self-imposed restrictions from the onset has proved to be an effective compositional strategy and effectively negates Schwartz's paradox of choice.

[7]The literal translation of *sfumato*, to evaporate like smoke, seems to fit this piece particularly well.

4.9 TenterHooks

Tenterhooks acts as the antithesis to *Djinn* in the sense that, where *Djinn* can be described as being an exercise in restraint in terms of the delicate nature in which the sound objects are treated, *Tenterhooks* is an exposition of bold gestures which tend interact with one another in a capricious fashion.

4.9.1 Process

This piece, perhaps more than any other work within in this portfolio, exemplifies the full gamut of my journey as a composer over these last four years. The source material was created in January of my first year of research and is comprised entirely of short recordings, ranging from 1 to 4 seconds in length, of me playing the disembodied comb from an old music box. Each of the 20 small clips were performed using a screw as an excitation source.

A file player, based on the looping instrument used in the *Ouroborus* pieces, was then used to trigger these sound files.[8] The exact timing and duration of these sound files were controlled using a pair of GUI sliders within the CsoundQt front end, each of which was then modulated by a dedicated pseudo-random counting device. The level, attack and decay characteristics of the amplitude envelope for each sound was also determined through the use of virtual chance operations. The audio output from this file playing instrument was then bussed to a set of auxiliary devices comprised of a reverb unit, a delay unit and a Csound version of an algorithmic beat processing device developed by Nick Collins for the audio programming software Supercollider (Collins, 2002).

In each of these auxiliary devices, the values for the majority of parameters were chosen using indeterminate processes. The range of values that each of these parameters could be assigned was gradually refined after many meticulous experiments with each device. The large number of pseudo-random processes in use within this system provide countless permutations which create a whole spectrum of idiosyncratic and nuanced possibilities of sound when viewed on a micro level, however the total systematic output,

[8] A full description of this particular compositional system is provided in Appendix C

when viewed at a macro level, is measured and circumscribed to a preconceived idea. Through the subtle tweaking of the maximum and minimum bounds for certain parameters, an approximate prediction could be made as to what the general sonic output from the system might be.

I generated a large amount of material using this method, a very small portion of which was used in one or two of the earlier compositions in this portfolio, but the majority of this source material lay untouched for over two years. I rediscovered the source material generated through this system when searching through an old hard-drive and decided to reimagine it using techniques that I had recently been experimenting with in Abelton Live. I had begun using Live instead of Pro Tools because of its OSC compatibility and I was having some success mapping VST parameters to incoming data from the LEAP motion sensor (see Appendix A.5). I refined over two hours worth of material into a source file of fifteen minutes. I then prepared five iterations of that source audio, repitching the material in octaves and minor thirds above and below the original material. A range of devices, including resonant filters, signal vocoders, various reverb and delays units, were inserted into the signal chain of each of these five tracks. A mixture of parameter automation and gestural input from the LEAP was used to control the behaviour of each device in the signal chain of the tracks.

4.9.2 Structure

After a number of performances of this material I began to pick out small motifs around which I could begin structuring the sonic landscape of the piece. These performances were once again a lengthy exercise in refinement but as I grew more familiar with the progression of the material phrases began to emerge from the interplay between the sound objects which highlighted and commented on the emerging motifs. This process served as a further distillation of the source material and from it I identified fourteen phrases, some of which had several variations of their own. These phrases, variations included, would serve as the sonic scaffolding for the final composition.

The overall shape of the final piece follows the progress of transformations that I made to the phrases and motifs. The structure is comprised of three sections. The first section is concerned with exposition of materials and processes. The initial motif is reiterated throughout this section in varying stages of development. Only very subtle transformations are performed in this first section, mainly comprising of adding delays and reverbs to the sound objects created by the *Ouroborus* looping instrument. One tool that I use quite a bit in this section and indeed in this piece overall is the freeze feature of the native reverb plugin in Live 9. This allows the user to capture the spectral content of a transient sound and pass it through an infinite feedback loop within the reverb algorithm. This particular feature inspired some of the instrument design in *Conatus*, the final corporeally driven piece of this portfolio.

Various iterations of these sound objects were repitched, both upward and downwardly, and subsequently superimposed on top of one another. The nature of the *Ouroborus* instrument is to create cyclical patterns of sound and as a result of this, I was inspired by the colotomic musical structures of southeast Asia (Becker, 1968). The serendipitous interactions of these superimposed sound objects were how the initial motifs within this piece were discovered. Further modulations were performed using some synchronous granular synthesis and the vocoder in Abelton Live 9.

The second section is introduced at 03:42 and is concerned with developing the previous phrases through synchronous and asynchronous granular synthesis. A portion of material, including the motifs established in the first section, were processed using the `fog` and `grain3` opcodes in Csound. Using the Csound frontend Cabbage, I was able to 'perform' manipulations of the material in real-time in reaction to some of the phrases created in the previous section. This section terminates with a dense culmination of sound objects at 05:53.

The final section is processed using the compositional system developed for *Conatus* (see chapter 4.10). I was struggling with how best to treat the material after the second section and found that many of my attempts to resolve the piece ended in passages that sounded either too contrived or too measured. In an effort to *play* the material more and think about it less I routed several motifs into the *Conatus* compositional system. I will discuss this system in more detail in the following section but for now it will suffice to say that the control of musical parameters in this system is governed by the Kinect and a glove interface that I developed. This allowed me to intuitively interact with the sonic material and circumvent any issues I was having with possibly over-intellectualizing the creation of musical gestures. The ability to simply react to, and shape a musical landscape through embodied means, resulted in the formation of a musical passage that faithfully represented the direction that had existed in my imagination but the execution of which had eluded me up until that point.

4.9.3 Conclusion

The manner in which I created the sonic materials for this composition is analogous to casting broad paintbrush strokes and waiting to see what images begin to emerge from the canvas. I took a maximalist approach to the initial compilation of sonic material, generating masses of sound objects and layering multiple audio tracks until musical phrases began to to reveal themselves.

Once certain motifs had been established, and I had developed a number of strong musical gestures, I began the process of refinement. This entailed sculpting and further processing the raw sonic material chosen from the initial source audio. The tools with which I carved out the majority of sound objects featured in this piece, were all developed with live performance in mind. These tools allowed me to interact with the sonic materials of *Tenterhooks* in a fluid and intuitive manner, and I feel that this piece exemplifies the physicality of performance that Michel Waisvisz referred to when speaking on aspects of the compositional process (Krefeld and Waisvisz, 1990).

4.10 Conatus

Conatus, the final live electronics piece written for this portfolio, was composed with the aim of recreating the level of collaboration and sharing of compositional strategies that a musician can experience when playing as part of an ensemble. To expand upon this notion, *Conatus* was an attempt to address the fact that electroacoustic composition is often a solitary exercise.[9] This is quite different from my experience of creating other forms of music through simply *jamming* with other musicians. This piece therefore was an attempt to marry the spontaneity present in the context of *jamming* with the fine level of control over sonic materials present in fixed media composition. It uses a compositional system which functions in a similar manner as those used in *Kinesia* and *Proprioception*. In addition to using the Kinect to track my movements, I included another interface within the compositional system with a view to grant the performer of the piece enhanced control over musical gestures. This came in the form of a glove fitted with various sensors (see Figure 4.19), all attached to an Intel Edison which in turn was sending packets of data to Csound via OSC. The goal of introducing this tool was to decrease the possibility false positives and missed triggers and therefore increase the robustness and reliability of the overall compositional system. I will discuss the glove in more detail later in this chapter.

Another major difference between *Conatus* and my previous live electronic pieces was that I would not be creating the source material myself. With *Kinesia* and *Proprioception*, all sound objects were created and synthesised as a result of the movements of the performer. In *Conatus*, all source material would be provided by Ror Conaty, a percussionist that I have collaborated with extensively in the past. In an echo of the conductor's role in an orchestra, I would not actually have the ability to create sound myself, I could only modify and manipulate the source material created by another musician.

[9] Although I have actively tried to circumvent this in other pieces within this portfolio

4.10.1 Source Material

The source material for *Conatus* is provided by an assemblage of various idiophone instruments which are played using both traditional and extended techniques. The instruments contained in this assemblage evolved over a number of months of rehearsals. The initial instruments used as source material consisted of a snare drum, a kick drum, toms, cymbals and hi-hats, all of which make up a traditional drum kit. Soon after beginning to rehearse using these instruments as source material, it became apparent that the traditional drum kit would be unsuitable for the purposes of this piece. The source material that it provided was far too bombastic for what I had in mind for *Conatus* and the playing style which Ror would slip into through force of habit created few opportunities to create the textures and sound characteristics that I enjoy working with. Slowly we began to remove elements such as the hi-hats and toms and replace them with bongos, tablas and a metallophone.

Figure 4.16: Finalised assemblage of source material

The snare drum was modified so as to provide a wider range of timbral possibilities. A hole was burnt in the center of the drum and through that hole a bass guitar string was fed, with a coin acting as a washer to prevent the string from coming all the way through the hole. This string could then be pulled taut to allow Ror to strike and bow the string at will. Various bells, chimes and mutes were also placed upon the top skin of the snare drum alongside a piezo contact microphone. This microphone was fed into an amplifier which I built based on a design by Alex Rice (Poff, 2015).

Figure 4.17: Prepared snare complete with bass string running through the center of the top skin

The pre-amp was built to take a high impedance balanced input signal and output a low impedance balanced signal. This impedance matching is responsible for preserving much of the low frequency response which would otherwise be lost if the piezo was plugged directly into the mixing desk. Another advantage of this pre-amplifier is that it balances the signal coming from the piezo meaning that the signal is far more robust and can be transmitted across larger distances than in its usual, unbalanced state. Due to the fact that the audio from the snare was now almost exclusively being captured by the contact microphone, Ror swapped out his regular drumsticks for a selection of skewers and chopsticks, all of which were far lighter and therefore had less chance of overloading the signal coming from the contact microphone than his regular drumsticks.

Several other percussive elements were added to the kit, including stacks of small cymbals and a cowbell. As time went on, idiophonic instruments were added and taken away until the kit was refined to our exact needs. Many of the more resonant drums from the original kit had been replaced by drums with a far shorter decay and the common

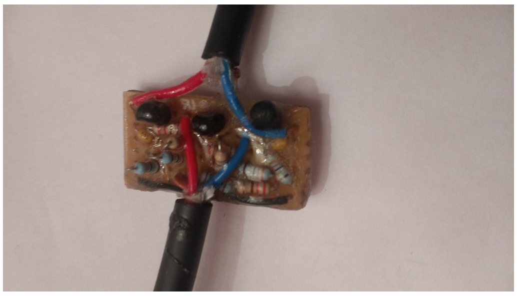

Figure 4.18: The pre-amp used with the contact microphone

coupling of crash and ride cymbal was swapped out for the cymbal stacks and a pair of cowbells.

Another point to note when discussing the creation of source material for this piece is the use of bowing as an excitation source. The timbres created when bowing cymbals is something that I have enjoyed for quite sometime and is a technique that has provided me with source material for some of my previous works[10]. I encouraged Ror to add a cello bow to his selection of skewers and mallets and we began experimenting with bowing cymbals and the tuned bars of his metallophone. I found the sounds created using this technique to be extremely interesting and they served to further expand the timbral possibilities provided by the source material.

As the instrumentation that would provide the source material became more refined, so too did the compositional system that I would be using to process the material. The goal was to be able to fluidly interact with the source material being provided by Ror as it was being played. I had been interested in implementing live sampling and loop manipulations using the Kinect for sometime and *Conatus* seemed like the perfect opportunity to investigate these ideas further. As I have mentioned in the conclusion of chapter 4.5, one issue that I was very much aware of was the unreliability of the Kinect as a standalone triggering mechanism. To address this issue I began to experiment with using wearable technologies to supplement the Kinect as a means for movement detection.

After observing footage of performances of *Kinesia* and *Proprioception* I noticed that most of the movement harnessed for musical means came from my arms and hands. In light of this, I decided to design a glove which would improve the overall robustness of

[10] *Cascando* and *Sinmara* are both examples of this within this particular portfolio although several of my earlier works also included textures created using bowed cymbals.

the compositional system. This would entail attaching an Intel Edison (see Appendix A.4) to the back of a glove. The glove would initially have a copper ring worn on the tip of each finger and the thumb (Figure 4.19). These copper rings would eventually be swapped out for conductive fabric (see Figure 4.20a). The copper on the thumb was attached to the +1.8V GPIO output of the Edison and each of the four fingers were attached to digital input GPIOs. The four fingers were pulled down to ground meaning that in their resting state they would return a value of 0 to the system. When the thumb came into contact with a finger, that finger would then receive 1.8V and return a value of 1 to the system. This provided me with four switches, each of which I could assign to the triggering of musical events or processes.

In addition to the four switches, I also made use of a 9DOF sensor attached to the Edison board (see Appendix 3.4.1). This facilitated the collection of data pertaining to the roll, pitch and yaw of the glove. Apart from providing three additional data sets which could be used to control further musical parameters, the 9DOF sensor also provided the glove with some level of modal functionality; when the glove is being used with the palm facing down a bank of four switches provided by the rings is available, turning the glove 180 degrees and having the palm face upwards enters a secondary mode, reassigning the functionality of each finger, essentially providing four additional switches.

(a) The top of the original glove with pouch for housing the Edison

(b) The underside of the original glove

Figure 4.19: Original glove design with copper rings

The Edison runs a Python script which collects the data from the GPIO pins and the 9DOF, and transmits that data via OSC over a wireless connection facilitated by a WiFi hotspot on a mobile phone. This enables connectivity without having to use a dedicated wireless router, reducing the amount of equipment required when travelling abroad to perform the piece.

To further refine the design of the glove, the copper rings were replaced with conductive fabric. The fabric was sewn into the existing material of another glove, each strip of fabric terminating in a male snap button. The hardwired connections from the Edison circuit were also swapped out for wires with female snap button connections (Figure 4.20b). This meant that the Edison module was now detachable, allowing for any modifications to the circuit to be performed with greater ease.

(a) Conductive fabric used to replace copper rings

(b) Modified circuit design with snap buttons

Figure 4.20: Improved glove design

The glove, coupled with the Kinect, provided relatively robust control over the compositional system. This was particularly desirable for this piece as I wished to have as much control over the system as possible and step away from the improvisatory nature so prevalent in the live performance of electroacoustic music.

4.10.2 Instruments

One idea that I wished to explore in this piece was textural counterpoint. When the source material provided short attack-decay trajectories it would be transformed into a graduated continuant and vice-versa. This style of transformation called for spectral freezing, infinite reverbs and other techniques for elongating the often transient nature of this particular source material. The following section will describe the exact processes employed to achieve these ends. The manner in which these processes are engaged is described in table 4.2 which follows the descriptions below.

Variable Loop

This was the software instrument that I included in the compositional system for *Conatus* and was inspired by instruments I had designed for both *Kinesia* and *Proprioception*. It involves capturing a loop of source material with one posture and then playing it back using another. The length of the loop is normalised and the start and end points of the loop are determined by the left and right hands of the performer respectively. This method of dynamically controlling the length of the loop means that visually the loop appears to be captured between the two hands of the performer, expanding and contracting as the hands are brought closer together and further apart.

Glitch Loop

In the case of this instrument, a loop is captured in a similar fashion as the Variable loop, but in this case the length of the loop is fixed and is defined by the amount of time that the recording function of the instrument is active. Once the loop is played back the timing is constant. The performer has the ability to change the volume of the loop and turn off and on the playback at will. There is also the option of cutting up the loop momentarily using the `bbcuts` opcode in Csound.

Sync Grain

This instrument uses an opcode within Csound that performs synchronous granular synthesis on a stream of audio stored inside a buffer. The functionality of the record and playback process is once again very similar to that of the previous two loop instruments, one posture records and writes audio to a buffer and through a slight modification of that posture playback is achieved.

The amplitude, grain pitch and time scaling, grain density and grain size of the sampled audio can all be modified by the movements of the performer. Elements of this instrument were based on a design by Victor Lazzarini, which include several mechanisms to ensure no out-of-bounds behaviour occurs (2014). For example, the maximum number of grain overlaps in the system is calculated by the product of the grain density and the frequency of grain generation. Exceeding the maximum overlaps can cause the system to crash, therefore there are measures in place within the instrument design to avoid this scenario.

Granule

Although this is another variation on a granular synthesis technique, the resulting audio output is quite different to that of the previous Synch Grain instrument. Granule makes use of the Csound opcode of the same name to create an asynchronous granular synthesis texture generator. The majority of parameters in this opcode cannot be modified in real-time but the amplitude, grain size and the density of the grains can be varied during the performance. While this does limit the influence that the performer has over the instrument, the nature of the process means that the range of textural possibilities available is quite broad. A small change in the source material will result in a completely different cloud of sound being generated through the `granule` opcode.

The grain size and duration are controlled by the roll and pitch of the performer's left hand respectively. As I have already mentioned, the `granule` opcode is limited in terms of real-time modulation but the performer can affect the pitch shift factor of the instrument by adjusting the initial position of their right hand along the x-axis when triggering the instrument.

Mince

This instrument makes use of an opcode within Csound called `mincer`, which performs phase-locked vocoder processing. Whilst experimenting with this opcode I came across a User Defined Opcode (UDO) by Joachim Heintz which works well with `mincer`[11]. This UDO is called `LpPhsr` and is used to create a time pointer when reading audio as a loop from a buffer. When used in conjunction with `mincer`, one can achieve a scrubbing effect similar to that of scratching a record. When the speed of `LpPhsr` is set to zero, the audio effect produced is that of a needle frozen to a single point on a record whilst still somehow producing sound. The amplitude and spectral characteristics of the sound that

[11]The LpPhsr UDO can be examined in detail at http://www.csounds.com/udo/cache/LpPhsr.udo

has been captured by `LpPhsr` and `mincer` are suspended in time and can be subjected to further manipulations.

The main modulation available to the performer with this instrument is the ability to control the rate of a sine wave oscillator, the function of which is to implement ring modulation on the signal captured from the source material. The roll of the performer's left hand is used to control the frequency of the sine wave oscillator. In addition to that, the audio produced with this instrument is fed through a band-pass filter, the band width of which is determined by the distance between the performer's left and right hands.

Infinite Reverb

This is an implementation of a feedback delay network (FDN) reverberation with a feedback level of 1. This feedback level creates a feedback loop, essentially resulting in an infinite reverb effect. Due to the nature of this effect, care must be taken to avoid feedback between the microphones used to capture the source material and the sound reinforcement system used to amplify the performance. The manner in which this instrument functions is based on the freeze feature of the native reverb in Abelton Live 9. When composing *TenterHooks*, I began to make use of the freeze button on this reverb and I thought it would be interesting to try and implement the same effect in Csound where I would be able use it within a compositional system that afforded greater performative control.

This instrument can be armed by making contact between the ring finger and thumb while the roll of the hand is greater than 180°. In order to feed audio into the FDN, the left hand of the performer must be in-line with their torso and their right hand must be above a predetermined threshold along the y-axis.

The amplitude of the instrument is controlled by the position of the performer's head along the y-axis and the cutoff frequency of a low-pass filter is determined by the position of the performer's torso along the z-axis.

Dynamic Reverb

The Dynamic Reverb instrument uses the same FDN as the infinite reverb instrument but in this case the feedback level can be modulated by the performer. The position of the right hand along the y-axis is used to determine this value, with the range being confined to be between 0 and 0.9.

Some spatialisation is achieved through the use of a digital emulation of a pair of Moog diode ladder filters. One of these filters is used for the left channel and the other is used for the right channel. The cutoff frequency of the left channels is the inverse of the right channel. This relationship creates a stereo panning effect that uses spectral content rather than temporal differences or discrepancies in sound intensity to create the illusion of stereophony.

PVFreeze

PVFreeze uses a phase vocoder within Csound to freeze the spectral components of an incoming signal without freezing the amplitude component of the signal. The resulting audio is a highly processed sounding version of the source audio that somewhat resembles the cathode ray recordings used by Ben Burtt for the sound of the lightsabers in Star Wars (O'Hara, 2015).

The spectrally frozen signal can be further modulated using the lateral movement of the right hand along the x-axis.

4.10.3 Structure

Although I had collaborated on previous works such as *Sound;Waves* and *Galilean Moons*, this was the first (for want of a better word) *true* collaboration of this portfolio. With *Sound;Waves*; the exchanging of ideas never passed beyond the conceptual stage, any decisions on the implementation, either technical or musical were largely left to me. In the case of *Galilean Moons*, although the discourse between the other composers and me was concerned with the musical frameworks in some cases, the end result is more of a triptych rather than single, unified landscape.

Finger	Roll	Switch	Instrument	Modulation Effect	Modulators
Index	$< 180°$	Latch	V. Loop	Loop start/end, Amplitude	L.Hand.x R.Hand.x, R.Hand.y
Index	$> 180°$	Latch	G. Loop	Glitch Effect, Amplitude	Index, R.Hand.y
Middle	$< 180°$	Latch	Grain	P.Scale, T.Scale, Dens, Freq, Amplitude	L.Hand.x, R.Hand.y
Middle	$> 180°$	Latch	Granule	Grain size, Density, Amplitude	Roll, Pitch, R.Hand.y
Ring	$< 180°$	Latch	Mince	Filter Width, LFO, Amp	Distance, Roll, L.Handy
Ring	$> 180°$	Latch	Inf. Reverb	LPF Cut-off, Amplitude	Torso.z, Head.y
Pinky	$< 180°$	Momentary	PVFreeze	Frequency Scale	R.Hand.x
Pinky	$> 180°$	Momentary	Dyn. Reverb	FB Level, LPF cutoff	R.Hand.y R.Hand

Table 4.2: Functionality of glove

With *Contaus*, the engagement with the percussionist, Ror, was almost exclusively in person, and the subject matter of conversations we had regarding the composition of the work were concerned with both conceptual and practical matters.

In order to efficiently converse about compositional strategies and ideas, we developed a system which acted like a shorthand for describing musical characteristics and behaviour. This shorthand seemed to manifest in three ways: firstly, we would could use other music as reference points to allude to a mood or style. One example of this would be how we began to demarcate the piece into three sections. Each of these sections, Eno, Reich and Autechre, were named after the composer or group whose style we were seeming to emulate (albeit quite vaguely at times).

Secondly, an onomatopoeic approach was used in which we would mimic the sonic content of a sound with our voices. An example of this would be the ululation of the sound *Gok* when referring to the sudden impact performed on the cowbell roughly two thirds of the way through the piece.

The third means of communication that we shared during the composition of this piece, was the use of graphics as a method of musical description. I created a graphical score soon after we began rehearsing with a view of conveying my ideas regarding the form and shape of the piece. The score itself can be seen in figure **??**. When comparing the general shape of the score one can observe the connection between the implied gestures on the page and the performed gestures on the audio recordings of *Conatus*.

It was important to me that this piece could be more or less recreated with the same structural shape and musical milestones each time it was performed. Having said this, I also wanted to facilitate the possibility of spontaneity within each performance. This was largely down to the fact that I had played with the percussionist for many years and we had developed a silent rapport when playing together. I have found the quality of this connection to be unique in my musical career and it would be remiss of me to ignore such a powerful musical connection.

Upon reflection, the most striking thing about the composition of this piece and its form, is how organically the overall scaffolding of the structure came into being. Over the course of several months of free improvisation with the technology and source material, certain motifs began to emerge. These motifs became structural milestones that we would seek to reach each time we performed the piece. One example of this was a musical gesture I have already mentioned which we referred to as *Gok*. This was created when Ror would simultaneously strike a cowbell and the kick drum, thus creating a strong attack-decay gesture. When this occurred I would engage the dynamic reverb instrument and raise my right hand to bring the feedback of the reverb to a level just below self oscillation. This musical gesture soon became a transitional section within the piece at which we would aim to arrive at a specified time during each performance. Knowing that the piece would end up in this state shaped the manner in which we constructed preceding musical phrases.

The technology used within *Conatus* had as much an effect on the structure of the piece as it did on the audio output of the compositional system. The design of software instruments and how they were activated had a direct impact on the performance of this piece. For example, the final section of the piece, which we dubbed the *Autechre* section, was greatly influenced by the fact that the *Dynamic Reverb* and the PVfreeze instrument were triggered using the same finger on the glove. This, in sonic terms, meant that moving from gated reverbs to the synth-like sound of the PVFreeze instrument was a trivial matter, all that was required was a simple rotation of the wrist. The somewhat arbitrary decision I had made, placing the trigger for both of these processes on the same finger, resulted in the definition of quite a significant portion of the third movement.

Score

The score for *Conatus* does not have the same descriptive symbols found in the score for Proprioception. Instead the symbols used are to be seen as largely interpretative; their main function is to provide a visual analogy to the overall shape and evolution of the piece. The score can be viewed as a long continuous image as was originally intended on the accompanying flash drive, but for ease of representation I have displayed it in this thesis in two halves, one on top of the other. The score should be read from left to right beginning with the top-most system.

As I have already stated, this score is intended to be interpretative for the most part but there are some explicit instructions present. There are two bars above the image of the score and one below, I will refer to these as bar one, bar two and bar three respectively.

Bar one contains a rough indication as to what section should be featured at a particular point. The three letters E, R and A refer to the names Eno, Reich and Autechre, which were used as shorthand descriptions by myself and Ror to describe the musical character of a particular section. Bar two provides information to the percussionist as to what exact instrument should be introduced or featured and bar three provides the equivalent information for the performance of the live electronics.

Some traditional musical terms have been used to describe the manner in which the percussionist should play the cymbals and then subsequently the moment in which they should switch to using their sticks and mallets. Traditional symbols and lettering have also been used to indicate the general shape of the dynamics throughout the piece.

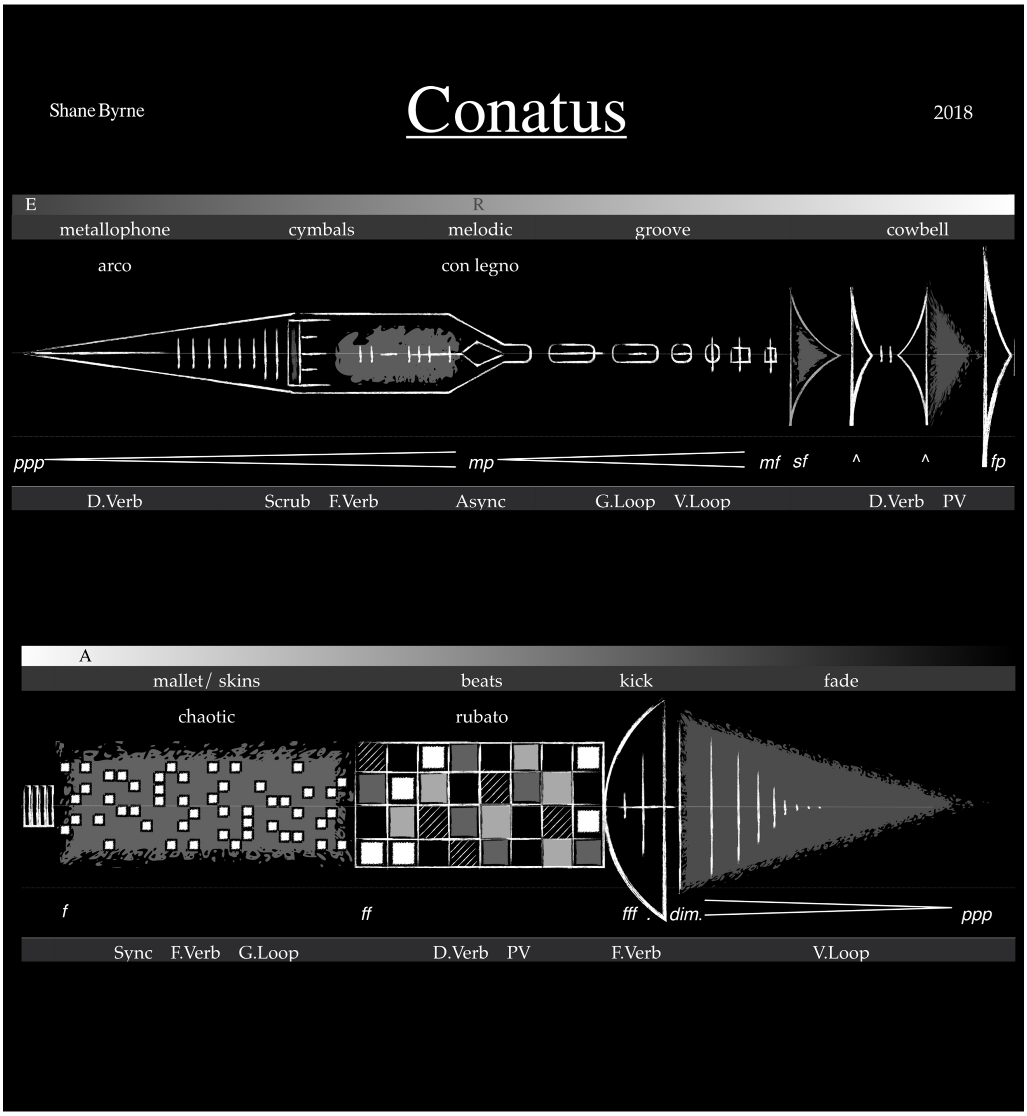

Figure 4.21: Conatus Score

4.10.4 Conclusion

Conatus was chosen to be the subject of a case study by a philosophy PhD researcher in UCD as part of his thesis. As a result of this, almost every interaction and rehearsal that we had was documented either textually, audibly or visually. This meant that the whole compositional process could be followed from start to finish and the mechanisms and strategies employed during its inception could be explicitly observed. Not only did this open up the possibilities for future analysis, but it also provided a constantly growing body of work from which we could further develop transient events and concepts that otherwise may been lost to memory.

This particular process proved extremely fruitful and approaching compositions in an almost ethnographic manner will now become part of my compositional my practise. Having the option to examine, in fine detail, the serendipitous musical gestures which may have occurred during a four hour rehearsal was instrumental in the final manifestation of this piece. The totality of the work done during the inception, development and performance of *Conatus* is another example of the process becoming integral to the art.

While *Tenterhooks* may represent the pinnacle of my methodological approach within the context of fixed media composition, *Conatus* represents the zenith of my live performance practise to date. It effectively addresses issues encountered in previous live performance pieces and further refines the mechanisms I employ in an effort marry musical gesture and physical gesture in the live performance of electroacoustic music.

5

Conclusion

The aim of this portfolio was to explore new ways in which HCI and mediation strategies can be used to expand the compositional and performative lexicon for makers of experimental electronic music. The theoretical frameworks introduced in the second chapter function as a scaffolding for the concepts and observations which are discussed in subsequent chapters. The work of authors such as Marc Leman and Rolph Inge Godøy help to situate this body of work in the context of current areas of interest in the field of electroacoustic composition with a view to advancing the current state of the art. References to the work of Small and Csikszentmihaly seek to situate the art of electroacoustic music in the broader context of the performative arts with a view to enhancing the potential for spectacle in the dissemination of future electroacoustic works.

The work presented in this portfolio represents my journey as a composer of electroacoustic music. During this time I have learned that to experiment and take risks is important if one wishes to avoid well travelled compositional paths, and uninclusive musical performances. Music is alive and it is ritual. The capacity for it to thrive as such, is down to the manner in which the composer interacts with their systems of musical composition. Although the medium itself is not exclusively the message, it certainly plays a defining role in its manifestation.

Much of time preparing the pieces in this portfolio was spent writing code, wrestling with gestural recognition and soldering circuits, which is not how I would have imagined my time being distributed when I initially set out to complete this PhD. As I have discussed in chapter 3, there was continuously a tension between what I believed my practise to be, that of the technician or that of artist. Completing this body of work and reflecting on my methodology has brought me to the conclusion that no such tension need exist. The exact methods that I use to create my art requires me to take an autoschediastic approach, to be a *bricoleur*. My artistic output is both multimedia and multimodal which by definition requires some level of bricolage.

A multimodal approach toward the creation of musical experiences has never been more achievable, thanks to advancements in mediation technologies and their affordability. I intend to continue investigating methods of constructing immersive musical experiences, augmenting compositional agents and further refining existing mediation strategies. Pieces such as *Conatus*, *Proprioception* and *Kinesia* have all just scratched the surface of what is possible for future performances of my musical work.

Appendices

Appendix A

Tools and Technology

"Life imitates art. We shape our tools and thereafter they shape us. These extensions of our senses begin to interact with our senses."

– John M. Culkin, *A Schoolman's Guide to Marshall McLuhan*

A.1 Software

Here I will briefly introduce two of the main software programming environments that I use throughout this portfolio. While any in-depth examination of either of these two programming environments goes beyond the scope of this thesis, I feel that it is important to briefly discuss how each of these programs operate in terms of syntax and functionality.

This brief discussion will serve to clarify any syntactical references I make to the software later on in this body of work and it will allow anyone unfamiliar with these programming environments to be able to decipher some of the code examples provided in this thesis.

A.1.1 Csound

Csound is an audio programming environment which can be used for the digital processing of audio material or for the synthesis of sound from first principles. It is descended from the MUSIC N series of digital synthesis programs developed by Max Mathews, beginning in 1957 with MUSIC I (Lazzarini, Yi, et al., 2016). This series of software environments continued to be developed by Mathews and over the numerous iterations of the MUSIC series, various tools were added to the computer musicians toolbox, many of which became staples in most of the modern audio programming software that is found today. The three main components introduced in the MUSIC series that defined the manner in which audio digital signal processing (DSP) is handled in Csound are the unit generator (UG), the instrument and the compiler, which Mathews called the acoustic compiler (ibid.). Csound itself was created in 1985 at MIT by Barry Vercoe(Vercoe and Ellis, 1990).

Unit Generators

The UG serves a building block for DSP and synthesis within Csound. It acts as a front-end for what can often be the complex inner workings of a function within Csound. In the context of Csound, these UGs are often referred to as *opcodes*. I will from now on use the term opcode when discussing UGs in this thesis.

To provide an example of one of these opcodes, I will discuss the implementation of a Schroeder Reverb, which consists of several all pass filters placed in series being fed into an array of parallel comb filters which are then sent into a mixing matrix (Smith, 2017). Rather than having to manually build this processor from first principles, the Csound user can instead invoke the reverb opcode, which does all the required processing but does so in a way that is hidden from the end-user. All the processing essentially takes place under the hood. Below in code excerpt A.1 is an example of the `reverb` unit generator.

```
arev reverb asig, krvt
```

Code Excerpt A.1: Implementing a Schroeder reverb in Csound

In the above example, it can be seen that on either side of the opcode are numerous variables which are denoted by one of the letters a, k, or i preceding the variable name. K and i-rate variables are scalars and are often used to control note parameters which change at either at a predefined control rate or a note initialisation rate respectively. A-rate variables are updated at the audio sampling rate and are used to store and recall the output of oscillators, filters and sound files (Lazzarini, Yi, et al., 2016). In code excerpt A.1, the inputs for the reverb opcode are placed to the right of the opcode and the outputs are placed on left. This is the most common syntax for the implementation of an opcode within Csound.

Instrument

The instrument in Csound is demarcated by `instr` and `endin` flags within the code. The instrument serves to chunk functions together in order to create a complete DSP/synthesis module. The example shown in code excerpt A.2 is an instrument containing a function which generates a sawtooth wave via the `vco2` opcode. The output of this opcode is then routed to a low-pass Butterworth filter. The output of this filter is subsequently sent to the default stereo outputs of the system using the `outs`[1] opcode.

```
instr 1
    kamp = 0.5
    kfreq = 540

    asig oscil kamp, kfreq
    afil butterlp asig, 470

    outs afil, afil
endin
```

Code Excerpt A.2: A basic Csound Instrument

Acoustic Compiler

The complier within Csound is based on the acoustic compiler that was introduced in MUSIC III (ibid.) and allows for efficient execution of synthesis programs and essentially

[1]This one of the aforementioned exceptions of syntax where the outputs are on the right of the opcode

grants the user the potential to create an unlimited amount of instruments for use in the creation or manipulation of audio.

Communication

One the greatest strengths of Csound is that is compatible across various different operating systems and platforms including Linux, OS, Android and HTML (Lazzarini, Costello, et al., 2014). This means that when using multiple devices, which may be running different operating systems, Csound can function in much the same way across these devices without having to edit or restructure the source code. Csound has always striven to retain backwards compatibility (Lazzarini, Yi, et al., 2016) which means that instruments that were created on an older version of Csound will still function in newer versions. This feature is extremely important to me due to the fact that I tend to use technological tools in my compositions. I simply cannot afford to rely on software developers to update the functionality of a specific piece of software.

One feature of Csound is its ability to facilitate multiple forms of communication between devices. This can take the form of MIDI, Open Sound Control (OSC), Bluetooth and serial communication. The potential for this audio programming environment to effectively and efficiently communicate with other software makes it an extremely useful tool for my compositional work, especially due to the fact that I often use an assemblage of devices and software within single projects.

A.1.2 Processing

Processing is a programming environment which has syntactical roots in the Java programming language (Reas and Fry, 2015). It was developed by Casey Reas and Benjamin Fry in 2001 as a prototyping environment and tool for teaching programming using graphics and visual feedback (Reas and Fry, 2004).

The software can export sketches as Java applets making it incredibly easy to share custom built programs with others within the Processing programming community. It is within this community and attitude of openness that the strength of the Processing programming environment really lies. While it is mainly geared toward the creation of vector images and image processing (Reas and Fry, 2004) it also has a large range of network communication possibilities. This, coupled with its large, active community of contributors makes Processing an extremely useful tool when attempting to access data from commercially available technologies such as the Xbox Kinect and the LEAP motion sensor, or when creating projects using microcontrollers such as the Arduino or the Beagle Bone.

Because of the large community of contributors, there is generally a high possibility that if it is at all possible to access raw data from devices and USB peripherals such as the aforementioned Kinect, it is likely that there will be some contributed library which can be used within the Processing environment to access that device. This was certainly true for the Kinect and the LEAP motion sensor, providing users like me, who do not consider themselves to be software developers, with a means for flexibly controlling these devices for their own artistic endeavours. The fact that Processing not only supports serial communication but also network communication in the form of the Open Sound Control (OSC) protocol, means that it was perfectly situated for use in my work where I use OSC extensively to send and receive messages to and from the Csound audio programming environment.

A.1.3 OSC

Open Sound Control (OSC) was developed by Matthew Wright and Adrian Freed to enable networked communication between computers, instruments and synthesisers(Wright, 2005). OSC is a binary message format which can be carried by any network technology (ibid.) meaning that it is incredibly useful when the user wishes to share performance information between several different musical devices or between applications running locally on a single device. With the aid of OSC, wireless communication between wearable or embedded devices and software sound synthesisers becomes a trivial task, opening up the potential for nearly any networked interface to be used in the context of musical performance.

A.2 Kinect

The Microsoft Kinect sensor is a computer vision input device which uses a combination of a depth sensor and an RGB video graphic array (VGA) camera to detect motion and recognise human skeletal structures. The original function of this device was to enable users to control aspects of Xbox 360 video games through the movement of their body. In this capacity, the Kinect was billed as an alternative method for video game interaction with a view to immersing the user in an active experience. Quoting from the original Microsoft press release (when the Kinect was still known as Project Natal), "Project Natal provides gameplay that gets you off the couch, on your feet and in the fun. Each "Project Natal" experience is designed to get players moving, laughing, cheering and playing together"(Microsoft, 2009). It is also described as a Natural User Interface (NUI) which is an interface that is essentially invisible to the user (Mann, 2007). In 2013 Microsoft released the Kinect 2.0, but for all projects in this portfolio I have used the original 2010 release.

The Kinect V1 has a depth sensor range of between 800mm and 4000mm (Microsoft, 2012a). The RGB camera has an image resolution of 640 x 480 pixels and a field of view of 62 x 48.6 degrees. The depth mage has a resolution of 320 x 240 pixels and a field of view of 58.5 x 46.6 degrees (Smeenk, 2014). The device itself measures approximately 280mm x 60mm x 75mm and is powered with its own power supply and connected to the host device via USB.

In late 2010 open source drivers were developed as part of a competition run by open-source hardware company Adafruit. This in turn led Primesense, the developers of the depth sensing technology used in the Kinect, to create their own open-source drivers and eventually the Open Natural Interaction Organisation (OpenNI) developed their own open-source software that could be used to read sensor data from the Kinect. It is this OpenNI framework that I use inside the Processing programming environment to extrapolate gestural data for use in my compositions. As a result of these developments, the Kinect can now be used on multiple operating systems as a generic NUI and has become a tool much favoured by artists working in the field of audio/visual media.

A.2.1 3D vision

The device itself consists of an RGB camera and an infrared depth sensor. The depth sensor is made up of an infrared projector and camera, developed by 3D sensing company Primesense. The depth sensor works by casting a speckled pattern of light into its range of vision. The sensor detects these speckles and cross correlates them with the speckled pattern cast by the IR projector. The 3D image is then reconstructed using triangulation techniques to achieve a convincing representation of the objects in the field of vision of the Kinect (Zhengyou, 2012).

Each pixel of the depth image is analysed by the Kinect software to recognise and recreate human skeletal structures. The Kinect can track up to 20 different points on a human skeleton (see Figure A.1) at a rate of 30 frames per second. It is this technology in particular that I found to be the most promising when envisaging ways in which the Kinect may be used to control musical parameters. The fact that the sensor is designed to detect multiple human skeletal structures from an array of different body shapes and sizes without minimal calibration necessary on the part of the end-user is particularly impressive.

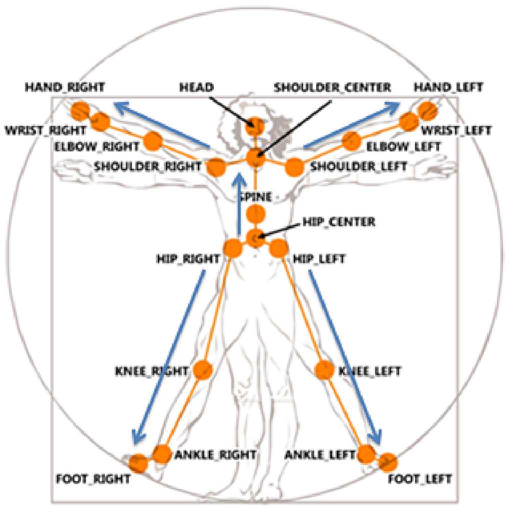

Figure A.1: Points of the body tracked by the Kinect sensor(Microsoft, 2012b)

One issue that I have encountered when using the depth sensor is that while the IR sensor is not susceptible to interference from domestic light sources, it does suffer from poor performance when under certain stage lighting and from any UV light sources such as sunlight. The speckled pattern generated via the IR projector tends to get washed out by these sources. This unfortunately means that the Kinect is not a viable option for motion sensing outdoors or any environment where there is a lot of natural light.

A.2.2 OpenNI and Processing

While there are several software environments that were developed to allow the user to extrapolate data relating to gesture and movement, including the applications such as Kinectare (Vik, 2014), The Wekinator (Fiebrink, Trueman, and Cook, 2009) and EasyGR (Ibañez et al., 2014), I elected to use the OpenNI framework within the Processing software environment to extrapolate the raw skeletal data from the Kinect. The code that I use is based on the OpenNI user test developed by Max Rheiner. This code collects the positions of 15 points of the user skeleton, interpolates these points, and draws the user within a 640 x 480 graphical window. When the user is detected their silhouette changes from a grey depth image to a blue one, providing the user with useful visual feedback.

I adapted the original code to calculate the velocity of the left and right hands respectively in addition to the preexisting point tracking methods. Another modification I made to the original code was to include OSC capabilities via the OSCP5 library for Processing. This enabled me to share the normalised, interpolated skeletal data with Csound via an internal network connection. Once the values within this data set were received by Csound, they could be used as arbitrary control values for any number of musical parameters.

A.3 Intel Galileo

The Galileo board was Intel's response to the growing interest in microcontrollers, open source hardware and embedded computing that began with the release of the Arduino in 2005 and was maintained through the release of subsequent single board computers such as the Beagle Board, MSP430 and the Particle Photon to name but a few. Rather than being in direct competition with Arduino, the Galileo is essentially an Intel Quark

System-on-Chip (SoC) (Ramon, 2014) paired with an Arduino Uno, all of which are mounted on a single board.

The Arduino functionality allows users to connect input devices such as switches, potentiometers and sliders, as well as output devices such as LEDs and speakers to the General Purpose Input and Output (GPIO) pins so that they may be used in software applications. The power of microcontrollers lies in the fact that they, "act as gateways between the physical world and the computing world" (O'Sullivan and Igoe, 2004). The Galileo was released under two versions, Gen1 and Gen2, both of which used slightly different processing architectures.

A.3.1 Applicability

The advantage that the Intel Galileo has over the Arduino and many of the other single board computers that I have mentioned, is that it is capable of running a Linux based operating system. As a result of this, the scope of interaction that the user has with the software is no longer confined to the Arduino integrated development environment. Professor Victor Lazzarini of Maynooth University developed a version of Csound that can run on the Yocto Linux distribution designed specifically for the Galileo, meaning that the task of physically interacting with the DSP and synthesis software of Csound became trivial (Lazzarini, Timoney, and Byrne, 2015).

A.3.2 GCsound

This specialised version of Csound became known as GCsound (ibid.). Using this specially designed distribution of Csound, the user can access the GPIO pins within the Csound programming environment by using a couple of lines of code.

Code excerpt A.3 demonstrates how a user might access an input device connected to analogue pin zero of the GPIO header. The opcode `chnget` is used to extrapolate data from a specified software bus, in this case that bus is the first analogue input on the board, A0. There are seven analogue input pins in total ranging from A0-A6.

```
ksig chnget "analog0"
```

Code Excerpt A.3: Accessing the first analogue input pin on the Galileo from Csound

Accessing digital input and output functionality is slightly more convoluted and requires the switching of several multiplex controls. Each of the 14 digital pins can be used as either inputs or outputs but this must specified in the Csound script using two Galileo opcodes: `gpin` and `gpout`. If the user wishes to access a GPIO pin as an input, the `gpin` opcode should be evoked and when the pin is required to be cast as an output then the `gpout` opcode should be used. Both Gen1 and Gen2 have different mapping schemes. Figure A.2 demonstrates the mapping scheme for the Gen1 board.

pin	mux selector, value	source/function
0	40, 0 40, 1	UART0 RXD (/dev/ttyS0) 50 (GPIO)
1	41, 0 41, 1	UART0 TXD (/dev/ttyS0) 51 (GPIO)
2	31, 0 31, 1	14 (Quark GPIO) 32 (GPIO)
3	30, 0 30, 1	15 (Quark GPIO) 18 (GPIO)
4	-	28 (GPIO)
5	-	17 (GPIO)
6	-	24 (GPIO)
7	-	27 (GPIO)
8	-	26 (GPIO)
9	-	19 (GPIO)
10	42, 0 42, 1	SPI1_CS (Quark) 16 (GPIO)
11	43, 0 43, 1	SPI1_MOSI (Quark) 25 (GPIO)
12	54, 0 54, 1	SPI_MISO (Quark) 38 (GPIO)
13	55, 0 55, 1	SPI_SCK (Quark) 39 (GPIO)

Figure A.2: Multiplexing on Gen1 board (Lazzarini, Timoney, and Byrne, 2015)

The above table shows that pins 4-9 are directly linked to their GPIO functionality without any multiplexing required. If however, the user wished to access the GPIO functionality of pin number 1, they would first have to declare pin 1 as a GPIO pin by changing the value of corresponding multiplex selector (which would be 41 in this example) to a value of 0 . The user can then access the GPIO functionality of pin 1 by using its corresponding GPIO value of 51. Code excerpt A.4 is an example of accessing the GPIO pins within the Csound environment. The instrument uses multiple analogue and digital pins to read a sound file, manipulate its pitch, reverberation time and delay feedback level whilst providing visual confirmation that the sound file is playing via an LED connected to GPIO pin 5.

```
instr Sfile

    kin gpin 28      ;read switch input from digital pin 4
    gpout kin, 17    ;switch status controls LED

    idel = 0.5       ;initialise the delay

    ;Assign parameters to input pins
    kamp    chnget  "analog0"
    krsize  chnget  "analog1"
    kfb     chnget  "analog2"
    kpitch  chnget  "analog3"
    kpitch += 0.5          ;offset pitch

    if kin == 1 then
        ;read soundfile
        asig diskin2 "Sfile.wav", kpitch, 0, 1
        asig *= kamp

        ;delay line
        adel delayr idel
            delayw (asig*kfb)+(adel*kfb)

        ;reverberation
        a1, a2 freeverb asig, asig, krsixe, 0.7
            outs a1+asig+adel, a2+asig+adel
    endif

endin
```

Code Excerpt A.4: Complete Delay/Reverb Instrument for Gen2 board

When using the Gen2 board the multiplexing scheme is slightly different (see Figure A.3). For example, if the user wanted to use digital pin number 1 as an output to control an LED, they would first have to access the multiplex selector that corresponds to pin number 1 (in this case 45) and set that value to 0 using the `gpout` opcode (see code excerpt A.5). This declares that the pin should be treated as a GPIO pin with access the Quark SoC. The user would then have to configure the pin to behave as an output by setting the multiplex value 28 to 0.

pin	mux 1, value	mux 2, value	dir	22k res	source/function
0	-	-	-	-	UART0 RXD (/dev/ttyS0)
	-	-	32	33	11 (Quark GPIO)
1	45, 1	-	-	-	UART0 TXD (/dev/ttyS0)
	45, 0	-	28	29	12 (Quark GPIO)
2	77, 1	-	-	-	UART1 RXD (/dev/ttyS1)
	77, 0	-	34	35	13 (Quark GPIO)
	77, 0	-	-	35	61 (PCAL9535A GPIO)
3	76, 1	-	-	-	UART1 TXD (/dev/ttyS1)
	76, 0	64, 0	16	17	14 (Quark GPIO)
	76, 0	64, 0	-	17	62 (PCAL9535A GPIO)
4	-	-	36	37	6 (Quark GPIO)
5	66, 0	-	18	19	0 (Quark GPIO)
6	68, 0	-	20	21	1 (Quark GPIO)
7	-	-	-	39	38 (PCAL9535A GPIO)
8	-	-	-	41	40 (PCAL9535A GPIO)
9	70 0	-	22	23	4 (Quark GPIO)
10	70, 0	-	26	27	10 (Quark GPIO)
11	44, 1	72, 0	-	-	SPI_MOSI (spidev1.0)
	44, 0	72, 0	24	25	5 (Quark GPIO)
12	-	-	-	-	SPI_MISO (spidev1.0)
	-	-	42	43	15 (Quark GPIO)
13	46, 1	-	-	-	SPI_SCK (spidev1.0)
	46, 0	-	30	31	5 (Quark GPIO)

Figure A.3: Multiplexing on Gen2 board (Lazzarini, Timoney, and Byrne, 2015)

There are also a pullup/pulldown resistor settings which enable the user to set the default state of the pin to be either high or low. This is controlled by setting multiplex selector 29 to a value of 0 for pulldown or a value of 1 for pullup. Code excerpt A.5 is an example which uses pin 1 as an output pin. In this Csound instrument built for use with the Gen2 board, a metronome is used to blink an LED once every second.

```
instr LED

    gpout 45, 0        ;access Quark GPIO
    gpout 28, 0        ;Set pin to output
    gpout 29, 0        ;Use pulldown resistor

    kcount metro 1   ;send a value every second

    gpout kcount, 12;set state of LED to metronome value

endin
```

Code Excerpt A.5: Blinking an LED on the Gen2 board

A.3.3 Communication

Network communication with the Galileo is possible through the use of the on-board Ethernet connection. The Galileo does not come with wireless functionality, for this a separate peripheral is needed. Establishing a connection between a host device and the Galileo can be achieved by using the ssh protocol[2] via a DHCP[3] server. This allows for fast transfer of files between devices and is a convenient method for prototyping new instruments and interfaces.

A.4 Intel Edison

The Intel Edison was the second generation of single board computers that Intel developed after the Intel Galileo. The Edison is specifically geared toward embedded computing and the internet of things (IOT) functionality. The emphasis on the IOT potential of this board meant that unlike the Galileo, the Edison came with in-built WiFi capabilities, immediately offering a connectivity advantage over its predecessor, itself only capable of wireless communication with the addition of a WiFi module. The footprint of the board is significantly smaller than that of the Galileo, measuring just 34mm x 24mm x 4mm, which indicates that Intel were attempting to increase the potential for use in embedded and wearable applications.

Although the Edison is smaller than the Galileo, it is far more powerful. Its CPU is made up of a dual processor, which consists of an Intel Atom x86 CPU running at 500MHz and a 32-bit Intel Quark processor running at 100MHz. The Galileo in comparison consists of a single core Intel Quark processor running at 400MHz. The RAM of the Edison has been increased by four times over the Galileo (Intel, 2013, 2014). This board is not only smaller and more compact in its in design than many other similar embedded computing devices, but the Edison also performed almost five times better than the Raspberry Pi and was over twice as fast as the Beaglebone Black in CPU benchmark tests (Hunt, 2015).

[2]This is a network protocol used to securely use network services over an unsecured network

[3]Dynamic Host Configuration Protocol (DHCP) is a means for allocating Internet Protocol (IP) addresses and other related information to a host within a client/server architecture

While this indicates that the Edison is indeed a large improvement over the Galileo in terms of wireless connectivity and processing power, one major disadvantage of this microcontroller is its GPIO connectivity. Due the fact that the Edison is so small, there are no GPIO pins on the standard board. If the user wants to access the GPIO functionality of the Edison, one option would be to buy an Arduino board which attaches to the Edison. This grants the user similar GPIO access as the Galileo, although this board has quite a large footprint (it is actually bigger than the Galileo). An alternative to that would be to use the smaller Edison Breakout board which has less GPIO access than the Arduino board but is significantly smaller. When debating which path to take, I came across a number of modules (or blocks) built by American based company SparkFun. One of these blocks was a GPIO module which provided equal amount of access to the GPIO pins as the Edison Breakout board but was truly modular, meaning that multiple blocks could be stacked on top of each other, extending the sensing possibilities of the Edison and providing easy access to battery power for mobile applications.

A.4.1 Peripherals

As I have already mentioned, there are a number of modules that can be used with the Edison in order to grant the user extended functionality for use in embedded and wearable applications. From the range of peripherals available from SparkFun, I used three in particular when creating some of the work in this portfolio. The final overall footprint of the Edison when all three of these peripherals are connected is still small enough for wearable or embedded applications, measuring just 70mm x 30mm x 18mm.

SparkFun GPIO board

This board is used to grant the user access to the GPIO functionality of the Edison board. This includes twelve digital pins, four digital pins with pulse width modulation capabilities, a VSYS power point, a 3.3V logic power point, a 1.8V logic power point and a ground connection.

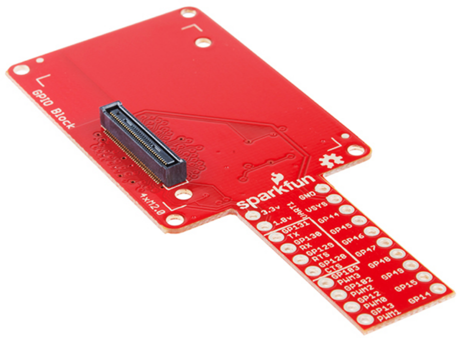

Figure A.4: Sparkfun Edison GPIO Block (Sparkfun, 2015)

SparkFun 9DOF IMU

This peripheral is an Inertial Measurement Unit (IMU) and is used to detect motion on 3D space. The 9DOF part of the title indicates that this particular board grants the user nine degrees of freedom (9DOF). These nine degrees are a triple axis accelerometer, a triple axis gyroscope and a triple axis magnetometer.

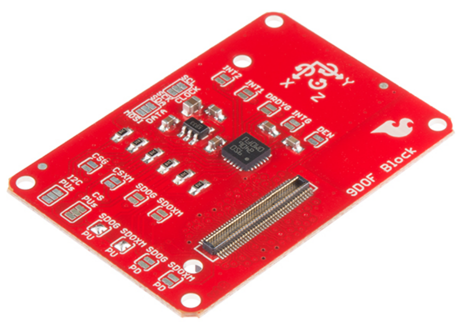

Figure A.5: Sparkfun 9DOF IMU Block (Sparkfun, 2014a)

Sparkfun Battery Block

This battery pack allows the user to power the Edison when access to mains power is not possible. It consists of a 400mAh lithium ion battery operating at 3.7 volts. To give an example of power consumption, when running a simple Python script on the Edison which reads accelerometer data and sends that data via OSC, the battery will provide power to the board for approximately 2.5 hours.

Figure A.6: Sparkfun 9DOF IMU Block (Sparkfun, 2014b)

A.4.2 Implementation

Although there was potential for developing an Edison specific version of Csound much the same way that one was developed for the Galileo, this was not the manner in which I used the board. In the case of the Edison I simply used it as a tool for collecting data rather than for audio synthesis or processing. This meant that the software running on the Edison was lightweight and data transfer was quick and free from congestion. For all pieces in this portfolio in which I used the Edison, its role was just to run simple Python scripts which collect data relevant to the movement of the physical components within that particular work and subsequently transfer that data to a host device via OSC.

A.5 LEAP Motion Sensor

The LEAP motion sensor is a small NUI that is designed to work specifically with hand movements and gestures. It comes in the form of a small USB peripheral measuring 80mm x 30mm x 12mm. It is intended to be used either by placing the sensor facing upright on a work surface or attaching it to an virtual reality headset. While it is primarily designed to recognise hands, it can also be adapted to be used with stick-like tools such as pencils and pointers.

Inside the small enclosure there are two cameras and three IR LEDs which operate outside the spectrum of human vision. The two cameras are fitted with wide angled lenses which provide the sensor with a field of vision extending to 60cm above the sensor (Colgan, 2014). The stereo vision capabilities of the sensor enables the accompanying software to triangulate the location of a hand and track it in 3D space. The software is calibrated to recognise not only the Cartesian coordinates of both the left and right hand, but also how many fingers are being held out as well as the roll, pitch and yaw of both hands. The software also has the ability to recognise several gestures including pointing, swiping and circular motion. The accuracy of this recognition is quite limited however and the software will often confuse one gesture for another.

A.5.1 LEAPintoCsound

Initially I began using the LEAP motion sensor in work that I had been doing with a colleague of mine, Simon Kenny. Simon had developed a library that allowed the LEAP motion to operate directly within the Csound programming environment(Kenny, 2016). This functionality was manifest in the form of several new opcodes that he wrote for Csound: `leaphand`, `leapgesture` and `leapsettings`. The first of these opcodes, `leaphand`, is used to read data from the LEAP motion sensor. It uses the following syntax:

```
kout1[,kout2,kout3]    leaphand    khand, kinfo
```

Code Excerpt A.6: The leaphand opcode

The variable `khand` is used to determine which particular hand the opcode should track. A value of 0 means that the newest hand will be tracked, while a value of 1 will mean that the oldest hand will be tracked. A value of 2 will track the left hand only and 3 will track the right hand only. Using a value of 4 will allow the user to track both hands simultaneously. The manner in which the second variable `kinfo` operates is dependant on the state of the `khand` variable. The number of outputs generated by the opcode is also dependant on `kinfo`. For example, if `kinfo` is set to 0, it will generate three outputs, x,y and z, which provide information pertaining to the cartesian coordinates of the tracked hand. If however, the value for `kinfo` is set to 1 it will return the number of fingers extended on the tracked hand. This will only produce one output. Table A.1 and table A.2 provide the mapping and functionality of single hand tracking and double hand tracking respectively.

Value for `kinfo`	No. of Outputs	Function
0	1	position (x,y,z)
1	3	number of fingers
2	1	velocity(magnitude of vector)
0	1	velocity(vector)
1	3	orientation(roll,pitch,yaw)
2	1	time alive(seconds)
0	1	activity flag

Table A.1: Functionality of opcode when `khand <= 3`

Value for `kinfo`	No. of Outputs	Function
0	1	distance between hands
1	3	distance between hands (x,y,z)
2	1	Angle of rotation on x,y plane

Table A.2: Functionality of opcode when `khand = 4`

Although basic gesture recognition was implemented in this library via the `leapgesture` opcode, the opcode itself proved to be quite unreliable and had a tendency to confuse gestures resulting in out-of-bounds values, sometimes causing the Csound front-end to crash. The reason for the lack of fixes for these bugs was the fact the opcode was never officially packaged with Csound and was being maintained solely by the original developer Simon, who had since moved on to other work[4]. When Csound version 6 was released, the library was not updated and ceased to function altogether. Thankfully the library has since been updated with most of its functionality restored.

The third opcode in the library, `leapsettings` is used for the modification of several internal settings for the calibration of the sensors software and hardware.

[4]Due to the fact that it was never officially packaged with Csound, it is very likely that I am one of the only people actually using this library

A.5.2 LEAP Motion for Processing

During the period when I was waiting for the library to be updated, I investigated other means for accessing the raw data from the LEAP motion sensor. I came upon a library written for Processing which collected similar data to the LEAPintoCsound library. Some notable exceptions include more robust gesture recognition and the inclusion of grab strength and pinch strength variables. I used this library when generating source material for some of the later fixed media pieces in this portfolio.

Appendix B

Interactive Pieces

> *"I want machines rather like computers to be an extension of the arm of the composer"*
>
> – Daphne Oram

B.1 EAREYEMOUTH

EAREYEMOUTH was an interactive installation which came about as the result of a collaboration with digital artist Simon Kenny. As I have mentioned earlier in this thesis, Simon is the author of the LEAPintoCsound library which facilitates the use of the LEAP motion sensor within the Csound programming environment. Having worked with this library for a number of months, refining several audio processes and testing the limitations of the LEAP, we both decided to put the library to use in the form of an interactive installation.

B.1.1 Inspiration

We were both interested in the potential of multimodal engagement and decided to base the aesthetic of the installation around three sensory organs, the mouth, the eyes and the ears. There were four mouths in total (Figure B.1c), each acting as an enclosure for a loudspeaker. These mouths were distributed around the installation space to facilitate quadraphonic output from Csound. A single eye (Figure B.1b), made from a repurposed glass bowl, and a single ear (Figure B.1a), cut from high-density foam using a CNC laser cutting machine, were placed in the center of the installation space. A small lapel microphone was placed inside the canal of the ear. This was used to capture audio information in the surrounding environment. The iris of the eye housed a LEAP motion sensor which was in turn connected to a laptop running Csound. The audio captured with the microphone in the ear was sent into Csound where it was then processed using several phase vocoder (PV) based opcodes. The parameters of these PV opcodes were linked to gestural data received by the LEAP motion sensor. The idea behind this initial version of the installation was that the participant could use the ambient sound of the space around them as source material for a live electronic improvisation. The LEAP was used to track the hand movements of participants, allowing them to create intricate and dramatic sonic transformations in real time through physical gestures.

B.1.2 Galway Autism Project

After exhibiting this work at the Irish Sound, Science and Technology Association Conference in 2014 and at Galway Culture night of the same year, we were approached by the Galway Autism Project (GAP) to bring the installation to their premises. They were interested in using non-tactile interfaces to create music and felt that some of their members would benefit from engaging with technology such as the LEAP motion sensor.

We traveled to Galway to begin working with the members of GAP, the majority of whom were 11 to 16 years old. We soon became aware that the sonic output from our system was not suitable for the members of GAP and were told that the sounds we used were "like clawing on the inside of my ears" by one of its members. This was the first time I truly understood the value of inclusive design and how apt the slogan, "nothing about us without us", is when designing interfaces for individuals with special needs.

Simon and I asked the group what particular sounds they enjoyed and the majority of members revealed that soft piano-style sounds and lush, low frequency pads were the sounds most favored. With this in mind, we set about redesigning the sonic output of the installation. The audio input received by the microphone in the ear was used to create an electric piano-style sound, the pitch and amplitude of which was determined using the `pvspitch` opcode within Csound. This opcode analyses a stream of frequency domain data and extrapolates the most significant partials within this stream to determine the pitch of a signal (O'Cinneide, 2005). The LEAP motion sensor was then used to trigger a low frequency pad sound whenever it detected the hand of a participant. The pitch of this instrument was also determined by audio input captured by the microphone. Low-pass and high-pass Butterworth filters were applied to the pad instrument, and the parameters of these filters were controlled by the velocity and position of the participant's hand.

This new design was warmly received by the members of GAP, who appeared to have a lot of fun playing on an "invisible" instrument. Following some subsequent revisions to address any bugs within the software, we submitted the installation for inclusion at the STEM Creative Tech Fest 2014 which was held in Google's Dublin headquarters. This allowed us to bring the installation to a whole new audience of children with special needs. The reception that the installation received was resoundingly positive. For this submission we were awarded the STEM prize for creative technology.

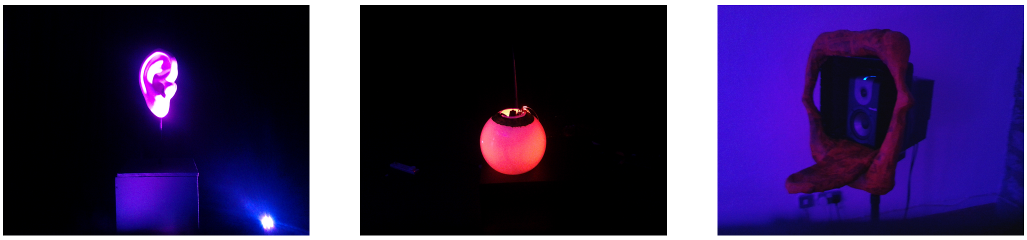

(a) The ear which houses a lapel microphone

(b) The eye with LEAP motion sensor

(c) A loudspeaker inside one of the four mouths

Figure B.1: Images from the EAREYEMOUTH installation

B.1.3 Conclusion

This project opened my eyes to the potential for music technology to be used in the facilitation of creative practices for those who are often excluded due to physical or intellectual disabilities. It became the catalyst for my subsequent work with groups such as the Drake Music Project in Northern Ireland and helped me realise the potential impact that my research could have in the world of inclusive creativity.

B.2 Sound;Waves

This interactive audio installation was created as part of a project called *Metamorphosis: Art as Research, Research as Art*, the goal of which was to highlight academic research in the public sphere through art installations inspired by academic research. For this project I collaborated with an oceanic engineer, Dr Paula Garcia Rosa, whose research interests include automation and control of industrial processes, renewable energy technologies and modelling, and the optimisation and control of ocean energy systems. In particular, Dr Rosa's focus was on one specific design of energy system known as an oscillating-body wave energy convertor (WEC) as shown in Figure B.2.

Figure B.2: An oscillating body WEC

The kinetic energy generated by the motion of this device on the surface of the water is converted into electricity via a number of dynamic processes (Garcia-Rosa, Bacelli, and Ringwood, 2015). Having collected sufficient information regarding Dr. Rosa's research interests, I set about trying to conceive of a way in which I could use the principles of her research to inspire and inform an interactive audio installation. I will now go on to discuss in detail, the final design I arrived at for this installation, how this design was implemented, and some observations I made whilst viewing people interacting with the work at the Metamorphosis exhibition.

B.2.1 Wave Creation

The first concept that I wanted to explore with this project, was the idea that the installation should be driven in some way, even if only symbolically, by wave energy. This meant that I would have to find a means of housing a significant body of water inside an exhibition space. For a time, I considered the idea of constructing a tank using acrylic sheeting but after investigating this option further I discovered that this would not only be an expensive solution, but also the pressure that the acrylic and the bonding adhesive would be under from any large body of water would be too great. The next alternative was to use a medium sized fish tank (measuring approximately 1000mm x 400mm x 600mm) and adapt it to suit the purposes of my installation.

My initial designs had incorporated the use of a motor or pump that could be controlled by the user in order to create motion in the water, but I felt that not only would the noise of the motor interfere with the sound being emitted from the installation (the total audible output of the installation was always paramount in design considerations for this piece), but it also reduced the immediacy and quality of interaction that a participant would have with the installation. The spectrum of potential kinetic input from the participant would be limited to turning a potentiometer or pushing a button which does not allow for many varying ways to physically interact with the piece.

Instead of employing the use of a motor or pump, I decided to use a more simple mechanical solution in order to create motion in the water. I created a paddle which was fixed to the bottom of the tank by an aluminium hinge. This not only served to create an immediate connection with the user's physical input into the system and audio output from the system, but it also facilitated the conversation between the user and the system in a natural and fluid manner which is essential in terms of designing interactive devices (O'Sullivan and Igoe, 2004).

The paddle was cut from a sheet of 6mm acrylic and then framed with strips of aluminium. The strips of aluminium where used to attach the paddle to the base of the tank and to add strength to the paddle; I anticipated that constant resistance from the water could possibly cause the sheet to break. A handle was then attached to the paddle which stuck up out of the top of the tank to allow ease of use for the participant (Figure B.3).

The hinge itself was attached to the tank using a strong waterproof adhesive. The result was that when the handle was pushed forward the paddle moved through the water, creating a surface wave. In the design of installations such as this, I feel that it is more interesting to allow the design of the installation to speak for itself and also to allow the participants to explore the piece for themselves. Therefore, I felt that it was important to try and make it immediately apparent how one should interact with the installation and aimed to make it explicitly clear that the participants should interact with the handle (Figure B.3).

Figure B.3: The finished design of the Sound;Waves water tank, complete with paddle and buoy

The installation at this point involved physical interaction and mechanical processes. I then decided to transduce this mechanical information into electronic signals in order to work within the digital domain. The main reason for this was due to the potential power and malleability of digital instruments in comparison to solely mechanically driven instruments (Rowe, 1993).

B.2.2 Detecting Motion

Once the participant had the ability to create waves in the tank, the next task was to come up with some way of digitally detecting motion in the water. In order to do this I employed the use of the Intel Edison (see Apendix 3.4). In conjunction with

Figure B.4: The Intel Edison connected to the 9DOFIMU

the Edison, I used an inertial measurement unit with nine degrees of freedom (9DOF IMU), which is a triple axis accelerometer, a triple axis gyroscope and a triple axis magnetometer packaged together in one block that can connect directly to the Intel Edison (see Appendix 3.4.1). Although an accelerometer is normally thought of as being used to measure the acceleration of a moving body, IMUs such as this can also be used to measure the tilt of an object along three axis, returning readings for an objects pitch, roll and yaw.

Both of these devices were then connected to a lithium battery block which meant that the motion sensing device could operate fully without the need for any physical connection to an alternative power source for up to 180 minutes. I then placed these three devices, which were at this point connected together as one unit (Figure B.4), into the housing of a floating pool light.

I had removed the battery compartment within the pool light and further hollowed out the interior of the internal enclosure to allow the Edison to fit inside. This pool light enclosure somewhat resembled the WEC that I was trying to represent and it also provided a watertight housing for the Edison IMU, which was essential for the device to function correctly (Figure B.5).

Figure B.5: The Intel Edison inside the unpainted floating enclosure

Although I had been running Csound directly on the Intel Galileo in previous projects (Lazzarini, Timoney, and Byrne, 2015), no sound synthesis or audio processing was taking place on the Edison in this particular project. I decided that rather than having a version of Csound running on the device, I would instead use the Edison as a means for collecting and transmitting data from the within the buoy to another, more powerful device which would host the Csound software environment. I adapted a python script, which had originally been written for the collection of pressure and humidity data (Barnett, 2015), to read information from the IMU pertaining to the roll, pitch and yaw of the buoy. This script also contained a means for sending the data as OSC messages to a remote host.

The idea behind all of this was to enable me to determine the position of the Edison (and by extension the buoy) in 3D space and then compare that reading to its previous position, thereby indicating how much the device had moved since the last reading. Once the kinetic data had been collected by the Edison and transmitted via OSC, it was read by a Csound script running on a separate device. In this case that device was a Macbook pro housed beneath the water tank. Both devices, the Edison and the Macbook, shared information over a private network, facilitated by a router which was placed underneath the tank in a hidden enclosure.

B.2.3 Csound Instruments

Within the Csound script I created four instruments, each given a specific function within the overall system. I will now discuss in detail these functions and how they were implemented within the musical system.

OSC Receiver

The first instrument dealt with receiving the data from the Edison and the IMU via the OSC protocol. The data was received from a port number that matched that of the Edison's transmission and address destination. This data was then scaled to fall within the ranges of values that would allow for functional manipulation of synthesis parameters within the musical system. For example, many of the parameters that are affected within the software instruments fall within a range of 0-1, so it was necessary to normalise the incoming data in order to make efficient use of the incoming values for the purpose of instrument manipulation.

Triggering Instrument

This software instrument was virtually powered by the movement of the buoy; the greater the difference between the previous position and the current position of the buoy, the greater the volume output of the instrument would be. Therefore, if the buoy was not moving at all there would be no audio output, whereas if the buoy was moving a great amount there would be quite a substantial (in terms of amplitude) audio output.

This particular instrument dealt with the triggering of instances of sound using the information collected from the OSC instrument. The readings of the current pitch, roll and yaw of the IMU device were compared to an array of previous readings for pitch, roll and yaw (which were updated approximately 345 times a second) and the amount of total motion was then determined by calculating the difference between them. This calculation was then used to inform a stream of impulses that acted as a trigger for the audio output; the greater the difference between the current and previous readings the greater the frequency of triggers per second.

There are also several pseudo random processes present within this instrument that allow for some deviation in the total duration of a single instance of sound generation. This meant that each single sonic event had a life span ranging anywhere from ten seconds to thirty seconds.

Generation Instrument

The second instrument was responsible for the actual synthesis of the audio output. Within this instrument, the timbre, amplitude envelope, spatial distribution and frequency information were defined. The synthesis parameters of this instrument were not only informed by the input from the user but also through a number of pseudo-random processes.

The reason I included this indeterminate element into the system was to ensure that no matter how long the user interacted with the installation, the audio output would constantly be in a state of flux. This ensured that the user could never be fully sure of the ensuing output from the system, inviting moments of serendipity and elements of play into the overall experience. It was also important to me as a designer that the musical system itself should have the potential to behave in such manner that the user approximately replicate musical gestures, whilst other certain nuances of the system were given space to exist and evolve organically (Davis, 2010).

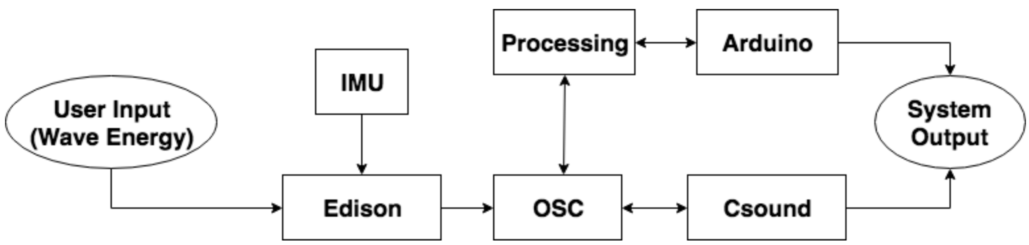

Figure B.6: A flow chart demonstrating the order of communication with the Sound;Waves system

Optimisation and Calibration

The audio instrument design itself consisted of a physical model made up of a bank of modal frequencies that could be played using a choice of three excitation sources. The resulting audio output resembles that of a small metallic bowl. The manner in which the instrument was excited was dictated through an indeterminate process using a pseudo random number generator. Each time an instance of an instrument was triggered, a three sided virtual coin was tossed, which then dictated how the instance of that instrument should be articulated.

In keeping with the idea of creating an analogy for Dr. Rosa's work, I included the possibility for the participant to alter or calibrate the sonic character of the instrument through additional user input in the form of three potentiometers, which were mounted on top of the lid of the tank. In order to be able to read the potentiometers within the musical system, I had to introduce an Arduino into the signal chain (Figure B.6). The communication between the Arduino and Csound was facilitated by the graphical programming environment Processing, which has libraries capable of gathering data from an Arduino and subsequently passing that data to Csound using the OSC protocol. This also provided me with the potential for additional feedback in the system, which something I will discuss later in this chapter.

The addition of user input offered a wider range of expressive potential for each instance of articulation within the musical system and granted the participant more control over the audio output. This combination of user control and indeterminate processes allowed for a more organic compositional system in which there was a degree of synergy between the participant and the musical system, therefore blurring the roles of performer and instrument, creating what Dr Simon Waters might describe as a "performance ecosystem" (2007).

Articulation One

The first instance of articulation of the bank of modal frequencies was that of white noise being used as an excitation source. This produced a sound rich in harmonics with a graduated continuant morphological archetype (Smalley, 1986). In the case of the first instance of excitation, the user was given the ability to alter the centre frequency of a second-order Butterworth band-pass filter. This filter is designed to operate in between a range of 60Hz and 2.6kHz.

Articulation Two

The second instance of an excitation source came in the form of a stream of impulses, the speed of which was dictated through an indeterminate process. Each time the instrument was triggered there were two impulse rates assigned to the instrument. The range of these intervals was between one hundred milliseconds and one second. Once both interval rates were determined, they were then each assigned to a variable, `intrvla` and `intrvlb` respectively.

A Csound function was then used to trace a line between the two interval values meaning that the output of impulses will either speed up or slow down depending on another indeterminate process. For example, if `intrvla` was two hundred milliseconds and `intrvlb` was five hundred milliseconds, and the pseudo-random process had dictated that `intrvla` would be the starting value and `intrvlb` is the final value, this would mean that the number of impulses used to excite the instrument would rise from two impulses per second to five impulses per second over the duration of that particular instance of the instrument. The user also had the ability to calibrate the character of this sound through the use of Q or resonance control. This parameter was assigned to one of the three potentiometers located on the top of the tank. The result of a narrow Q produced a highly resonant, bright, metallic sound whereas a wide Q resulted in a more muted sonic output.

Articulation Three

The third and final means of excitation created a sound with an attack-decay morphological archetype (Smalley, 1986) which resembled the sound produced when striking a large metallic object with a soft beater or mallet. The participant could alter this sound by introducing a low frequency oscillator (LFO). The frequency of this LFO could be altered through the manipulation of the third potentiometer on the lid of the tank.

The indeterminate processes, alongside the calibration of certain aspects of the instruments and the fact the instruments themselves were driven by external forces, resulted in the creation of instruments that could be described as expressive, playful and surprising.

There was scope for repetition of certain musical gestures but there was also room left for serendipitous moments when interacting with the instrument. I feel that this style of reciprocal interaction lends itself to the idea of the installation itself having agency, and perpetuates the concepts of encouraging new results from composed material as proposed by composers such as John Cage (Cage, 1961).

Visual Feedback

Although the main idea behind the installation was rooted in audio composition, I also wanted to create some visual feedback within the installation. Just as the audio engine of the installation was driven by the motion of the waves, I felt that it would improve the experiential worth of the installation to drive this visual feedback with wave motion also. In order to facilitate this, I extended the functionality of the Arduino sketch that I had integrated into the system to read the values of the potentiometers on top of the tank.

I purposely chose to use an Arduino Mega for this project simply because of the large number of analog GPIO pins available on the board. The reason for using analog GPIO pins was that I wanted to be able to gradually fade the intensity and colours of an array of light emitting diodes (LEDs), whilst simultaneously being able to read analog information from the potentiometers. While there are ways to achieve this using other smaller models of Arduino, such as using pulse width modulation with the digital pins or multiplexing to obtain the use of more outputs, I felt that this was the most pragmatic way of achieving my goals. Just as the audio output was causally linked to the motion of the buoy in the water, so to would the output of the LEDs.

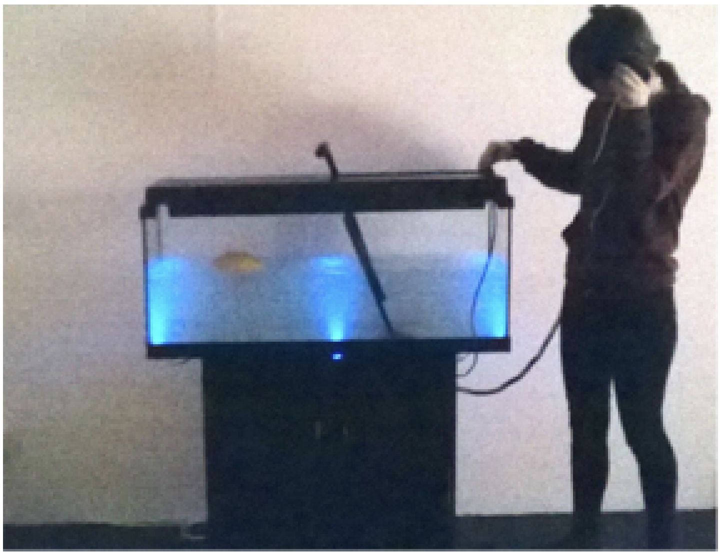

Figure B.7: An example of the Visual Feedback provided by the LEDs

In the Processing sketch I added a few lines of code that sent certain cues to the Arduino when the amount of motion of the buoy in the water exceeded a certain threshold. These cues were then used to control the intensity of the lights. When there was no movement, the intensity of the lights were at their lowest meaning that there was no visual feedback. Inversely, when there was a lot of motion in the water the LEDs shone brightly.

The LEDs that I used were RGB LEDs, meaning that the lights were capable of producing a wide spectrum of colours. Since there were three primary colours in use in this system and I had provided the participant with three potentiometers with which to shape the audio output, I decided that it would make sense to assign a colour to each parameter in the system. The LFO was mapped to the colour blue, the centre frequency was mapped to the colour red and the Q was mapped to the colour green. This provided the participant with further clues to the fact that they were having an effect over the entire system through their calibration of the instrument.

This confirmation of interaction is something that is extremely important when dealing with subtle changes in the sound, especially when the participant may be to a certain extent "untrained" in terms of listening to electronic music. Just because one may not understand fully the effect they are having on the audio output does not mean that they should not be able to appreciate the impact that they having on the installation. This line of thought is the reason that I chose not to label any of the potentiometers on the top of the tank. I felt that it might be alienating to someone that is not schooled in some of the terminology of electronic music to be encountered with three labels with the letters: CF, LFO, Q. I felt that it was much more natural to allow the user to discover what the dials did for themselves without any preconceived notions or abstract cues.

B.2.4 Conclusion

Over the duration of the exhibition in which this installation was featured, I observed that the overall setting of the piece had a large role to play in how the participants contextualised the sounds that they were hearing. When I was initially designing the sound of the instrument I had asked several colleagues for their opinions regarding the character of the sounds used in the system and whether or not these sounds were pleasing or displeasing in any way to them. During this time I received a variance in reactions which ranged from describing the interesting metallic timbres to the crude amplitude envelopes that were present at various manifestations of the instrument. When prototyping this installation I did not receive any comments which related to the oceanic context that the concept of the piece was derived from (at this time the sound was just presented as an abstract entity without any reference to the ocean, a buoy, or the overall project). However, when I presented the sound of this instrument as part of the complete installation, all the contextual references and comparisons that participants expressed to me were in some way associated with the ocean.

The first such observation came from a participant who had described the sound as being reminiscent of "a lonely buoy bobbing on the waves" even though he had found it difficult to situate the sound so clearly when he previously listened to the prototype, when the sonic output had not been associated with the buoy or the tank.

One other description that stood out to me as being quite poignant came from a man who had at one point in his life been a deep-sea diver. He said that after being immersed in the sound for a few minutes, with his eyes closed and headphones securely fixed on his head, he felt as if he was back under the water, diving and exploring the seabed.

He said that there was something about those particular sounds that had triggered memories of diving, and that it was highly representative of what he would experience when underwater. I myself have never been diving and I have done no research into what a diver hears when he or she is submerged beneath the water so I can neither confirm or deny that there may be a relationship between the sounds produced by the instrument and those sounds experienced by a person deep underwater. However, I think that this provides an interesting example of the potential for multimodality to create an immersive experience for a participant of such an installation.

B.3 Pete the Bee

This project came about as the result of taking part in the Big Ears *Designing Inclusive Interactions* workshop which was curated by Dr Koichi Samuels of Queens University, Belfast and facilitated by Dr Brendan McCloskey of the University of Ulster. This project was hosted in Queens University and run in conjunction with the Drake Music Project Northern Ireland.

B.3.1 Big Ears 2015

The basic idea behind the Big Ears project was to assemble several groups of musicians, designers and makers with a view to develop bespoke instruments for musicians with physical disabilities. Each group was given three days, during which time they would be assigned a musician to work with, discuss the specific functionality required from the proposed instrument and design and build that instrument to the best of their ability. On the evening of the third day, the musician would then showcase the instrument in an ensemble performance in the Sonic Arts Research Centre (SARC) at Queens University (McEvoy, 2015).

I was part of a team of three designers. Our task was to create an instrument for a musician who enjoyed playing the piano but who found that moving from one chord shape to another in time with the music proved to be difficult. It took her quite some time reorientate her fingers into the correct shape required for a particular chord voicing and the frustration that this caused meant that playing the piano was not as enjoyable as she would have liked (ibid.). Our solution to this problem was to create a keyboard that would give her the ability to achieve polyphony with the use of only one button at a time. Along with this, we also included an infra-red sensor that the musician could use to articulate and manipulate the dynamics of the triggered audio.

In general, resources were quite scarce for this workshop so rather than buying pre-fabricated switches for our project, we were encouraged to create our own switches using conductive foam and copper tape. The basic idea behind these switches was that two strips of copper would be connected to a microcontroller, in this case an Arduino Uno.

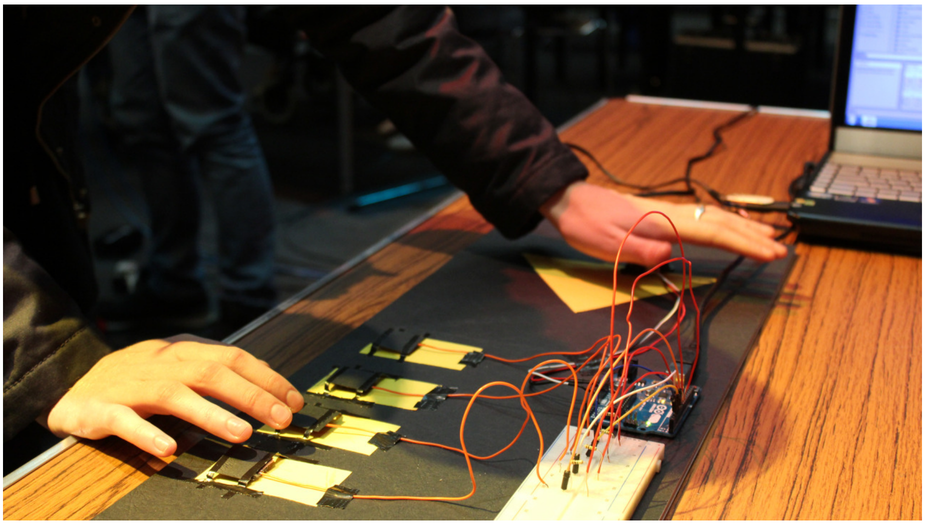

Figure B.8: Pete the Bee, a Max for Live hardware interface

One strip of copper would be connected to an analogue input and the other would be connected to an output pin operating at a logic level of 5 volts. The default state of the copper strip connected to the input pin would be boolean low (a value of zero). When this copper came into contact with the other strip, the state would change to boolean high (a value of one), due to the connection to a positive voltage. The conductive foam placed in between the two strips of copper changed the functionality of the sensor from a discrete on/off switch to a force sensitive resistor (FSR). Mapping the incoming sensor data from this switch to the amplitude of a midi note granted the user dynamic control over the articulation of the note. This autoschediastic approach to devising HCI solutions was completely new to me at the time and would go on to have a profound effect on how I approached similar HCI problems. A total of four of these improvised input devices were then connected to the Arduino Uno which in turn relayed midi messages to Ableton Live via Max for Live. The infra-red distance sensor was added to the instrument to grant the user further control of the wet/dry reverb mix.

B.3.2 Pitfalls of Inclusive Design

The event itself was sponsored by Ableton, which meant that the final musical output had to take place in the Ableton Live DAW environment. It also meant that the designers all had to use Max for Live in order to extrapolate sensor data from the microcontrollers used in the various instruments. Although I had never used Max before and I was much more comfortable using Csound for creating and processing audio, I found the experience to be an extremely useful exercise in adaptability.

Due to the time constraints and the obligation to work exclusively with Abelton Live, not a lot of possibilities were available to design the functionality of the software instrument from first principles. Instead, much of the focus of this project was centered around the physical functionality and design of the instrument. The main lesson that I learned from this experience was that when feasible, always consult with the end user at as many points throughout the design process as possible. This sentiment is often reiterated within the paradigm of inclusive design and to ignore such considerations can often result in inappropriate and unusable design(Newell and Gregor, 2000).

The musician that we were working with expressed her frustration at being asked to work with interfaces whose design seemed to make sweeping assumptions about her level of intelligence. She and several other musicians stated that most of the devices and interfaces she used on a day to day basis seemed to be designed for pre-school children (McEvoy, 2015). The assumption that her mental capacity was somehow linked to her physical condition was something that constantly irritated her. Therefore, she was very much involved with the aesthetic considerations of the instrument, both sonically and visually. It was as a result of her design requests that the instrument was given the name Pete the Bee (figure B.8).

B.3.3 Conclusion

One major flaw with this interface was the fact that it required Abelton Live in order to function. This in turn required the use of a relatively powerful laptop, a licence for both Live and Max for Live, and someone with knowledge of the software in order for it to operate. In addition to this, if any minor changes needed to be made to the physical component of the interface, a basic knowledge of electronics and possibly the Arduino IDE would also be required. This all meant that while the musician could play the instrument as part of the ensemble concert, she would likely never have the opportunity to play it again even though it was designed specifically for her.

B.4 Joystick

In this section I will discuss the design and construction of a digital instrument made specifically for a musician with cerebral palsy. I will also touch on how microcontrollers such as the Intel Galileo can be used in conjunction with audio programming environments to create bespoke digital instruments. This second attempt at designing a bespoke interface attempts to address some of the shortcomings mentioned in the previous section.

B.4.1 Big Ears 2016

One year after I had taken part in the first Big Ears workshop I was asked to return to Belfast to participate once again in designing and developing a bespoke musical interface. There were one or two changes made to the format of the workshop that set it apart from the previous year. Firstly, the projects were no longer to be undertaken in groups. Instead the complete design, development and implementation was to be handled by just one designer per instrument. The second major difference was that each designer now had several months in which to design and build an interface/instrument for the musicians that they were partnered with. The third deviation from the form of the previous year was that there were no longer any stipulations regarding the software used in the design of the instrument.

These changes not only meant that I could take a lot more time to carefully develop an interface specific to the needs of the musician I was partnered with, but I could also work in an audio programming environment that I was used to, namely Csound. Based on the experience I had accumulated in designing embedded instruments using the Intel Galileo (see Appendix A.3) this additionally meant that I could focus on creating an interface/instrument that could operate as a standalone device, not requiring any programming or DSP experience from the end user, thereby circumventing the issues I had with the project from the previous year.

B.4.2 Portability

Drawing from the research of Dr Victor Lazzarini (Lazzarini, Timoney, and Byrne, 2015), I was able to focus my attention on designing this instrument using Csound running on embedded devices such as the Intel Galileo. The advantage of using a device such as the Galileo to run Csound is that it allows the user the facility to connect input and output devices directly to the board itself, allowing communication with these inputs and outputs from directly within the Csound programming environment. This offers the possibility of streamlined hardware interfaces, granting the designer more choices in terms of portability and robustness. The functionality of Csound on the Intel Galileo is discussed in detail in the chapter A.3.

B.4.3 Hardware Design

The musician that I was partnered with for this project was a nine-year-old boy with cerebral palsy. Having spoken with the musician and his care assistant, the physical limitations of the musician and what criteria I needed to consider when building the device became apparent. The musician had very limited control of his hands and arms. His care assistant had told me that he had recently begun using his right hand to manipulate a joystick in order to control the movement of his wheelchair, and she mentioned his level of control over this device was increasing with time. I felt that this would be a good opportunity to create an interface that would allow the user to build up a greater level of control over devices that he would be using in a day to day context. With this in mind, I decided to design the interface/instrument based around a joystick controller similar to the one that he was learning to use.

This immediately gave me two control parameters to work with, the x-axis and y-axis of the joystick. I wanted to avoid adding any buttons to trigger an instance of the instrument as the user's control over his left hand was limited. It was important for me that the user should be able to fluidly interact with the instrument and with as little potential for frustration as possible. I felt that to rely on a second hand for instrument input was to add an unnecessary level of uncertainty to the system and invite

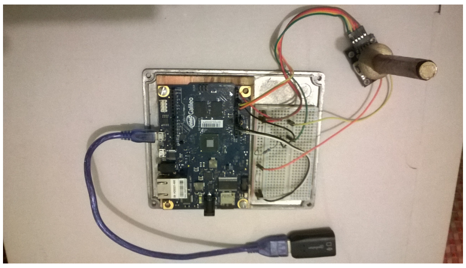

(a) Elevation of the Joystick interior

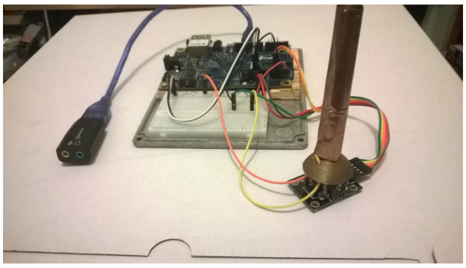

(b) Isometric view of Joystick interior

Figure B.9: Images of the internal workings of prototype Joystick Instrument with the 3D printed Joystick shaft

the possibility of inappropriate instrument design(Newell and Gregor, 2000).

To address this issue, I decided to make the joystick touch sensitive so that when the user made contact with the shaft of joystick an instance of the instrument would be triggered. This was achieved by creating a homemade FSR using copper tape and velostat. Velostat is a conductive material that exhibits reduced resistance once pressure is applied to its surface, making it ideal for the construction of a homemade FSR(Adafruit, 2013).

One strip of copper was wrapped around the shaft of the joystick and a wire was soldered to this strip. A layer of velostat was then wrapped around this strip of copper and then another layer of copper was wrapped around this. A wire was also soldered to this outer strip of copper. The first wire was connected directly to ground on the Intel Galileo and the second wire was connected to analog input 5 (A5) on the board. This wire was also connected to a pull-up resistor. Velostat acts as a weak conductor, meaning that when the joystick was squeezed the voltage going to A5 was increased. By setting a threshold within the Csound script, I was able use this behaviour to trigger an instance of the instrument whenever the joystick was touched.

B.4.4 Instrument Design

The design of the software instrument was based around a simple additive synthesiser using four oscillators. To allow for variation between each iteration of the instrument, the waveform of each oscillator was randomised using a pseudo-random process within Csound, creating timbres which varied slightly with each new iteration. A similar process was also applied to the amplitude of each oscillator. The function of this was not necessarily to provide indeterminate musical material, as the effect that these processes had on musical parameters such as amplitude was relatively tenuous. Instead, these pseudo-random processes were used to create subtle nuances within the instrument to counteract the predictable nature of a highly deterministic instrument design.

The four oscillators were summed and filtered using an emulation of a Moog diode ladder filter configuration. The cut-off frequency and resonance of this filter were controlled by the x and y axis respectively. In order to add another level of expressivity to the interface/instrument, I decided to further exploit the touch sensitive functionality of the joystick. I created a modulating oscillator with a view to implement some simple FM synthesis. I set the index of modulation to correspond to the amount of pressure being applied to the joystick.

In the original prototype, I used a Playstation 3 joystick module as the joystick for the interface, which meant that not only did I have a touch sensitive joystick with an x and y axis but I also had the use of a switch which was engaged when the stick was pressed down. Pressing down on this switch instigated a form of ring modulation to further enhance the spectral and harmonic properties of the instrument.

The final part of the instrument design was to offer the musician some control over choosing discrete pitches when performing on the instrument. I was aware that this instrument would eventually be used as part of an ensemble. With this in mind, I decided to incorporate some level of tonality into the instrument. I avoided using the continuous control of one of the axis as a means for selecting notes, as this portmanteau effect sounded over used and clichéd. I also had my doubts over the musician's ability to fully control the voicing of the instrument and the filter at the same time without rapidly jumping from one pitch to another.

Rather than a having the fundamental pitch being a control-rate parameter within Csound, meaning that the pitch of the instrument could be modified in much the same way one might modify the pitch of a trombone, I instead made it an i-rate parameter within Csound. This meant that the pitch of the instrument was determined by the

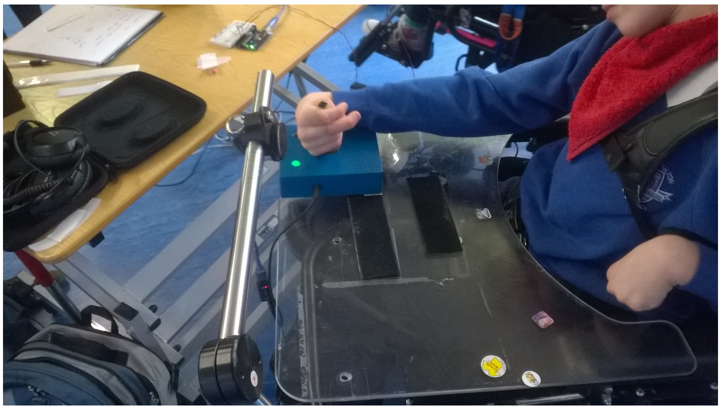

Figure B.10: Testing the Prototype Joystick

joystick's initial position on the y-axis. The advantage of this was that once a note had been triggered, it could then be held and the expressivity of the filter coupled with the FM modulation could be explored without fear of jumping through numerous pitches.

B.4.5 Implementation

One month before the instrument was to be completed, I travelled to Belfast to meet up with the musician. The first thing that I noticed was that I had assumed a greater level of control from the musician than he actually had and it seemed that he may not be capable of fully exploiting the subtleties of the instrument. When the musician squeezed the joystick there seemed to be some coarticulation (Godøy, Jensenius, and Nymoen, 2010) associated with the action of squeezing, meaning that whenever he grabbed the joystick he involuntarily pulled his hand toward his chest. This resulted in the 3D-printed shaft of the joystick becoming detached from the module, even though I had secured to the module with a strong adhesive. It also became apparent that the method of choosing notes was unable to be efficiently implemented by the musician. One positive that I took from this meeting was the fact that the musician seemed to enjoy the audio output from the instrument. After discovering the level of robustness that this interface required and the exact level of control that the musician had, I returned to the drawing board. I ordered a joystick that was almost identical to the one that the musician was accustomed to using on his wheelchair and began to redesign the note picking device of the instrument.

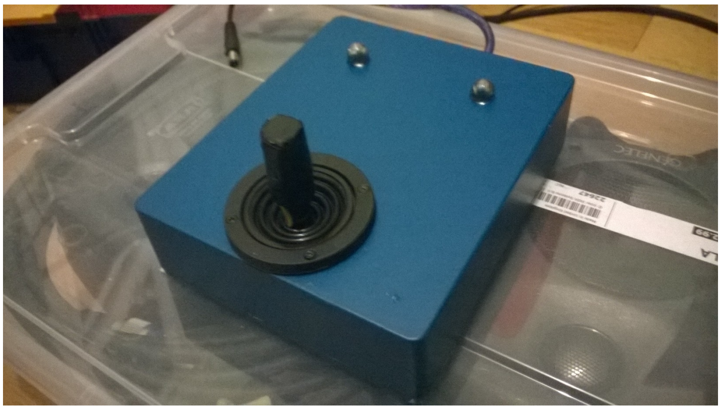

Figure B.11: The finalised Joystick Instrument

I eventually found a joystick of similar design to the one that the musician had on his wheelchair. It was capable of continuous control (which was surprisingly hard to find, most joysticks appear to be comprised of four discrete switches rather than two potentiometers). I prepared the shaft in the same manner as the previous joystick, wrapping it in a layer of copper tape, followed by a layer of velostat and finally another layer of copper tape. This allowed me to make the instrument touch and force sensitive as before.

There was however one element that I had to sacrifice and that was the switch which was part of the joystick module. This new joystick did not have a switch integrated into its design and as I wanted to avoid introducing too many controls to the interface, I decided to leave the switch out. This meant that the new design did not have the ring modulating functionality of the prototype, which was a shame but considering it was to be used in a tonally driven ensemble improvisation, I thought the inharmonious spectral content provided by the ring modulation would not be missed too much.

(a) The green light indicates that Csound has booted and is running

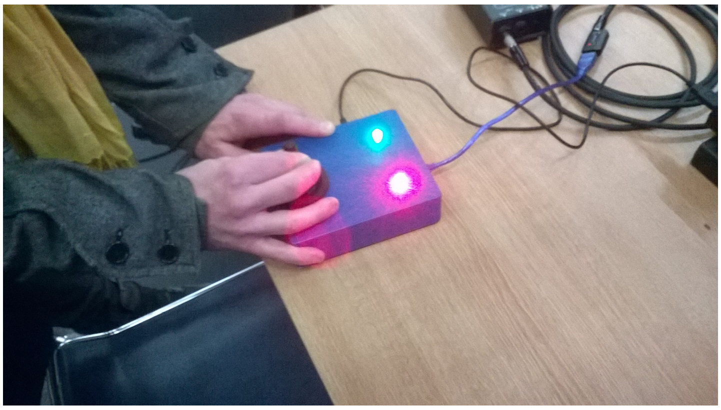

(b) The red light indicates that the instrument is active

Figure B.12: The Joystick in use at the performance in SARC

B.4.6 Conclusion

Unfortunately, the musician was unable to attend the ensemble performance in the SARC where the instrument was to be revealed due to illness. Instead, I performed as part of the ensemble under the direction of Dr Brendan McCloskey. The performance itself was scheduled to take place as part of the annual Sonorities Festival hosted in Queens University. I performed an improvised piece with the ensemble and was satisfied with the range of musicality possible from simply manipulating a single input device.

I did not get the opportunity to present the instrument to the musician in person but I have stayed in contact with the CEO of Drake Music Northern Ireland who has since informed me that he received his instrument and was delighted with how it performed. At the time of writing this thesis, the instrument is still functioning without error and without any further intervention from myself, almost two years after handing it over to the musician.

This project was a great success in terms of addressing the unresolved design challenges of the previous Big Ears project. The instrument required no expert knowledge in order for it to be used, no additional hardware apart from an audio connection for it to be heard and the design was robust enough to withstand at least two years of playing. There are certainly elements that can be improved on in this particular design, for instance a built-in amplifier and loudspeaker would enhance the portability of the instrument but overall I was pleased with the outcome.

B.5 Button Game

After my experience working on the Big Ears project, I realised the potential effect my research could have in the context of inclusive creativity. It seemed that many of the issues preventing people with special needs from creating music were often related to physical restrictions. These restrictions however were not on the part of the users, rather the designers of musical interfaces. With the skills I had accumulated over the course of my research, I felt that I could circumvent these issues and offer some of these would-be musicians a chance to engage in creative practices on their own terms, not being excluded because of poor interface design. I began to realise that the seemingly obscure knowledge I had accumulated over the years studying music technology could be put to good practical use.

This led me to get involved in music therapy sessions in a special needs school located in south Dublin. My task here was to involve the students in the process of creating music as opposed to their previous experience of music therapy which was to just passively sit and listen as someone played songs for them on a guitar. In order to achieve this I devised a simple generative synthesis program written in Csound that would compose music from minimal cues provided by the user.

B.5.1 System Design

The algorithm I used ensured that the music would always be in tune, the notes would always be in time, and the system would not require high levels of skillful engagement from the user in order to achieve satisfying musical results. The balance between reproducible results and sufficient variation in the sounds used was achieved through the use of indeterminate processes within Csound. For example, the exact pitch of the first instrument was determined through a pseudo-random process. A virtual eight-sided dice functioned as a note picking device, choosing one of eight notes within a just-intoned major scale.

B.5.2 Instrument Design

In total I created four instruments, all of which occupied a different spectral space. The first instrument was a harmonically rich pad sound which was based on a previous iteration that I had used within the EAREYEMOUTH installation. I had spent quite a bit of time developing an instrument that would be pleasant for a listener who may be particularly sensitive to aural stimulus, so with that in mind I decided to recycle the instrument as it was to be used with a similar context. The next two instruments were percussive in their amplitude envelope. One of which was a low-mid frequency instrument loosely modelled on a wood block. This was filtered and the initial attack was slowed down just enough to retain the percussive effect whilst softening the overall amplitude envelope. The second percussion instrument was loosely modelled on a chime and was created using a simple frequency modulation techniques. The fourth instrument was created to serve as a bridge between the graduated continuance of the pad and the percussive character of the wood block and the chimes. The initial attack was semi-percussive, perhaps close in character to the envelope of a piano but with a longer decay than either the wood block or the chime.

B.5.3 Physical Interaction

For the purpose of physically interfacing with the compositional system I made use of large BIGmack buttons (Inclusive Technology, 2010) which the students were already accustomed to using as input devices (see Figure B.13. These buttons are roughly about 170mm in diameter, made from plastic but designed to robust. Each button was a different colour, red, blue, green and yellow and I used this difference as inspiration when designing the compositional system.

B.5.4 Communication

I wrote a simple program within the Processing environment which would display a circle on a screen in front of the user. The colour of this circle would change intermittently from green to red to blue to yellow in a pseudo-random order. Each colour corresponded

Figure B.13: The buttons used to interface with the compositional system

to a instrument within Csound. Using OSC, I sent data pertaining to which button was pressed from Processing to Csound.

When a particular colour was present on the screen, the user would have the option to either turn the corresponding instrument on if it was currently off, or inversely turn the corresponding instrument off if it was on. For example, if red was shown on the screen, that meant that the user had the opportunity to either engage or disengage the pad instrument.

B.5.5 Conclusion

At first most of the students followed the cues on the screen as they appeared, turning on and off instruments when prompted. However, over time the students became aware of the exact effect they were having on the composition and began to make decisions as to whether or not they wanted a particular sound to enter into or be removed from the piece that they were writing. Not only did this mean that the users were actively engaging in the process of composing music, but as it turned out their hand-eye coordination and reaction time was improving as well (McDermott, 2016). This was perhaps due to the fact that they only had a small window of time in which to introduce or remove an instance of an instrument before the colour on the screen changed.

I was involved in these sessions over the course of three months and worked closely with the classroom assistant and music therapist to monitor the progress of each user. The reaction from each of the users was different with some enjoying the sounds more than others, however the overall feedback from the users and their carers was overwhelmingly positive. A short piece composed by one of the users is included in the flash-drive which accompanies this thesis.

Appendix C

Ouroborus

C.1 Ouroborus

This is the first work that I specifically wrote with live performance in mind. It is actually more accurate to refer to this work as a compositional system as opposed to a composition in its own right, for reasons which I will discuss later in this commentary. It is due to this that I have not included this piece in the main portfolio. I do however feel that this piece is significant enough to be included within the appendix.

Initially *Ouroborus* made use of two technologies that I had previously been experimenting with: the LEAP motion sensor and a hardware interface powered by the Intel Galileo. What I set out to achieve was to create a compositional system that would enable me to perform electronic music in a live setting whilst retaining the same level of detail and compositional control that was present in my fixed media pieces.

The idea behind this piece was to generate source audio from a six-string electric bass guitar and then send that source audio into a compositional system hosted in the Csound programming environment. The source audio could then be manipulated through a mixture of gestures tracked by the LEAP motion sensor, in addition to receiving sensor data from the Intel Galileo. The compositional system would be informed by input from the performer and the performer in-turn would be influenced by the output of the compositional system. Pseudo-random processes were used when looping the source material, giving the piece an indeterminate structure and generating a sense of interaction between the composer and the compositional system.

C.1.1 Process

Initially, I had used the LEAP as the sole method of sonic manipulation by mapping sensor data to musical parameters, but I soon discovered that there was significant deficit in the range of processes I could express using gestural motion alone. The entire range of expression afforded by the LEAP could only be realistically mapped to the parameters of one or two separate processes before those mappings began to encumber one another. For example, if control over the amplitude of process A was mapped to the location of the right hand on the y-axis, the right hand cannot be used to control any other process without the amplitude of process A being affected. It may have been possible to control a multitude of processes and parameters by employing some sort of modal behaviour to the LEAP, such as adding conditions describing which process to affect depending on what gesture had just been performed. However it has been noted that this kind of interface design can hinder the clarity and intuitive use of a system (O'Sullivan and Igoe, 2004) and I felt that it would go against the intentions behind the work to employ a compositional system that further abstracted the physical expression of the performer from the musical output.

In order to circumvent this issue, and in an effort to add another level of depth to the processing of the source material, I introduced the Galileo hardware controller (Figure C.1a) into the setup for this piece. Since the Galileo board (see Appendix A.3) has an ethernet connection I was able to communicate with the laptop acting as the server within this compositional system via the OSC protocol (see Appendix A.1.3). This controller afforded access to three potentiometers, one toggle switch and a photoresistor. Using the three potentiometers, it was possible to change the master gain of three digital audio processes which were used in the composition: synchronous granular synthesis, spectral manipulations performed using a phase vocoder, and temporal operations using several delay lines and looped audio. This afforded me continuous control over each process with

(a) The Galileo hardware controller

(b) The MTVSL, a programmable effects unit

Figure C.1: Hardware controllers used in Ouroborus

the LEAP, whilst retaining separate control over the presence of each process in the mix.

After reviewing several performances of this piece I felt that it did not fulfill my expectations of what I had intended the work to sound like. Although I recognised that the piece was intended for live performance and therefore scrutiny of the isolated audio should be conducted with care, I still felt that the range of sounds available to me during these early performances and the control I was afforded over the sonic properties of the material, was quite limited. In order to tackle this issue, I introduced one more piece of hardware into my assemblage of sensors and controllers.

The third piece of hardware that I introduced into the compositional system was a hardware controller I had built for my MA which I dubbed the MTVSL (Figure C.1b). This had no influence on the operations of the musical parameters within the Csound environment hosted on the server. It was instead used to manipulate the incoming source audio before it reached the stage of processing by the LEAP. The MTVSL is essentially a programmable digital multi-effects unit. Inside the enclosure is an Arduino Mega which is connected to a Raspberry Pi. The Raspberry Pi is programmed to run a Csound script upon boot and this Csound script is designed to communicate with the potentiometers and switches on the surface of the effects unit. The manner in which each of these potentiometers affects the sound is defined within the Csound script running on the Raspberry Pi. The processes that I implemented in the context of this composition were: a version of an adaptive fm (ADFM) synthesis unit (Lazzarini, Timoney, and Lysaght, 2007), a schroeder reverb, a variable delay and a rhythmic glitch effect implemented using the `bbcutm` opcode in Csound.

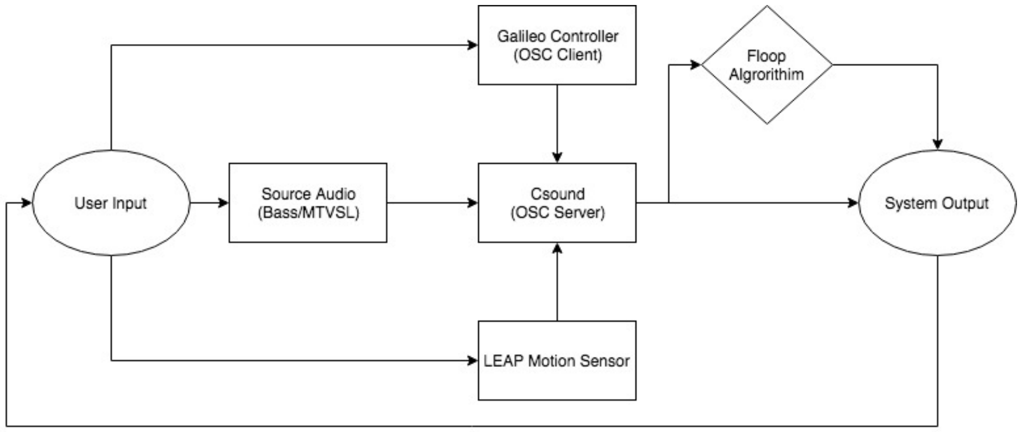

Figure C.2: Flow chart demonstrating the compositional system of Ouroborus

Although the introduction of the MTVSL opened up a whole new range of DSP possibilities, it also introduced quite a bit of noise into the system. While this did initially bother me, I decided to use this noise as a musical tool and an example of this can be heard at 06:10 in that particular performance. The complex waveform generated by noise within the system served as an ideal source to create a rich harmonic texture, upon which all other sound objects could be placed. A mixture of variable delay, ADFM synthesis and synchronous granular synthesis is used to create this particular texture.

```
instr floop

;chance operations
        icoin random 1, 10
        icoin = int(icoin)
        gidur random 1,5
        iamp random 0.3, 1
        ipitch  random 0.2, 2
        iloopstart random 0, 0.8
        iend random 1, 1.8
        iloopend = iloopstart+iend
        icrossfade random 0.001, 0.1

;use audio from buffer
        ifn = 1

;select whether loop runs forward, back and forth
or in reverse
        imode random 1, 3
        imode = int(imode)
        ipan random 0, 1

;Looper with declipping envelope
        asig flooper2 iamp, ipitch, 0, 1, 0.05, ifn,\
```

```
        0, imode
        aenv transeg 0.0001, p3*0.25, -2, 1, p3*0.75,\
        -2, 0.00001

        if icoin%7 == 0 then
            gadel = asig+gadel
        endif

        garev = asig+garev

        aL, aR pan2 asig, ipan
        aL*=aenv
        aR*=aenv

;master output controlled by Galileo Controller (gkin3)
        gaOutL = aL*gkin3
        gaOutR = aR*gkin3

endin
```

Code Excerpt C.1: Looping instrument with pseudo-random variables

In addition to the MTVSL, I began to include an Ebow as a means for exciting the bass strings. An Ebow is a handheld electromagnetic device which is designed to lengthen the sustain of guitar strings (Marossy, 2007). This provided me with the means to free up one of my hands so it could be used with the LEAP.

C.1.2 Structure

Ouroborus was always intended to be an improvisation governed by a loose structure. The structural aim was simply to provide an exposition of materials and subsequently develop this exposition according to indeterminate output of the compositional system.

In order to provide this indeterminate output, I designed a Csound instrument which would constantly write the incoming audio to a buffer of a fixed length. This buffer was then accessed using the `flooper` opcode, which was designed to read audio from a buffer and loop it based on a user defined start-time, duration and crossfade length (Lazzarini, Yi, et al., 2016). In the context of this piece, I applied pseudo-random processes to parameters such as the playback speed, amplitude and the direction of playback (see code excerpt C.1). The effect that these processes had on the audio output of the system was used to inform successive musical decisions on the part of the performer. This system of reciprocating feedback is the source of the title of this compositional system. I purposefully use the term compositional system, due to the fact that the actual organisation of musical ideas in this piece is ever changing, it is the system which facilitates the expression of these ideas which remains fixed.

C.1.3 Conclusion

Ultimately I was somewhat disappointed with the final outcomes of this investigation into gesture based composition and performance. I felt that the improvisatory nature of this work coupled with my unrefined control over the system led to an abundance of redundant variations, lacking in any great complexity (Lehmann and Kopiez, n.d.). However, the experience that performing *Ouroborus* gave me was invaluable. This compositional system gave me the opportunity to perform my electroacoustic music in a live context for the first time (not including live diffusion of fixed media), allowing me to experience first-hand, the tension that can exist between the performer of an unusual digital instrument and the audience. It also taught me the value of including fail-safes, kill-switches and intelligent interface design into a compositional system. No matter how many holes within the system I feel have been plugged, there will always be a new problem which appears right before a performance. These errors often cannot be foreseen due to the sheer volume of variables within a compositional system such as this. The only way to prepare for problems arising at inconvenient moments is to provided the performer with a means for absolute control over the system output. This will ensure that infinite feedback loops, out of bounds values and other systematic anomalies can be silenced when necessary. A modular approach to building the compositional system also helps with debugging if any issues do arise. After this work, I tried to incorporate a modular approach to building other compositional systems. Although the last live performance of this work was in late 2015, I still consistently use the compositional system I developed for *Ouroborus* when generating material for fixed media compositions.

Appendix D

List of Portfolio Works

Track No.	**Title**	**Year**	**Duration**	**File Type**
01	Sinmara	2015	09’54”	wav audio
02	Kinesia: The Third Law	2015	06’57”	mp4 movie
03	Cascando	2016	04’31”	wav audio
04	Galilean Moons: Callisto	2016	03’40”	wav audio
05	Proprioception	2016	06’52”	mov movie
06	Disintegrate	2017	09’33”	aif audio
07	iAmAtEase	2017	04’51”	wav audio
08	Djinn	2017	07’48”	wav audio
09	Tenterhooks	2018	08’28”	wav audio
10	Conatus	2017/18	15’02”	mp4 movie

Appendix E

Flash Drive Contents

- **Thesis - Electronic Copy**
- **Portfolio of Compositions** - See Appendix D
- **Supplementary Material**
  - Compositional Material
    - *Cascando*
      - Cascando_Tokyo_Rehearsal
    - *Conatus*
      - Conatus_AVBody'18(excerpt)
      - Conatus_AVBody'18.wav
      - Conatus(Alt)
      - Conatus.csd
      - Glove.py
      - skelton_dance_full_Glove2.pde
      - Conatus_Score
    - *Galilean Moons*
      - GalileanMoons_Full.wav
      - 5.1
        - Callisto.C.wav
        - Callisto.L.wav
        - Callisto.LFE.wav
        - Callisto.Ls.wav
        - Callisto.R.wav
        - Callisto.Rs.wav
    - *Kinesia*
      - Kinesia_Mermaid
      - Kinesia_(iFIMPaC)
      - Kinesia.csd
      - skelton_dance_full2.pde
      - Kineisa_Score
    - *Ouroborus*
      - Ouroborus.wav
      - OSCsend3.csd
      - Ouroborus-Improv.csd
    - *Proprioception*
      - Proprioception_Noisefloor'17
      - Proprioception_SoundThought'17
      - Proprioception.csd
      - skelton_dance_full_TorsoFix_2.pde
      - Proprioception_Score
    - *Sinmara*
      - RandOsc.csd
  - Interactive pieces
    - *Button Game*
      - button_game_4.wav
      - Button_Game.csd
    - *EAREYEMOUTH*
      - EEM_CultureNight.mp4
      - EEM_Google.mp4
    - *Joystick*
      - drakeFinal.csd
      - BigEars_interview.m4v
      - Joystick_Early_Demo.mov
    - *Sound;Waves*
      - WEC.csd
      - Sound;Waves_Demo.mov
      - WEC_Demo.mov
      - WEC.py

Appendix F

Supplementary Material

F.1 Compositional Material

The materials in this folder consists of additional material which is provided to support and contextualise the compositions within the main portfolio.

F.1.1 Cascando

Cascando_Tokyo_Rehearsal

This file contains a short video of rehearsals for the *Cascando* performance in Tokyo, Jun'16.

F.1.2 Conatus

Conatus_AVBody'18(excerpt)

A short video showing an excerpt from the performance at the AV Body symposium, Huddersfield, Jun'18.

Conatus_AVBody'18.wav

Full audio recording of the performance at the AV Body symposium, Huddersfield, Jun'18.

Conatus(Alt)

Alternative take of *Conatus*.

Conatus.csd

Csound code.

Glove.py

Python code for the Glove.

skelton_dance_full_Glove2.pde

Processing code for the Kinect.

Conatus_Score

Image file of graphic score.

F.1.3 Galilean Moons

GalileanMoons_Full.wav

The entire *Galilean Moons* Piece, complete with sections from the two other composers involved in the project.

5.1

Mono stems of a 5.1 mix of *Galilean Moons: Callisto*.

F.1.4 Kinesia

Kinesia_Mermaid

Promotional video for *Kinesia: The Third Law*. Shot on the same day as the version in the portfolio in the Mermaid Theatre, Bray, July'16.

Kinesia_(iFIMPaC)

Performance of *Kinesia* at Sounds Like This festival in Leeds, 2016. This was the first time that I performed the piece myself.

Kinesia.csd

Csound code

skelton_dance_full2.pde

Processing code for the Kinect.

Kinesia_Score

Image file of graphic score.

F.1.5 Ouroborus

OSCsend3.csd

Csound code for the Galileo hardware interface.

Ouroborus-Improv.csd

Main Csound code for the *Ouroborus* compositional system.

F.1.6 Proprioception

Proprioception_Noisefloor'17

Performance at Noisefloor, Stoke-on-Trent.

Proprioception_SoundThought'17

Performance at Sound Thought, Glasgow.

Proprioception.csd

Csound Code.

skelton_dance_full_TorsoFix_2.pde

Processing code for Kinect.

Proprioception_Score

Image file of graphic score.

F.1.7 Sinmara

RandOSC.csd

Csound code used to generate some of the source material in *Sinmara*

F.2 Supplementary Material

The materials in this folder consists of additional material which is provided to support and contextualise the interactive pieces described in this thesis.

F.2.1 Button Game

button_game_4.wav

A composition by one of the users from St.John of Gods created using the Button Game compositional system

Button_Game.csd

Csound code

F.2.2 EAREYEMOUTH

EEM_CultureNight.mp4

Footage from the installation at Culture Night, Galway'14

EEM_Google.mp4

Footage from the installation at Creative Tech Fest, Dublin'14

F.2.3 Joystick

drakeFinal.csd

Csound code

BigEars_interview.m4v

Interview taken during the early stages of the instrument build at Queen's University

Joystick_Early_Demo.mov

Early demo video of the piece with original, 3D printed joystick shaft

F.2.4 Sound;Waves

WEC.csd

Csound code

Sound;Waves_Demo.mov

Short demo video of the installation

WEC_Demo.mov

Short demo video of the Edison communicating with Csound over OSC

WEC.py

Python code for the Edison

Bibliography

Adafruit (2013). *Pressure-Sensitive Conductive Sheet (Velostat/Linqstat)*. `https://www.adafruit.com/product/1361`. [Accessed 03/18].

Barnett, R. (2015). *RTIMULib2/Fusion11.py*. GitHub. URL: `https://github.com/richardstechnotes/RTIMULib2/blob/master/Linux/python/tests/Fusion11.py`.

Becker, J. (1968). "Percussive Patterns in the Music of Mainland Southeast Asia". In: *Ethnomusicology* 12.2, pp. 173–191.

Beckett, S., S. Lawlor, and J. Pilling (2014). *The collected poems of Samuel Beckett: a critical edition*. New York: Grove Press.

Boulez, P. (1986). "Technology and the Composer". In: *The Language of Electroacoustic Music*. Ed. by S. Emmerson. London: Macmillan, pp. 5–14.

Cage, J. (1961). *Silence : lectures and writings*. Middletown , Connecticut: Wesleyan University Press.

Campbell, J. (2008). *The hero with a thousand faces*. Novato, CA: New World Library.

Carson, C. (1996). *Last Night's Fun: a book about music, food and time*. London: Pimlico.

Clarke, M. (2000). "FOF and FOG Synthesis in Csound". In: *The Csound Book*. Ed. by Richard Boulanger. Cambridge, MA, USA: MIT Press, pp. 293–306.

Cohen, L. (1992). *The Future*. New York: Colombia Records.

Colgan, Alex (2014). *How Does the Leap Motion Controller Work?* URL: `http://blog.leapmotion.com/hardware-to-software-how-does-the-leap-motion-controller-work/`.

Collins, N. (2002). "The BBcut Library". In: *Proceedings of the International Computer Music Conference*. Göteborg, Sweden: ICMC.

Croft, J. (2007). "Theses on liveness". In: *Organised Sound* 12.01, p. 59.

Csikszentmihalyi, M. (2009). *Flow: the psychology of optimal experience*. New York: Harper Row.

Curtis, A. (2015). *The Century of The Self*. London, UK: Universal Studios.

Darmon, E. (2007). *Pierre Henry: The Art of Sounds*. [Video File]. URL: `https://www.youtube.com/watch?v=Zsb1w-3iFa8`.

Davies, O. (2003). "The Nightmare Experience, Sleep Paralysis, and Witchcraft Accusations". In: *Folklore* 114.2, pp. 181–203.

Davis, T. (2010). "Complexity as Process: Complexity-inspired approaches to composition". In: *Organised Sound* 15.02, pp. 137–146.

Delap, G. (2016). "At the Cutting Edge of Unobtainium: Collaborating with the Next Generation Sound Synthesis Project". In: *Metamorphosis: Collaboration in Research and Arts Practice*. Smithfield, Dublin.

d'Escriván, Julio (2006). "To sing the body electric: Instruments and effort in the performance of electronic music". In: *Contemporary Music Review* 25.1-2, pp. 183–191.

Dillon, J. and C. Doolan (2017). *What is Folklore?* URL: `https://soundcloud.com/folklore_podcast`.

Donnarumma, M. (2012). *Hypo Chrysos*. URL: `https://vimeo.com/37921373`.

Donnarumma, M. (2016a). *Corpus Nil*. [Video File]. URL: `https://vimeo.com/152710490`.

Donnarumma, M. (2016b). *Xth Sense*. URL: `http://marcodonnarumma.com/works/xth-sense/`.

Donnarumma, M. (2018). *Eingeweide*. URL: `https://vimeo.com/298134433`.

Durkheim, E., C. Cosman, and M. Cladis (2001). *The Elementary Forms of Religious Life*. Oxford, England: Oxford University Press.

Emmerson, S. (1986). "The Relation of Language to Materials". In: *The Language of Electroacoustic Music*, pp. 17–39.

Emmerson, S. (2016). *Living Electronic Music*. London: Routledge Taylor & Francis Group.

Esposito, P. and T. Kasai (2017). “Butoh Dance, Noguchi Taiso, and Healing”. In: *Oxford Handbooks Last accessed 02/18.*

Fiebrink, R., D. Trueman, and P.R. Cook (2009). “A metainstrument for interactive, on-the-fly machine learning”. In: *Proceedings of the International Conference of New Interfaces for Musical Expression (NIME’09).*

Fischman, R. (2012). *”Ruraq Maki”*. [Video File]. URL: `https://vimeo.com/55093629`.

Fischman, R. (2014). “A Manual Actions Expressive System”. In: *NOVARS Matinee*. Manchester, UK: Springer.

Garcia-Rosa, P., G. Bacelli, and J. Ringwood (2015). “Control-Informed Geometric Optimization of Wave Energy Converters: The Impact of Device Motion and Force Constraints”. In: *Energies* 8.12, pp. 13672–13687.

Gibson, J. (2015). *The Ecological Approach to Visual Perception: classic edition*. New York, NY: Psychology Press.

Glassie, H. (1982). *Passing time in Ballymenone*. Bloomington: Indiana University Press.

Godøy, R.I (2003). “Motor-Mimetic Music Cognition”. In: *Leonardo* 36.4, pp. 317–319.

Godøy, R.I, A.R Jensenius, and K. Nymoen (2010). “Chunking in Music by Coarticulation”. In: *Acta Acustica united with Acustica* 96.4, pp. 690–700.

Hunt, D. (2015). *Raspberry Pi, Beaglebone Black, Intel Edison–Benchmarked*. `http://www.davidhunt.ie/raspberry-pi-beaglebone-black-intel-edison-benchmarked/`. [Last accessed 02/18].

Ibañez, R., Á. Soria, A. Teyseyre, and M. Campo (2014). “Easy gesture recognition for Kinect”. In: *Advances in Engineering Software* 76, pp. 171–180.

Inclusive Technology, Ltd. (2010). *BIGmack Buttons*. `http://www.inclusive.co.uk/ablenet-bigmack-p2039`. Last Accessed 04/18.

Intel (2013). *Intel Galileo Board*. `https://ark.intel.com/products/78919/Intel-Galileo-Board`. [Last accessed 02/18].

Intel (2014). *Intel Edison Compute Module*. https://ark.intel.com/products/84572/Intel-Edison-Compute-Module-IoT. [Last accessed 02/18].

Kenny, S (2016). *LeapInto*. https://github.com/digithree/leap-into. [Last accessed 02/18].

Klein, N. (2008). *The Shock Doctrine: The Rise of Disaster Capitalism*. Toronto: Vintage Canada.

Kopplin, M, C Bartlett, J.H Holland, and C Iten (2008). *Flickwerk: The Aesthetics of Mended Japanese Ceramics*. Münster: Museum für Lackkunst.

Krefeld, V. and M. Waisvisz (1990). "The Hand in the Web: An Interview with Michel Waisvisz". In: *Computer Music Journal* 14.2, pp. 28–33.

Lazzarini, V. (2014). *Session 09: Granular Synthesis*. Lecture Notes, Software Sound Synthesis, Maynooth University. Delivered 02/12/2014.

Lazzarini, V., E. Costello, S. Yi, and J. ffitch (2014). "Development Tools for Ubiquitous Music on the World Wide Web." In: *Ubiquitous Music. Computational Music Science*. Ed. by Pimenta M. Keller D. Lazzarini V. Cham: Springer, pp. 111–128.

Lazzarini, V. and J. Timoney (2011). "The ModFM synthesis vocoder". In: *International Computer Music Conference (ICMC) 2011*. Huddersfield, UK: ICMC.

Lazzarini, V., J. Timoney, and S. Byrne (2015). "Embedded Sound Synthesis". In: *Proceedings of the Linux Audio Conference 2015*. Mainz, Germany: Johannes Gutenberg University.

Lazzarini, V., J. Timoney, and T. Lysaght (2007). "Adaptive FM Synthesis". In: *DAFX-07 the 10th Int. Conference on Digital Audio Effects*.

Lazzarini, V., S. Yi, J. ffitch, J. Heintz, Ø. Brandtsegg, and I. McCurdy (2016). *Csound*. Cham, Switzerland: Springer International Publishing.

Lehmann, A. and R. Kopiez (n.d.). "Revisiting composition and improvisation with a historical perspective". In: *ESCOM 2002 Proceedings*. URL: http://escom.org/proceedings/ESCOM2002/sources/Pdf/symposium/Lehmann.pdf.

Leman, M. (2008). *Embodied Music Cognition and Mediation Technology*. Cambridge: MIT Press.

Liberman, A. M. and I. G. Mattingly (1985). “The motor theory of speech perception revised”. In: *Cognition* 21.1, pp. 1–36. DOI: 10.1016/0010-0277(85)90021-6.

Linderoth, S. (n.d.). “Appendix E. Modal Frequency Ratios”. In: *Csounds.com - Csound Manual*. [Last accessed 02/18].

Mann, S. (2007). “Natural Interfaces for Musical Expression: Physiphones and a Physics-based Organology”. In: *Proceedings of the 7th International Conference on New Interfaces for Musical Expression*. NIME ’07. New York, New York: ACM, pp. 118–123.

Marossy, P.J. (2007). *The Technology of the Ebow*. https://web.archive.org/web/20120415220647/http://www.diyguitarist.com/DIYStompboxes/EbowTech.htm. [Accessed 02/18].

McCurdy, I. (2012). *Choir.csd*. http://iainmccurdy.org/CsoundRealtimeExamples/Cabbage/Instruments/Synths/Choir.csd. Csound/Cabbage code.

McDermott, K. (2016). “Assistive Technology in Therapuetic and Health Settings”. unpublished.

McEvoy, A. (2015). *Big Ears 2015 Evaluation Report*. https://www.dropbox.com/s/zgpvk8h05z2ziol/BigEars2015Evaluation.pdf?dl=0. Last accessed 04/18.

Microsoft (2009). *Xbox 360:Project Natal 101*. https://web.archive.org/web/20120121223600/http://download.microsoft.com/download\\/A/4/A/A4A457B3-DF5D-4BF2-AD4E-963454BA0BCC/ProjectNatalFactSheetMay09.zip. [Last accessed 02/18].

Microsoft (2012a). *Kinect Sensor*. https://msdn.microsoft.com/en-us/library/hh438998.aspx. [Last accessed 02/18].

Microsoft (2012b). *Tracking Users with Kinect Skeletal Tracking*. https://msdn.microsoft.com/en-us/library/jj131025.aspx. [Last accessed 02/18].

Mindflux (2002). *Essential Reality - P5 Glove*. URL: http://www.mindflux.com.au/products/essentialreality/p5glove.html.

Moore, A. and D. Moore (2011). *Sonic Art: Recipes and Reasonings*. https://www.sheffield.ac.uk/polopoly_fs/1.152862!/file/sonicart_recipesandreasonings.pdf. [Last accessed 02/18].

Murphy, D. (2011). "Hate Couture: Subcultural Fundamentalism and the Serbian Black Metal Music Scene". PhD thesis. Maynooth University.

Newell, A. F. and P Gregor (2000). ""User sensitive inclusive design"— in search of a new paradigm". In: *Proceedings on the 2000 conference on Universal Usability - CUU 00*.

Norman, D. (2013). *The Design of Everyday Things*. New York: Basic Books.

O'Cinneide, A. (2005). "Introducing PVSPITCH: a pitch tracking opcode for Csound". In: *Sounds Electric '05 Conference*.

O'Connor, Roisin (2016). "Anish Kapoor gets his hands on 'pinkest pink' after being banned from use by its creator". In: *www.independent.co.uk*. URL: https://www.independent.co.uk/arts-entertainment/art/news/anish-kapoor-pinkest-pink-blackest-black-paint-war-a7497751.html.

O'Sullivan, D and T Igoe (2004). *Physical Computing: Sensing and Controlling the Physical World with Computers*. Boston, MA, United States: Course Technology Press, p. xxv.

Pellegrino, G. Di, L. Fadiga, L. Fogassi, V. Gallese, and G. Rizzolatti (1992). "Understanding Motor Events: A Neurophysiological Study". In: *Experimental Brain Research* 91.1, pp. 176–180.

Pfeifer, R. and J. Bongard (2007). *How the body shapes the way we think: a new view of intelligence*. Cambridge: Massachusetts.

Poff, Z. (2015). *Alex Rice Piezo Preamplifier*. https://www.zachpoff.com/resources/alex-rice-piezo-preamplifier/. [Last accessed 03/18].

Purcell Verdun, L. (2016). *What Is Vocal Fry? A Speech-Language Pathologist Explains*. online video. URL: https://www.youtube.com/watch?v=4L7-9N1xQZA.

Ramon, M.C. (2014). "Intel Galileo and Intel Galileo Gen 2". In: *Intel® Galileo and Intel® Galileo Gen 2: API Features and Arduino Projects for Linux Programmers*. Berkeley, CA: Apress, pp. 1–33.

Reas, C. and B. Fry (2004). "Processing.Org: Programming for Artists and Designers". In: *ACM SIGGRAPH 2004 Web Graphics*. SIGGRAPH '04. Los Angeles, California: ACM, pp. 3–. ISBN: 1-58113-896-2.

Reas, C. and B. Fry (2015). *Getting started with Processing*. San Francisco: Maker Media.

Risko, E. F. and Sam J. Gilbert (2016). "Cognitive Offloading". In: *Trends in Cognitive Sciences* 20.9, pp. 676–688.

Rowe, R. (1993). *Interactive music systems: machine listening and composing*. Cambridge, MA: The MIT Press.

Ryan, J. (1991). "Some remarks on musical instrument design at STEIM". In: *Contemporary Music Review* 6.1, pp. 3–17.

Schachner, J.C (2012). "The Body in Electronic Music Performance". In: *Proceedings of the 9th Sound and Music Computing Conference (SMC'12)*. Aalborg University, Copenhagen: SMC, pp. 194–200.

Schechner, R. (2015). *Performance Theory*. London: Routledge.

Schwartz, B (2004). *The Paradox of Choice: Why More is Less*. New York: Ecco.

Small, C. (2010). *Musicking: The Meanings of Performing and Listening*. Middletown: Wesleyan Univ. Press.

Smalley, D. (1986). "Spectro-morphology and Structuring Processes". In: *The Language of Electroacoustic Music*. Ed. by S. Emmerson. London: Macmillan, pp. 61–93.

Smalley, D. (1997). "Spectromorphology: explaining sound-shapes". In: *Organised Sound* 2.2, pp. 107–126.

Smeenk, R. (2014). *Kinect V1 and Kinect V2 fields of view compared*. `https://smeenk.com/kinect-field-of-view-comparison/`. [Last accessed 02/18].

Smith, J. (2017). *Physical Audio Signal Processing: for Virtual Musical Instruments and Digital Audio Effects*. W3K Publishing.

Sonami, L. (2000). *"Why dreams like a loose engine"*. [Video File]. Roulette Intermedium. URL: https://vimeo.com/11316136.

Sonami, L. (2010). *lady's glove*. URL: http://sonami.net/ladys-glove/.

Sparkfun (2014a). *SparkFun Block for Intel® Edison - 9 Degrees of Freedom*. https://www.sparkfun.com/products/retired/13033. [Last accessed 02/18].

Sparkfun (2014b). *SparkFun Block for Intel® Edison - Battery*. https://www.sparkfun.com/products/13037. [Last accessed 02/18].

Sparkfun (2015). *SparkFun Block for Intel® Edison - GPIO*. https://www.sparkfun.com/products/13038. [Last accessed 02/18].

Toshiharu, K. and K. Parsons (2003). "Perception in Butoh Dance". In: *Memoirs of the Hokkaido Institute of Technololy* 31, pp. 257–264.

Vercoe, B. and D. Ellis (1990). "Real-Time CSOUND: Software Synthesis with Sensing and Control". In: *The Oxford Handbook of Innovation*. Ed. by S. Arnold, Computer Music Association, I.C.M. Conference, and G. Hair. Oxford: ICMC Glasgow, pp. 209–211.

Vickhoff, B. (2008). *A perspective theory of music perception and emotion*. Gothenburg: University of Gothenburg.

Vik, C. (2014). *Kinectar:Your MIDI-Kinect interface*. http://www.ethnotekh.com/software/kinectar/. [Last accessed 02/18].

Vonnegut, K. (2006). *Player Piano*. New York, NY: Dial Press Trade Paperbacks.

Waisvisz, M. (2004). *No Backup*. STEIMTUBE. URL: https://www.youtube.com/watch?v=U1L-mVGqug4.

Waisvisz, M. (2006). *The Hands*. URL: http://www.crackle.org/TheHands.htm.

Waters, S. (2007). "Performance Ecosystems: Ecological Approaches to Musical Interaction". In: *Proceedings of Electroacoustic Music Studies Network (EMS'07)*. Leicester, United Kingdom.

Wilkins, Bryan (N.D). *RE20 Broadcast Announcer Microphone w/ Variable-D by Electro-Voice*. URL: https://www.electrovoice.com/product.php?id=91.

Willerslev, Rane (2007). *Soul hunters: hunting, animism, and personhood among the Siberian Yukaghirs*. Berkeley: University of California Press.

Wishart, T. (2000). "Sonic Composition in Tongues of Fire". In: *Computer Music Journal* 24.2, pp. 22–30.

Wright, M. (2005). "Open Sound Control: An Enabling Technology for Musical Networking". In: *Org. Sound* 10.3, pp. 193–200.

Xenakis, Iannis (1958). *Concret PH*.

Zappa, F. and P. Occhiogrosso (1999). *The real Frank Zappa book*. New York: Simon Schuster.

Zhengyou, Z. (2012). "Microsoft Kinect Sensor and Its Effect". In: *IEEE Multimedia* 19.2, pp. 4–10.