Two novel performance pieces intended to explore musicality within gestural mapping and game-data interpretation.

Ian Caballero

Dissertation submitted towards Masters by research degree.

University of Huddersfield

Department of Music, Humanities and Media

September 2016

Contents

Abstract

A project exploring two methods of interacting with music. As part of this project, two performance pieces are submitted, each with their own purpose-built interfacing system. Gestural mapping is carried out within these systems to interpret data within these performances. These pieces are examined and evaluated and the interfaces are compared to conventional instruments. The first piece, “Leaps”, makes use of non-tactile hand tracking technology to derive its data. Hand gestures are mapped into musical output. The second piece, “Chess” derives musical expression from a game of chess. Data about the game is collected in real time and this is instantly mapped into sound. The pieces are discussed in detail within the following document.

Introduction

This project aims to explore the possibilities of mapping gestures in various contexts to a sonic output, in order to create realistic expression and musicality within the context of a performance. The project is based around two interfaces created on Cycling 74's Max 6 (Max Msp, 2016). The interfaces interpret input from performers in their own unique way, each providing a distinct take on mapping gestural input. The two pieces in question are "Leaps" and "Chess". The idea behind collating the two pieces is to explore these specific and distinct forms of interaction with music. These interfaces are coupled with a score for performance and should be thought of as part of the piece itself.

"Leaps" explores the use of non-tactile interfacing within performance. It places the performer in an unfamiliar circumstance in which the usual tactile sensory feedback of an instrument is removed. In brief, Leaps is a piece which is controlled with both hands with no tactile interfacing required. Performers rely on the visual feedback (a screenshot of this can be seen in page 14, Figure 2) shown on a screen to interact. Gestures are repurposed based on the visual feedback so performers must maintain a line of sight. Additionally, the interface has a menu-based system that limits the speed of specific sonic developments to the performer's ability to navigate through the menu. The result is generally a slow-changing piece. Usually controllable parameters such as pitch and expression become much harder to access accurately. Furthermore, by using actions that are likely to have an obvious meaning for the user and audience (such as a "circle" motion [Leaps video timestamp: 00:15:01] to control rate of playback), a clear link between actions and sound is established. That is to say, the interface's own gestural mapping metaphorically reflects the sonic output it controls.

"Chess" takes its mapping in a completely different direction. The piece consists of two performers playing out a full game of chess. Each move is instantly noted by a third party (off-stage) which then outputs a sonic interpretation of the game. The interface aims to follow the game's main pivotal points (such as opening moves, captures, strategic retreats, check, and checkmate) and represent these in sound. The focus of this piece is on the tension that is created within a game of chess. Less obvious data such as thinking time is also used to create or alleviate tension in the music accordingly [Chess video timestamp 00:05:04]. Whilst in conventional instruments, the sonic output is based on a performer's direct investment in creating the sound; in "Chess", the interface's reactions to moves is the way sound is achieved. Indeed, this is the solitary form of input with which the performer can influence the direction that the piece will take. This restricted form of input still provides plenty of data for interpretation; however, the conflict between creating a good performance and playing the game is a challenge that the performers must overcome. This is also an area to be explored.

Research aims

Throughout the course of this project, I aim to explore and discuss interaction between the performer and the sounds produced within the pieces submitted, and how it contributes to the musicality of the performance. I will be discussing mapping the human-computer interaction in order to replicate the intuitiveness in interaction that is regularly found in an acoustic instrument.

Two pieces are to be composed alongside a Max/MSP patch which will contain all aspects of data collection, mapping, and sound generation. These pieces will look at different aspects of creative interaction between performers and their instruments.

I will also be exploring how these interfaces fit within the existing models and definitions of "an instrument". As part of this, I will evaluate the validity of referring to the interfaces within the pieces created as “instruments”.

Context of project

The works submitted as part of this project draw influence and inspiration from various sources and experiences. These are discussed within this section. Both “Chess” and “Leaps” are pieces that can be performed by people with no musical background, as an in-depth knowledge of music theory is not at all required. This will hopefully enable people to explore their creativity without the need to become skilled at a particular musical instrument. This idea was used in the development of Propellerhead’s “Players” software (“Propellerhead,” 2016) feature in which a single press of a note can trigger an entire range of musical outcomes of a potentially complex nature.

Non-tactile interfaces for music have been available for a long time; however, the advent of easily available interfacing technologies has made the process much more accessible to the average composer. Works exploring the potential of gestural control can also feature the use of wearable technology as tracking hardware ("the Gloves", 2016). These interfaces share the common factor of attempting to mimic actions commonly found in day-to-day life as metaphors for the sounds they produce. An example of this concept can be found in Imogen Heap’s “The Gloves” ("the Gloves", 2016) in which she is able to control parameters such as filter cutoff by turning her hand as if she was turning a physical dial. Earlier instruments have shown this similar user-interaction such as Michael Waisvisz’s “The Hands” (Krefeld & Waisvisz, 1990) in which the user wears and holds a controller in which small movements can be tracked and used to manipulate sound. Whist “The Hands” is considered by its creator an instrument, I cannot see “Leaps” as this, because whilst the sonic output can be manipulated at the will of the performer, there is fundamentally, a limited number of sounds that each section is based on. The gestural interface of “Leaps” is based around familiar actions so that they can be easily understood for intuitive performance. A similar concept was used on Steinberg’s “Cubase iC Air gesture control technology” (Steinberg, 2013) where these familiar gestures (like hold palm up to stop a track) were used to facilitate the workflow of producers. Motion tracking technology is now very readily available. Microsoft’s “Kinect” ("Kinect - Windows app development", 2016) has led to the creation of many non-tactile interfaces for music, and as more of these are being made, exploration into how to create adequate musical expression with this interface becomes more and more relevant. The NIME paper “Creating Musical Expression Using Kinect” (Yoo, Beak, & Lee, 2011) looks at this in the context of their specific interface. One of the challenges faced when creating “Leaps” was to create an environment for the user that feels intuitive and one that allows for free and varied expression.

Taking data from non-musical systems and adapting it for the purpose of musical output is also not a novel idea. “Birds on the Wires” by Jarbas Agnelli (Agnelli, 2009) was created from a single photograph of birds perched on power cables. Their positioning was noted by the composer and each point was subsequently mapped to a musical note using a one-to-one paradigm in which every note corresponds to a specific predetermined note. Systems utilising much more complex mapping include Bartholomäus Traubeck’s “Years” (Traubeck, 2011) in which Bartholomäus uses a small camera to translate the rings of a tree to MIDI data and subsequently to the sound of a piano. This system (whilst more complex in its mapping) does not feature a live performance element. That is to say, the data being converted is always the same, rendering these pieces as a form of playback for a predetermined recording or score. With “Chess”, the sonic output is based on the live game as it is played. This expands the possibilities in terms of what input data the system can interpret, as it also takes player thinking-time and player reaction-times into account. It also opens up the piece to the interpretation of the performers, in this case - infinitely - as there are more than 100 billion possible piece combinations in a forty-move chess game. The interface of chess is visually similar to that of the "Reactable", an instrument developed at Universitat Pompeu Fabra. ("Reactable | Music Technology Group", 2016) The similarity lies in the way the user interacts with a table-top interface. Performers place objects on the "Reactable" which then shows a clear visual representation of what the user is doing on its touch-screen interface. This directly contrasts the user interaction in "Chess", where the performer is limited to the game of chess as a control input, and there is no further expression for performers to add other than the moves themselves. Perhaps the most relevant game-based performance happened in 1968, when a game of chess was played between Marcel Duchamp and John Cage at the Ryerson Theatre in Toronto. This game was a performance piece to headline the "Sightsoundsystems Festival" ("Marcel Duchamp & John Cage Play Musical Chess", 2016). This game was a performance using a special electronic chess board. The board contained sensors on each square that sent electrical signals to a system which altered the amplification levels of musicians performing around the room. Marcel Duchamp was once quoted to have said: "I have come to the personal conclusion that while all artists are not chess players, all chess players are artists." ("Marcel Duchamp & John Cage Play Musical Chess", 2016) This piece holds a lot of relevance for the early development of this project, as it provided inspiration to create a piece that not only controls the levels of music but in fact generates an entire sound world. The piece also holds relevance for the future development of "Chess" because of the potential to someday move on to using a purpose-built chess board to input data.

Leaps

The interface

"Leaps" is an interface for music intended to remove the typical tactile aspects of interacting with a musical instrument. This lack of reference point and tactile feedback generally found within an acoustic instrument, places the performer in a situation where they must explore unfamiliar details within their own performance. The interface hardware consists of a computer and a Leap Motion receiver (Leap Motion, 2016) which sends data to the computer. The hardware uses two cameras sensitive to infrared light, which is constantly emitted from three in-built LEDs at a wavelength 850 nanometres. The two cameras assemble a three-dimensional plot of hand positions, from which accurate X, Y, Z data of each joint in the hand can be collected. This data feeds the Max 6 patch which in turn, outputs audio utilising four channels in a 5.1 setup with sub management. (The videos submitted alongside this are stereo recordings). The system functions well in low lighting as the cameras particularly seek lighting at a wavelength of 850 nanometres, which is emitted from the Leap Motion itself. This feature also allows for maintaining accurate readings in the context of dynamic performance lighting.

Whilst the interface utilises no physical controls, it also takes the unconventional approach of delivering visual feedback for the performer. This feedback has been found to be crucial in the performer’s awareness of their place within certain parameters. The visual feedback is not visible to the audience but is there to aid the performer as well as to display the score.

Human-interface interaction

The gestures:

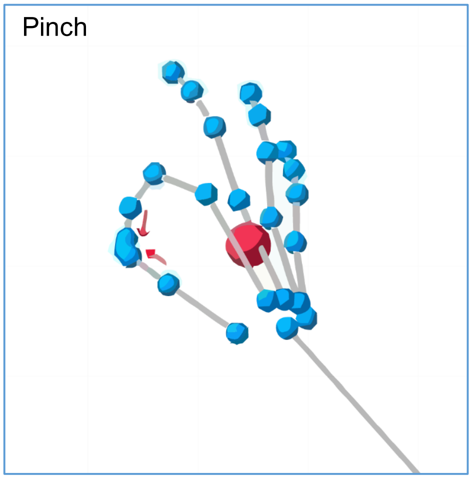

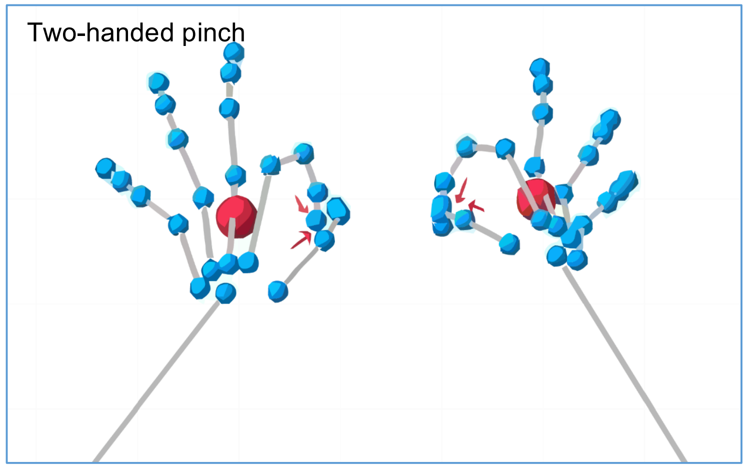

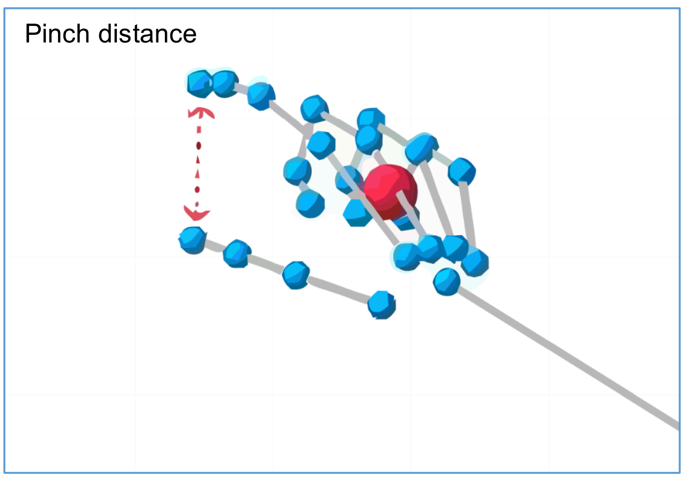

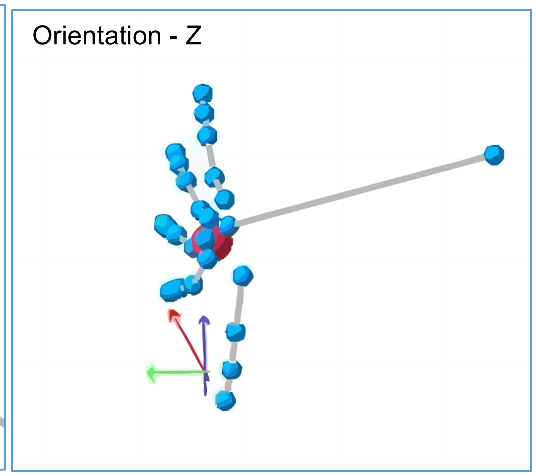

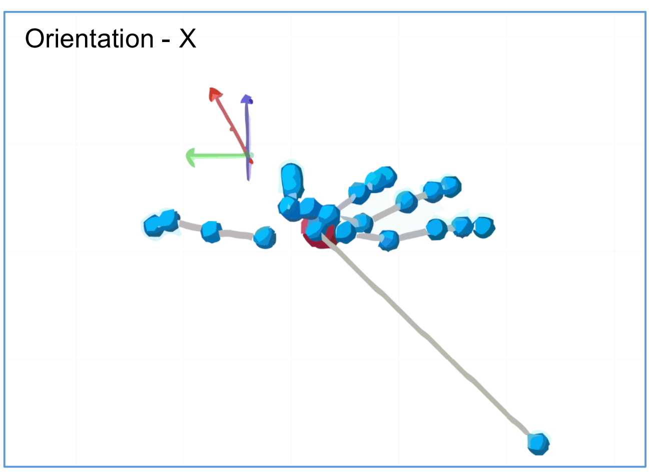

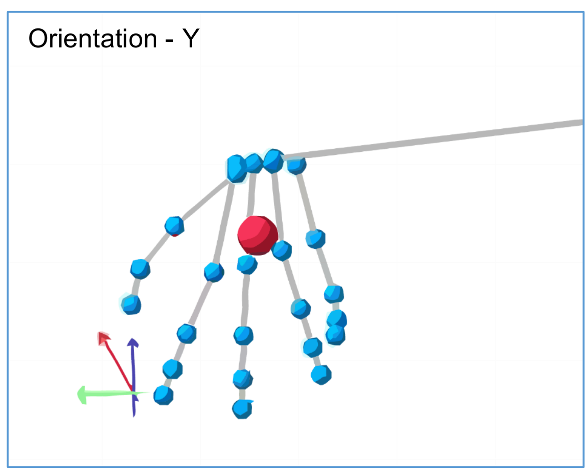

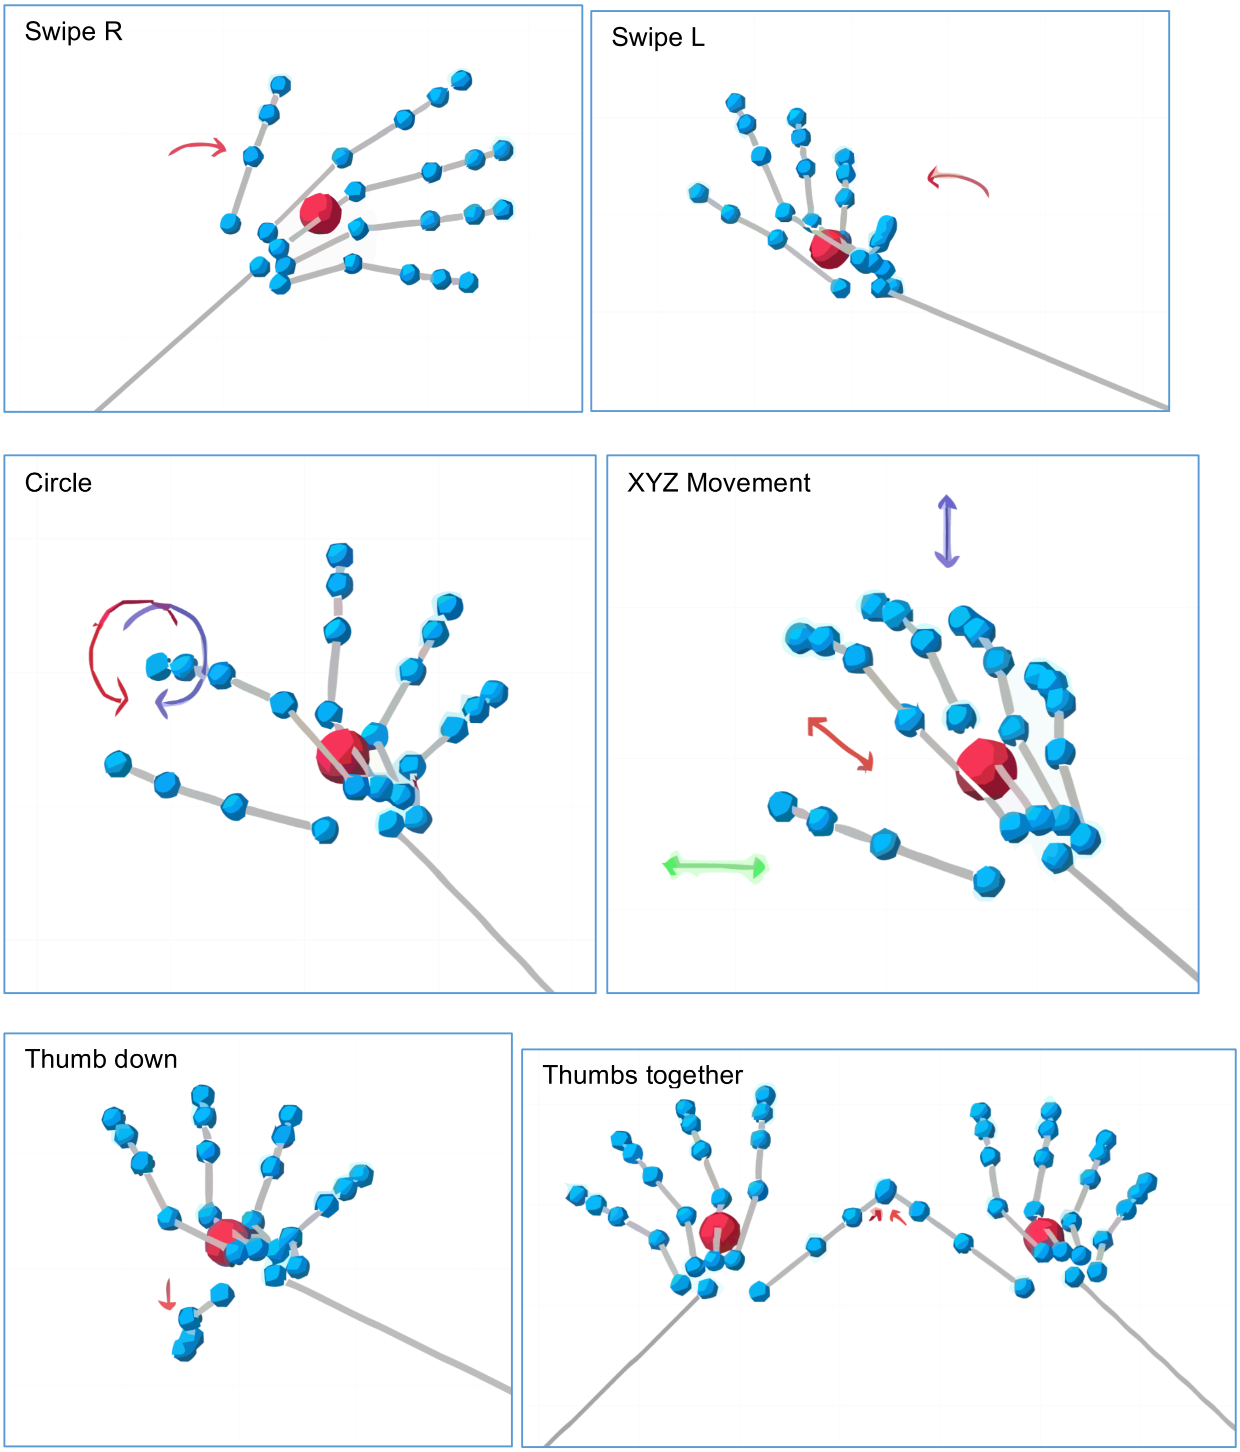
Swipe R
Swipe L
Circle
XYZ Movement
Thumb down
Thumbs together

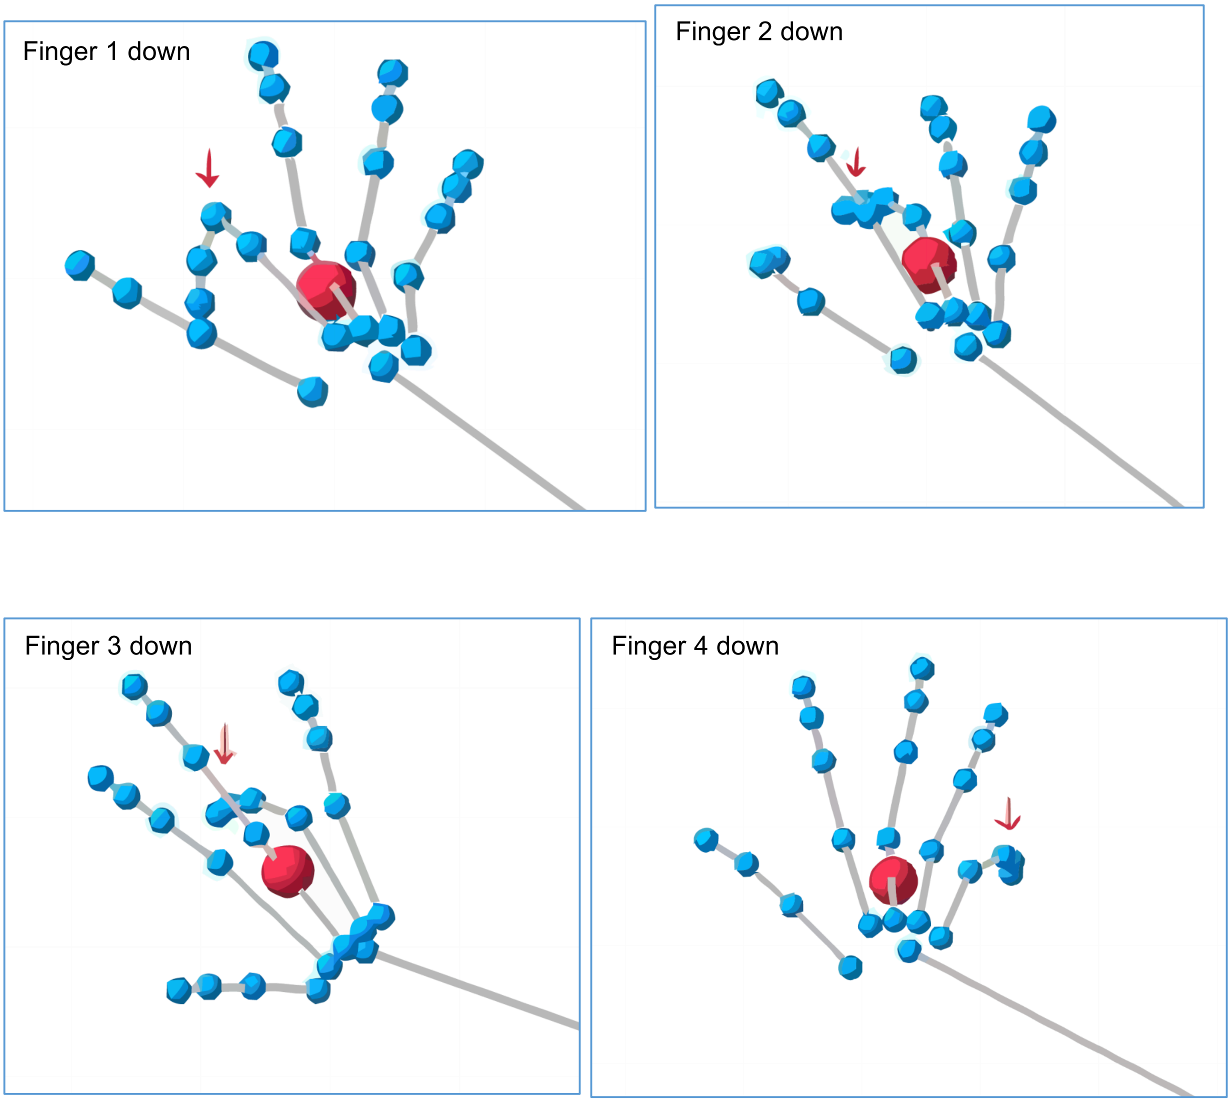

Figure 1: This shows all the gestures that “Leaps” can interpret. These gestures control different parameters depending on when they are used in the piece.

The non-tactile interface presents elements of familiarity for both the performer and their audience. The simplicity of removing apparatuses with which to interact, can create freedom not previously explored. This familiarity however, is entirely dependant on the mapping. Familiar gestures found in day-to-day computer interfacing such as "swipe", "pinch", and "turn" (see diagrams above) are used in "Leaps" to contextualise the mapping in to gestures that the average person can relate to. This is done not only for the performer to interact in a more intuitive way but for the audience's benefit as well, as they are able to recognise clearly when sounds are being produced.

The performer is presented with a navigation panel with different symbols, each representing the ability to control a specific parameter group within the sound-world. The performer can navigate between these in multiple ways but predominantly by placing both thumbs together for just under three seconds. The reaction time for this feature is a balance of my own judgement, the absolute fastest that the piece should move in terms of musical pace, and the interface's stability in terms of tracking with the Leap motion. The system's interface is as much a part of "Leaps" as the score is; therefore, the menus are arranged in order of their use within the score, with the earliest desideratum presented first. When the performer follows the score presented, the piece lends itself to fluid interfacing, as all the menus are in the correct order and there is no time spent scrolling.

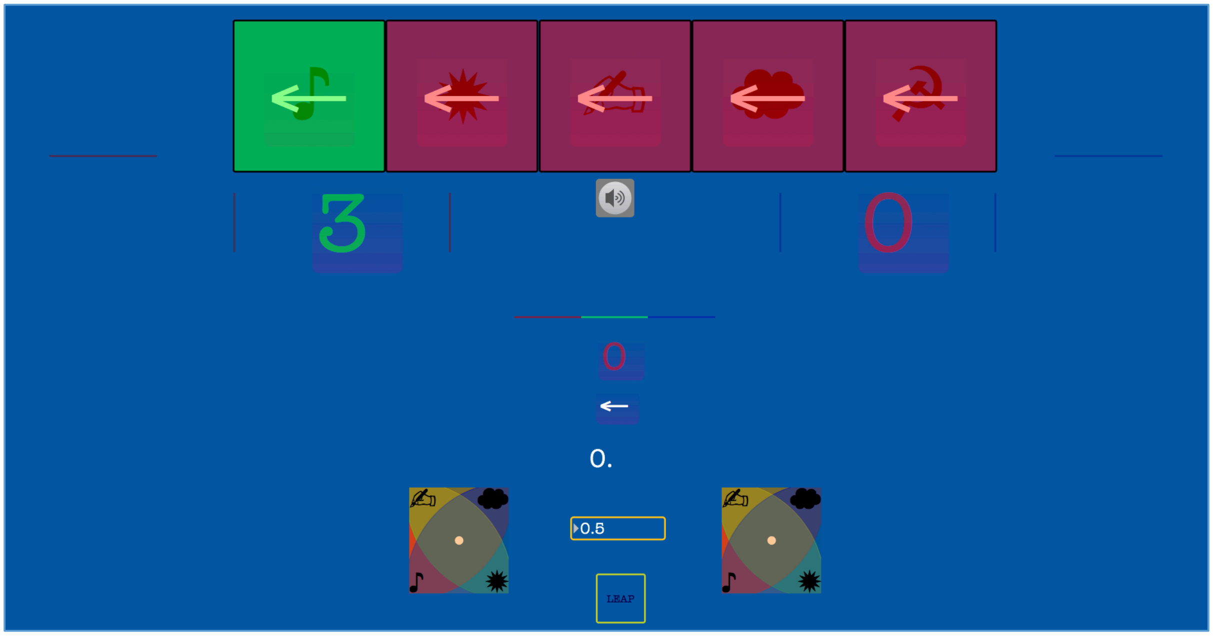

Figure 2: shows what the performer sees when they use “Leaps”.

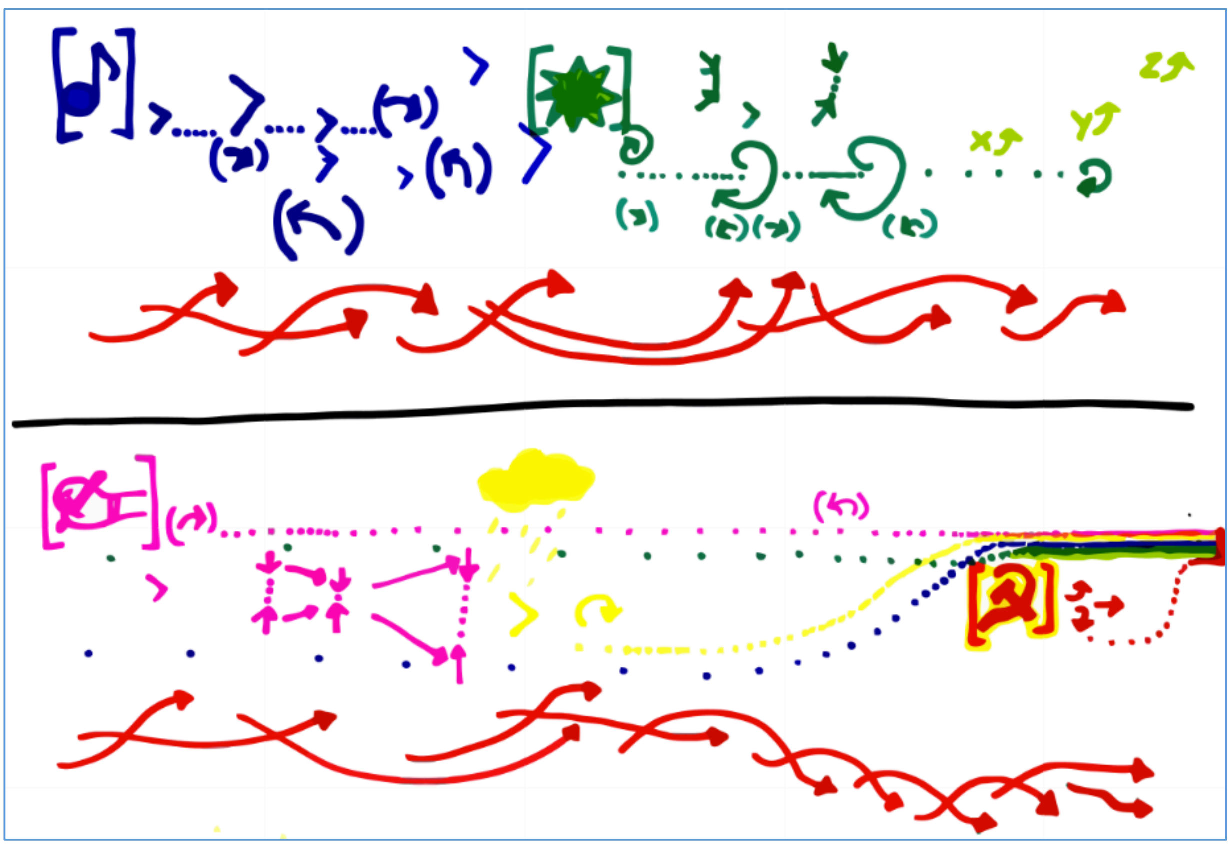

Figure 3: Shows the full score provided for “Leaps”. Gestures are graphically portrayed to give an indication of structure but the time domain is not specified. This allows the performer the freedom of improvisation within a score. Each symbol corresponds to one shown on the interface, so the user knows where to be at any given time.

As an instrument

The definition of an instrument can be interpreted in a vast variety of ways. A musical instrument is an object which we interact with in order to make noise. This simple interpretation has been expanded upon in an equally simplistic way to categorise instruments into "brass, percussion, strings and woodwind". This outdated categorisation overlooks the electronic instruments developed in the last century.

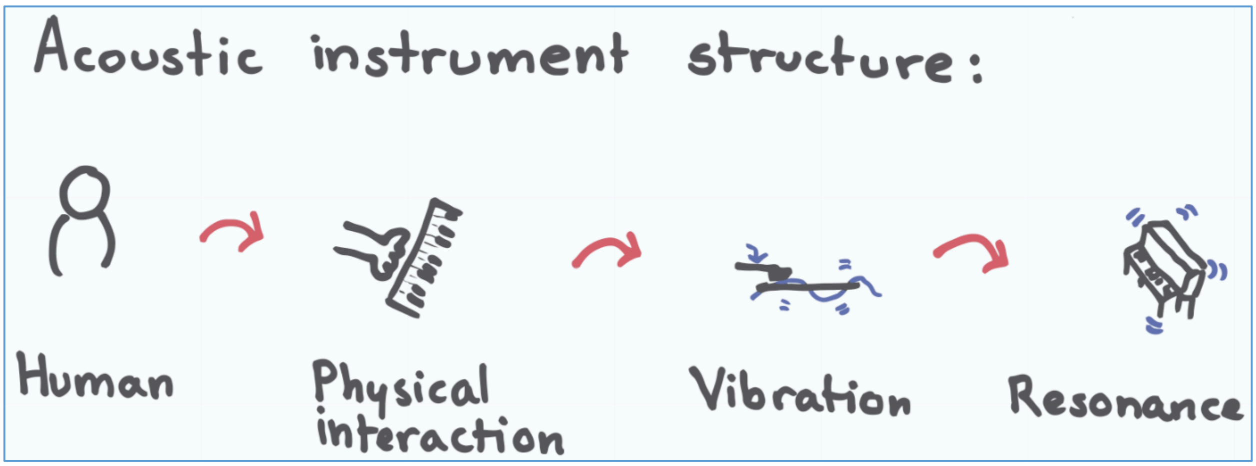

Figure 4: The fundamental structure of an acoustic instrument.

The model above is consistent amongst a vast variety of acoustic instruments. This model is somewhat disrupted by the advent of electronically-based instruments. Indeed, perhaps the most notable non-tactile instrument to mention is Léon Theremin's invention: "the Theremin" ("Leon Theremin playing his own instrument", 2008) This instrument does not fit the above model in more ways than one, the most obvious one being the lack of tactile interaction required to create sound. Furthermore, we begin to see a separation between components that trigger the sound, components that interpret these triggers, and components that output sound. A proposed model within which all of these recent developments (as well as the theremin and other early synthesisers) can be categorised, can be seen below:

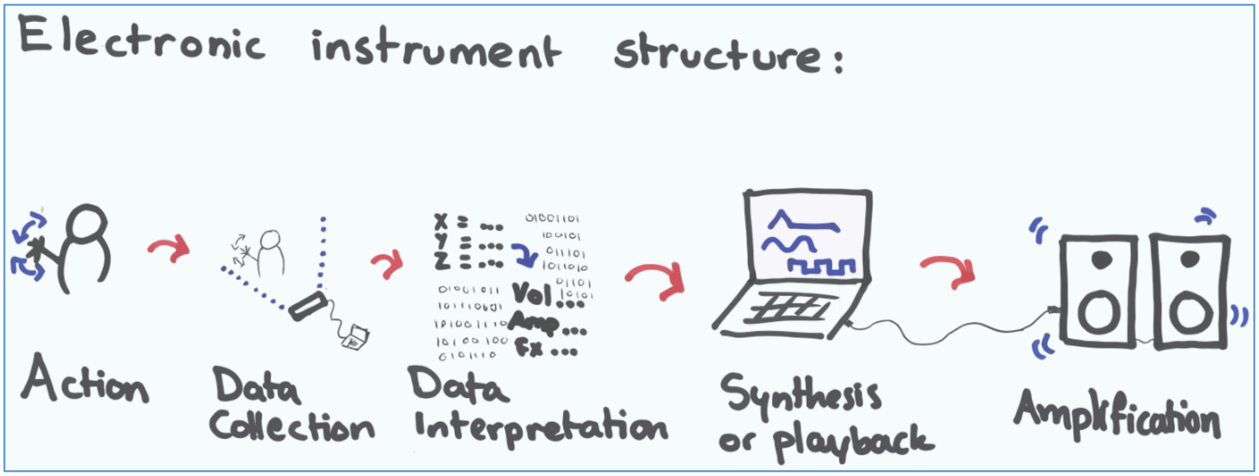

Figure 5: Shows the fundamental structure of an electronic instrument.

It is worth considering this new model alongside the old one when defining an instrument. In the case of “Leaps”, defining the interface as an instrument is somewhat problematic. Whilst an instrument can generally be seen to be a tabula rasa for a musician to utilise, only limited by the physical capabilities of the instrument, this cannot be said for “Leaps”. The interface’s primary function is to perform the score it was built to, and deviation from this score can at times be counterintuitive. In addition, the sound world created by “Leaps” is predominantly relevant to the atmosphere created by following the score provided. This leaves the “Leaps” in between the definition of an “instrument” and a “piece”. Though my intentions for “Leaps” have always been to create a standalone experience, I have maintained an open mindedness towards the integration of new systems in the future, and the possibility of “Leaps” to go on to be used as an instrument in a variety of contexts with limitless potential. Whilst it is hard to describe “Leaps” as an instrument, its applications can be re-purposed with minimal effort. The Max 6 environment makes it simple to compartmentalise individual sections of the programming. The format of interfacing and mapping can be connected to new sonic output in a relatively straightforward way, expanding the future flexibility of the interface.

Non-tactile interface

“Leaps” maintains the performer’s focus on their actions and how the computer is interpreting these actions. The aim of using metaphorical mapping is for this to be intuitive. It would be extremely confusing to present a performer with a non-tactile system in which the mapping and the sonic output bear no relationship. “Leaps” also presents the performer with operational paradigms to follow. These facilitate the learning experience and allow the user to formulate a schema of the actions required to get a desired output. These operational paradigms can be found in a lot of modern technology such as the common location of a home button at the bottom of a mobile phone, or the use of navigation menus, usually situated at the top of a page. These not only allow you to navigate without much thinking, but also constantly show where you are within the structure of the system you are interfacing with. “Leaps” delivers these operational paradigms in the form of the visual feedback when navigating through the different sections of the piece. The performer is always clearly aware of where they are in the piece. This is important, as the position of the navigational menu completely alters the sonic function of individual physical gestures.

The performer’s awareness of which gestures relate to which sounds is crucial. This awareness will come from practice but there are also systems in place to aid the user in coherently controlling the performance. The visual feedback provided enables the user to be sure of where their hands are in relation to the control parameter that they are aiming to control. Early versions of “Leaps” featured no visual feedback for the performer. Controlling parameters such as “height” or “horizontal distance from sensor” was difficult to do accurately, as establishing a definitive point was extremely impractical. The different gestures available, as seen on figure 1 above and figure 6 on pages 19 and 20, are all visually represented to the performer.

Polytextural sound

“Leaps” allows the performer to control multiple textures at once. The format of the piece allows the user to set sounds in motion and control when they stop. This means the performer can overlap these sounds to create thick, rich textures which will continue indefinitely. The performer can remove their hands from the instrument and safely assume that the sound will continue to be generated until their next intervention. To balance this polytextural mode, the performer is also given control of certain sounds which will only trigger with certain actions. An example of this is the the continuous drone sound (originally a sample of a bell) that can be heard at the beginning of the piece [Leaps video, Timestamp: 00:00:26], or the slightly out of tune glitches that can be heard when the performer moves over to the third section [Leaps video, Timestamp: 00:04:20] “Leaps” also gives the performer the ability to set in motion a continuous percussive element. This bridges the gap between the textural modes, and the one-shot modes. A diagram below shows how these modes and how the user interacts with them.

Leaps audio modes

Bells:

Pitch Ofset = Buttons

Trigger = Pinch (L)

Bass:

Start = Swipe(R)

Stop = Swipe(L)

Tempo = Pinch distance

Samples:

Trigger on/off = Swipe (L)/(R)

Start location for playback = Pinch distance(L)

Toggle loop = Pinch (R)

Length of loop = Pinch distance (R) (only within 3s of "Pinch (R)")

Pitch Ofset = Circle size & time

Speed of pitch shift = Speed of circle

Delay time = Orientation Z

Delay feedback = Orientation X

Feedback level. = Orientation Y

One-shot samples:

Trigger = Pinch (L)

Drums / Percussion:

Start = Swipe (R)

Stop = Swipe (L)

Speed = Pinch (R)

Additive:

Trigger L = Pinch (L)

Trigger R = Pinch (R)

Automation Switch = Circle speed

Stop all additive = Swipe (L)

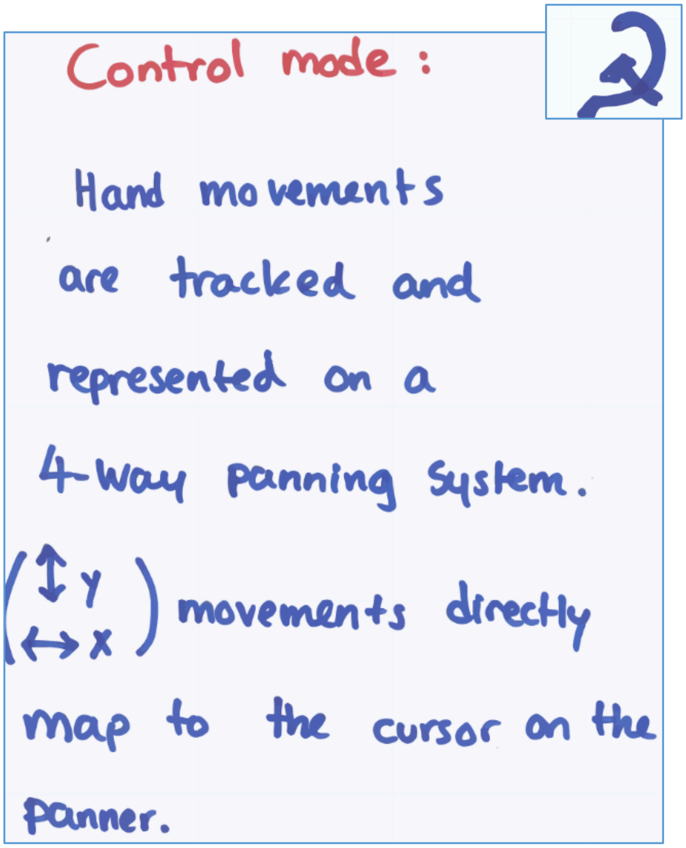

Figure 6: This shows what the modes do in relation to the symbols. It is a useful reference as it will aid in the understanding of the structure explained within the text.

The final stage of the piece sees an eight-way amplitude control for each main textural element of the piece. This is visible at (Leaps video, Timestamp: 00:12:40). This amplitude control produces very obvious gestural controls to a spectator -up is louder and down is quieter. These controls are a way of molding the texture created to a desired point before ending the piece. This final control of the sounds is fundamental in allowing the performer decide how the piece's levels are balanced. Minimal instruction is given at this point in order to emphasise the variety that comes from every performance with "Leaps". This polytextural control is something not at all found in acoustic instruments. Referring back to Figure 4, one can see that this system struggles to fit this model. This forms part of further considerations regarding the future applications of "Leaps" as an instrument.

Achieving expression

“Leaps” integrates expression controls into every aspect of the performance. Most notably, the user is almost always in control of the output volume, via a direct mapping of their left and right hand. This corresponds to the left speaker pair and the right speaker pair respectively. This control does not push the output volume too erratically, as it situated pre- reverb. This means that sporadic movements of the hand have a much more organic sound to them, rather than just sharply changing the sound all together. Other controls for expression changes include the velocity at which certain actions are carried out. For example, the faster a circular movement is done when controlling the second section [Leaps video, Timestamp: 00:02:57], the higher the pitch of the sample playback. This pitch shift is done in predetermined intervals which gives the shift an intentionally “not-smooth” transition between pitches. This fits with the rest of the sound world’s frequent glitching.

Mapping

When mapping the hand gestures and movements within “Leaps”, I focused on two main ideas:

- To enable the performer to intuitively interpret the score.
- To contextualise the audience’s experience of the piece and allowing them to make a mental link between the action and the sounds.

Performers are given much more scope to be intuitive with their gestures when they can instinctively recognise what effect they will have on their sonic output. Throughout the piece, the context of certain gestures might change. For example, swiping right might at one point initiate a bass sound, whilst at another point, it might trigger the start of a drum machine. These gestures remain the same throughout the piece in terms of how they affect specific sounds, however the part of the sound world they affect changes in different sections.

The “Leap motion” interface has a very fast response time for tracking. Generally, any lag is negligible in the context of performance, however there are certain considerations to be made in the subject of fluidity. With an interface such as “Leaps”, interpreting multiple gestures at once can be a complicated task. Time buffers are in place within every gesture so that there is no overlap when collecting data. For example, when setting up for a “pinch” gesture, the thumbs from both hands might at some point share the same XYZ coordinate. This would initiate a completely different gesture, one that does not come at the time desired. In order to avoid these miscalculations, the aforementioned time buffers create a small but significant delay in interpreting gestures. This gives the system time to work out how to best interpret the gestures being presented by the performer. Whilst these delay times might be perceived as imperfections, I see them as a part of the idiosyncrasies of the piece. With the exception of the “thumbs” gesture (in which the slightly longer wait time is also in place to break up the piece musically), generally, the time delays are in the range of 50 to 100 milliseconds and pose no real obstruction to the flow of the piece.

As mentioned, the interface is based around the very piece it was made for. When putting “Leaps” in the context of a new score, or a fully improvised performance, there are some drawbacks to consider. Firstly, the navigation is slow to use when wanting to rapidly go between sections that are not directly adjacent to each other. This would create a gap in the performance that would not necessarily be desired. This format of control can be compared to the controls of an organ’s valves. These are not always the easiest to access, and the organist must learn how to quickly access these in order to get the best out of the instrument. From the start, the piece utilises clear mapping between actions and sounds. This clarity is crucial for the audience to be aware of what is happening and retain their interest. As the piece develops, there might be times when the performer’s gestures do not directly relate to the sounds heard, or they do not seem to be affecting the sonics of the performance. This is because occasionally, the performer will be navigating or setting up for certain sounds to occur.

The structure

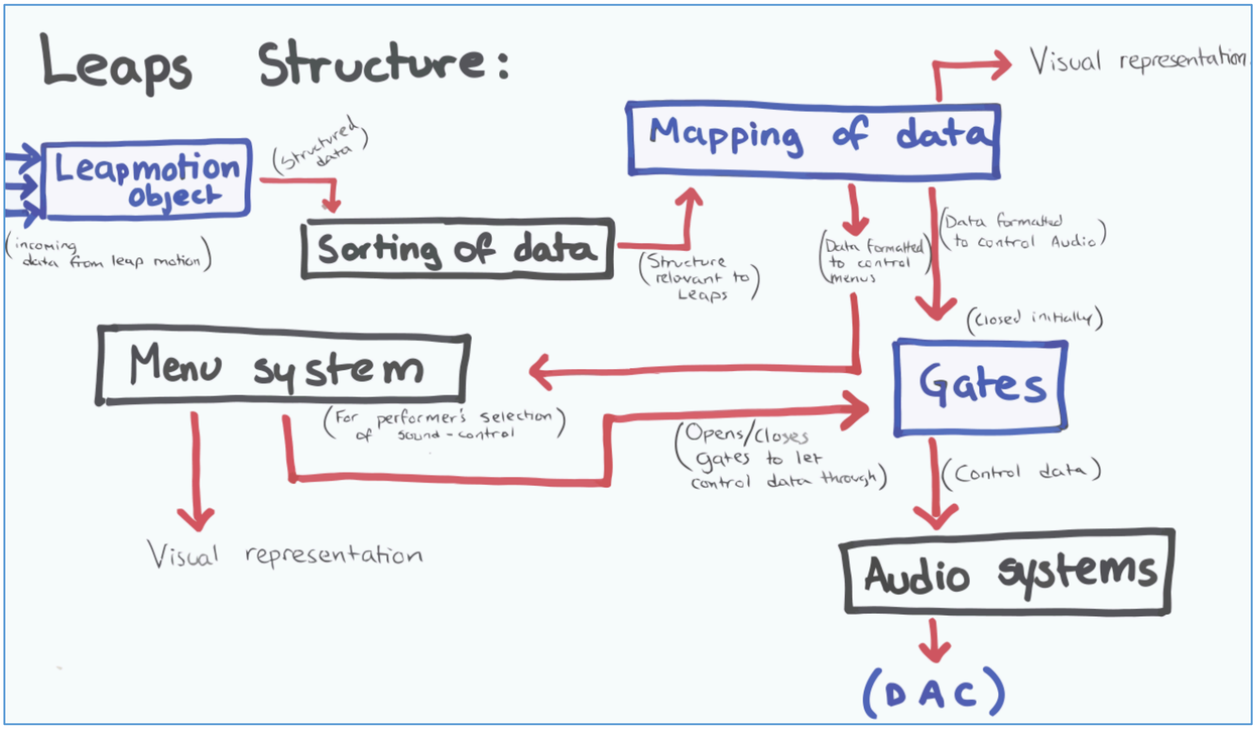

Figure 7: Shows a simple representation of how the programming is structured. It shows in the simplest of ways, how gates are opened and closed to vary the sonic output.

Sonic variety

When performing with “Leaps”, one will always hear a slightly varied output every new time. This comes down to three factors:

- The user’s understandable lack of ability to recreate the same exact hand positioning twice. This is because the performer lacks a reference point.

- The random seed numbers used to generate the sound world. This emulates the subtle intricacies of an acoustic instrument played in varied conditions.

- The seeds based on detailed user interaction. These are a way of creating variation based on user actions. In this case, these are more noticeable variations that directly relate to what the user is doing.

In terms of performance, this means that of course, the piece will be different every time. These three factors combined make for a drastically varied performance every rendition, as can be seen within the videos submitted. The desired effect relies on the fact that one never quite knows what the piece is going to produce. These pieces as a whole are about the lack of familiarity one has with them. This is highlighted not only by the piece but by the way the interfaces are built to respond. Because of this, the performer and the audience will share a certain perspective, as they shall not always know exactly what to expect. This ambiguity also encourages improvisation and innovative interpretation of the score. The aspect of variety does however, have implications when defining the piece's interface as a potential instrument. It is hard to justify an interface as an instrument when the sonic output is never fully certain.

Technical overview

Getting data from the "Leap Motion" hardware was facilitated by the "Leapmotion Max external" (Françoise, 2014). The external takes in a regular bang input, and outputs all position data at the time of the bang. This data comes in a format that is easy to route within Max 6. The object allows Max 6 to query a certain parameter and to output the parameter's multiple values, usually in the form of X, Y and Z values. Once Max has these raw values, the data can be interpreted to match the relevant gesture. Data interpretation can be simple, or it can lead to some complicated patching, depending on the desired output.

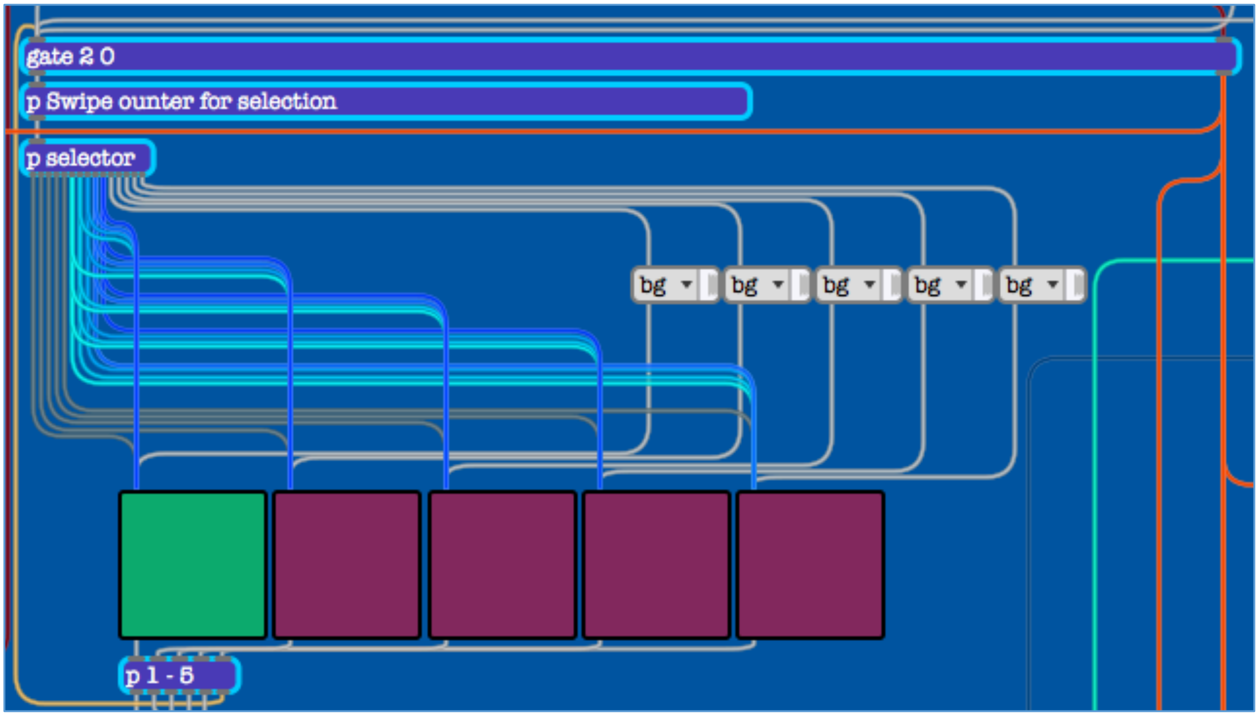

Figure 8: Shows the menu system within the programming interface alongside some UI elements.

The menu system is the performer's main method of navigating through their different sound options. This menu system is triggered by the "thumbs" gesture within figure 1. The menu swipe action activates after just under three seconds of holding the "thumbs" gesture. This then gives the user the ability to swipe to a different control section. If the user has landed on the section they want to remain in, the system will lock into that position after roughly a second of swipe inactivity. The whole process is contextualised for the performer with a colour system: red shows inactivity, orange shows current selection, and green shows the menu system in its locked state.

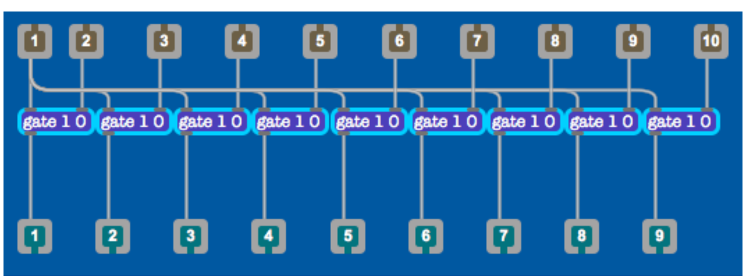

Figure 9: Shows a series of gates used within leaps.

The gate system routes the control data to different audio systems, depending on what section is active in the menu system. These systems are intentionally separated so that they can be expanded upon when “Leaps” undergoes further development. The gates system only allows for one audio system to receive data at any given time. This means that there is no confusion as to what system is currently under control.

Bells

This system consists of a recording I made at Wakefield cathedral of one of their bells ringing. This sound is convolved with a sine wave and filtered as shown below.

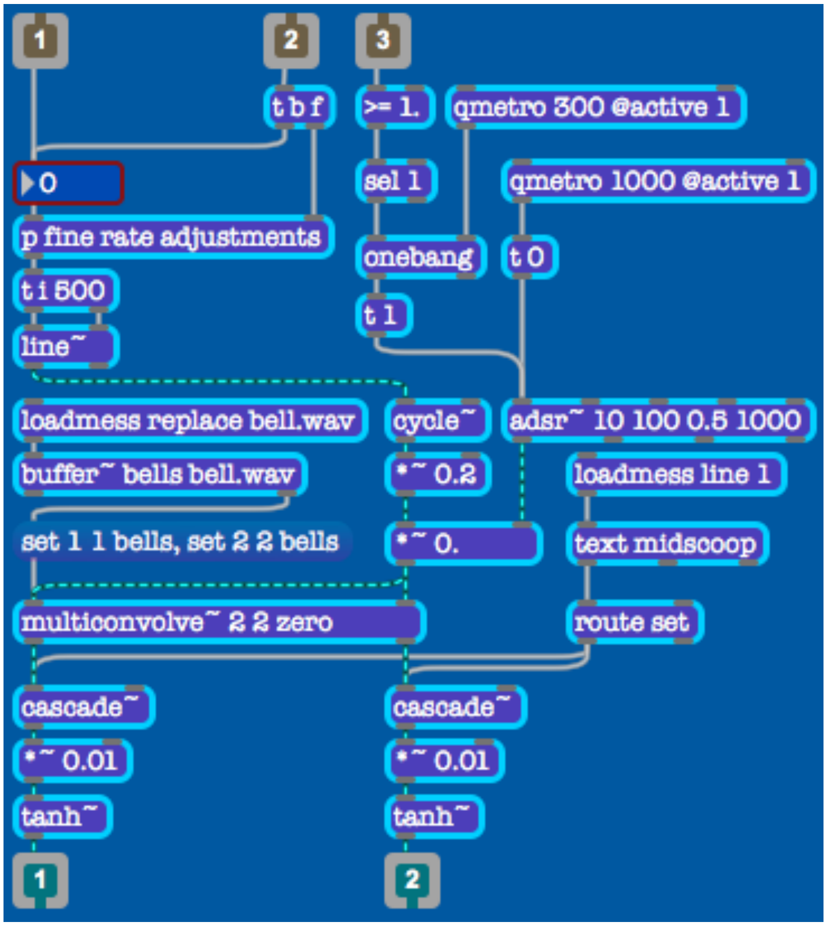

Figure 10: Shows processing done within the Bells sub-patch.

The sound takes in the control parameter “pinch” which outputs a different pitch value depending on the motion of the individual “finger tracking” parameter. The resulting sonic experience is an electronic bell sound with a very soft attack and decay. This is accentuated by the master reverb on the main output. The reverb used on this, and indeed also within the “Leaps” patching is Alexander Harker’s and Pierre Alexandre Tremblay’s “HISSTools convolution reverb. (Harker & Tremblay, 2016) “Bells” also has the ability to control the “Bass” system. These two systems work well together in terms of the contrast they establish; therefore, I thought it best to allow the user to operate the controls of these two systems simultaneously. The “bass” system is basically an additive mono synthesiser that incorporates the sound of pink noise, and clicks at the start to make the output more organic. The synthesiser emulates the sound of heavy wind, or the sea when played back in the lower frequency range, and starts to embody a more tonal sound when pushed up to the higher end of its frequency range. The system is triggered by a definitive swipe of the hand towards the right of the performer. This is easily reversed by swiping in the opposite direction. This action is instantly recognisable. Furthermore, the attack of the “bass” system bears a strong contrast to the slow attack “bells” system that it is paired with. This highlights the drama of the “bass” sound even more when it is introduced.

Samples

This system is a sampler with various controls. It features the ability to commence playback at the command of a “swipe” gesture, which is mapped in the same way as the “bass” system.

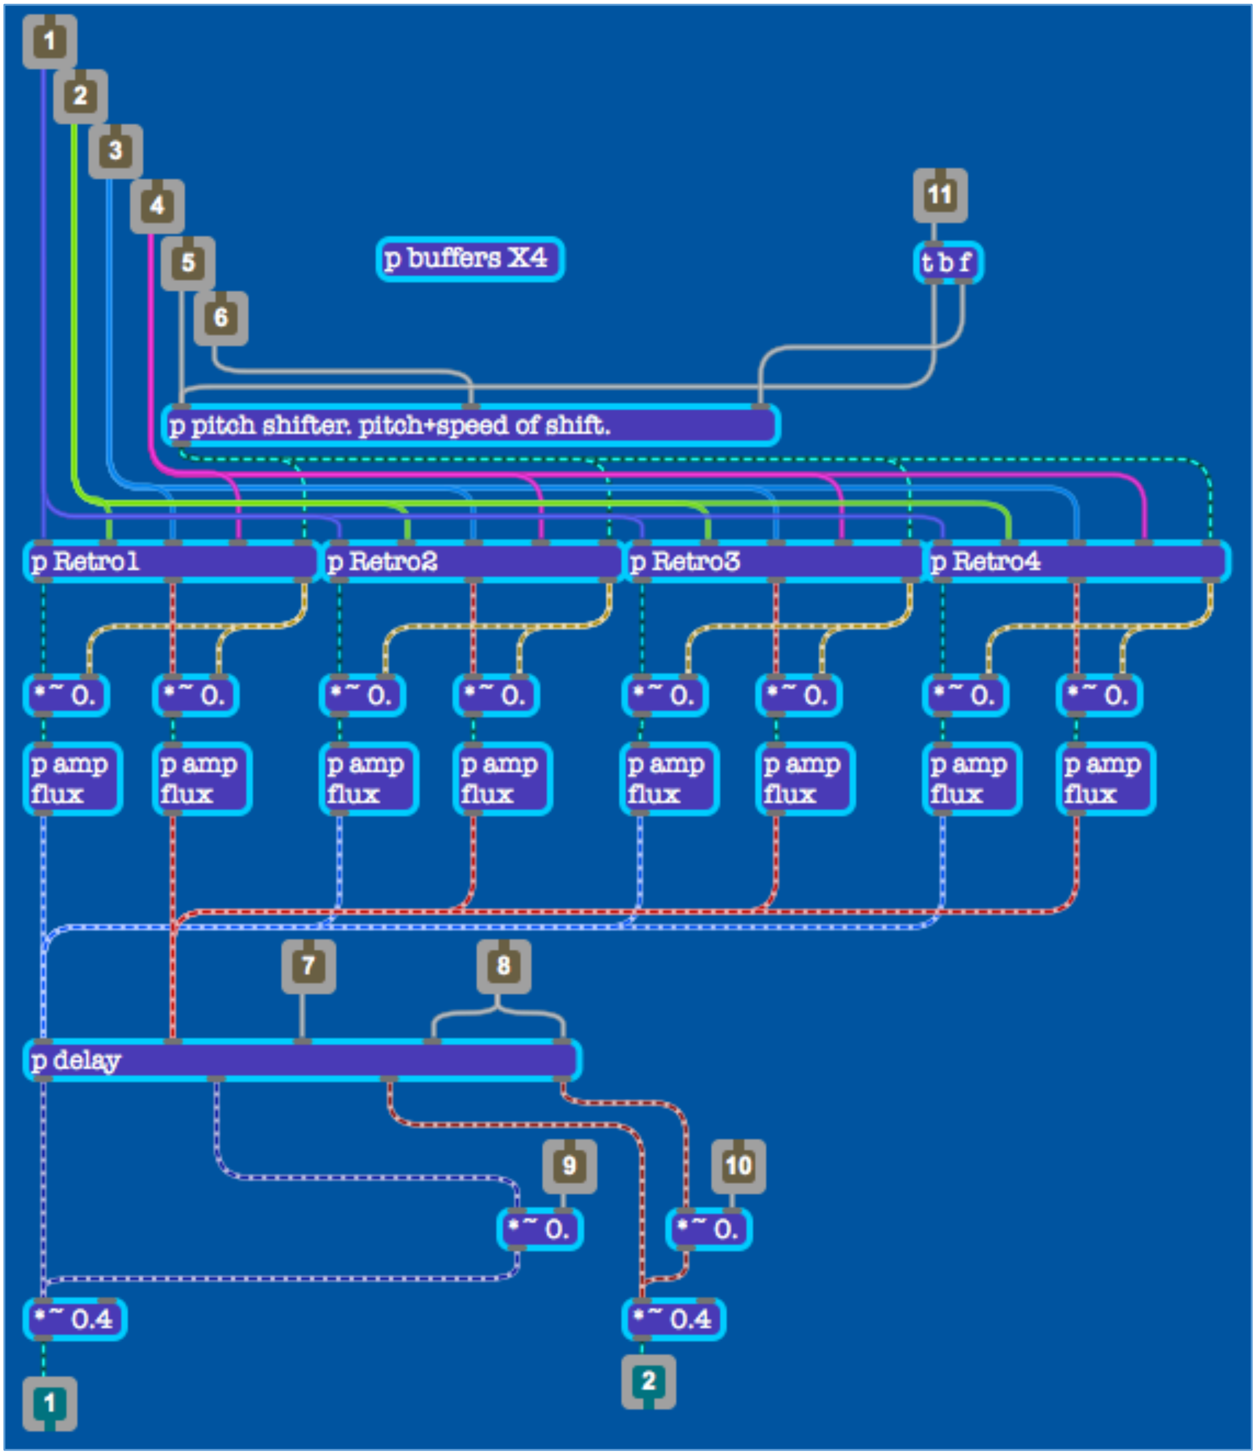

Figure 11: A look inside the retro leaps sampling structure.

The system consists of four potential samples which can be loaded in any of 28 different combinations within four possible (identically controlled) samplers. Once these samples are played, they are passed through an amplitude controller that is constantly fluctuating, creating dynamic movement between the four samples. These samples are never at rest. The constant movement between these 28 possible combinations creates scope for the system to deliver a varied and unique performance. The most prominent feature of the "samples" system is the pitch shift that can be achieved by the "rotation" gesture. The speed of the rotation directly correlates to the time it takes the pitch shift to reach its desired destination. The duration the performer does this gesture relates to the pitch at the end destination of that ramp. This is a metaphorical gesture designed to replicate the winding action that one would see in the playback of a tape machine. The system is able to create glitch-like effects such as skipping and looping. One intuitive function that this system uses is the ability to loop when performing the "pinch" gesture with the right hand. The user is then given 1000 milliseconds to decide what the length of this loop should be by gesturing with their index finger and thumb, as if they were gesturing a measurement of length to another human. Whilst this mapping will be obvious to the performer, it will not necessarily be immediately understood by an audience. This is because the mapping relates to something that the audience does not expect to happen. Throughout the piece, I attempt to incorporate a healthy mix of both of these types of mapping.

Automation

The automation system is intended as a single-shot way to trigger a sound. The sound in question is built up of a sample being simultaneously manipulated in four different ways. These four methods are:

- Frequency modulation
- Pitch shift
- Stutter effect
- Amplitude modulation

These four processes are simultaneously played-back when the sound is triggered, and a range of amplitude modulation presets are activated. There are twenty possible outcomes, each controlling the stereo output of the four manipulation methods. That is eight amplitude controls working simultaneously when the sample is triggered. The playback of this system, whilst unpredictable, is quite recognisable when performed. This is because it is of a similar nature to the “bells” system heard at the beginning of the piece; however, it features a much faster attack time which directly contrast the bells heard at the beginning. This quick-attack sound is emphasised by the ability to simultaneously control a drum machine system within this section, which also has quite fast-attack sounds. This system features the ability for the performer to control the speed of the drum playback with the “pinch” gesture described in samples, and the ability to start and stop the playback using the familiar “swipe” action. This drum machine develops continuously throughout the piece, adding and removing triggers at various points. This means that the output will be vastly unpredictable whilst the structure of the performance remains under the control of the performer.

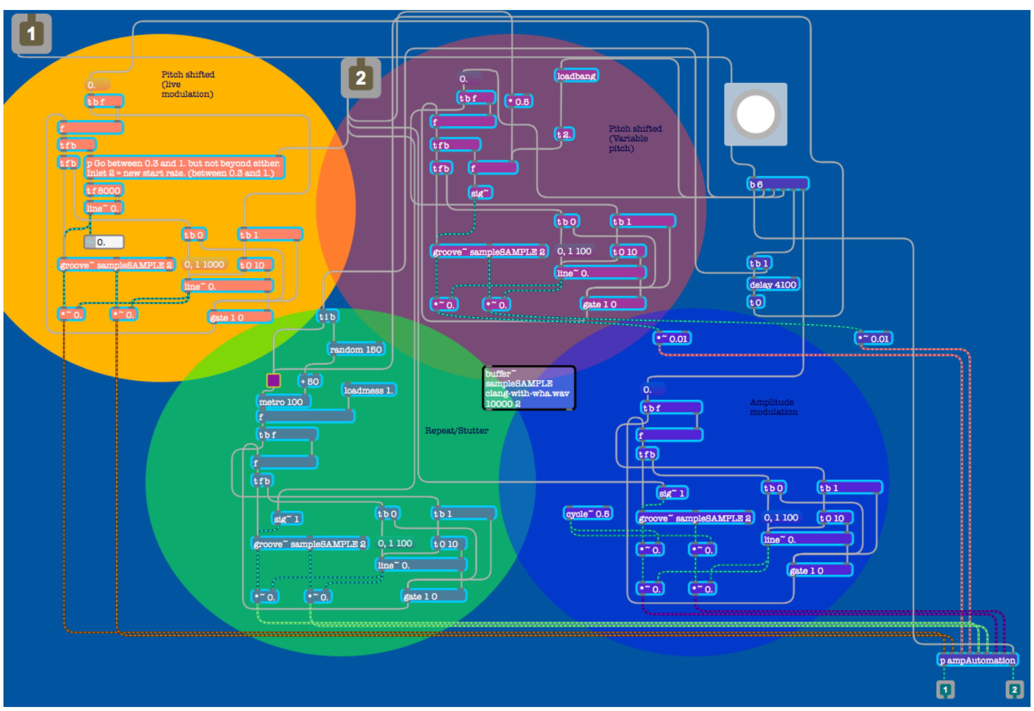

Figure 12: A look inside the automation control patcher.

Waves

This section features an additive synthesiser that continuously outputs various harmonics (of varied wave-shapes) of a fundamental frequency provided. These various harmonics pass through an amplitude modulation system that continuously alters the volume of each individual harmonic. The resulting sound creates a haunting bed of constantly moving sounds. This system changes the direction of the piece dramatically when it is introduced towards the end of the piece. Up until this point, the piece is polytextural yet minimalistic in nature. This drone of movement rushes the pace of the piece as it nears completion. It creates a recognisable sound similar to that which have been heard before, to enter the final section of the composition.

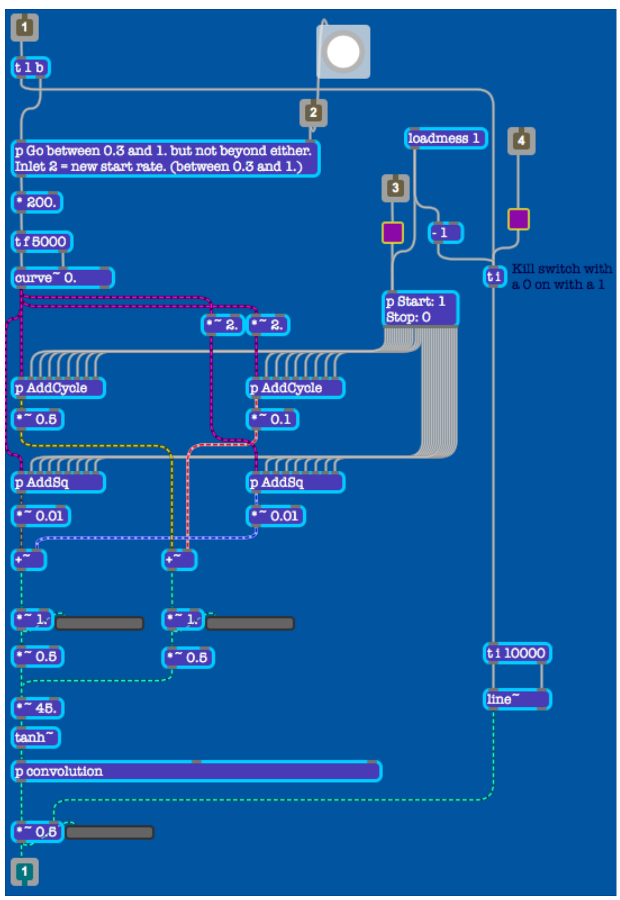

Figure 13: This shows the patching within the "waves" section of the piece.

The mixer

In addition to these systems, there is a mixer that features at the end of the piece. This system consists of two four-way mixers that can be controlled using the height and horizontal data from the user's palm location. This creates an instantly intuitive mapping that requires very little explanation in practice.

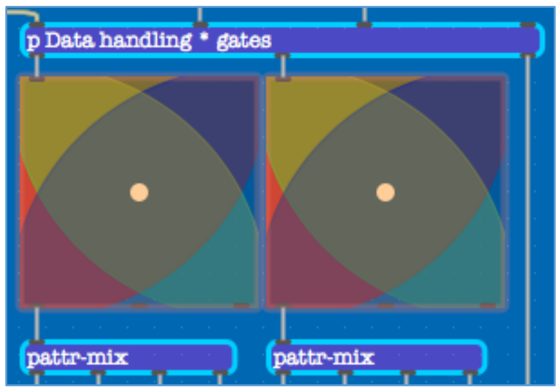

Figure 14: This shows the mixing interface which controls multiple faders.

The master pitch control

The entire sound world derives its pitch from a common master pitch that keeps the various sound systems from creating clashing harmonies that would contradict the intended character of the piece. Whilst all systems share the same seed value, all systems also have the ability to deviate within specified intervals. The system allows the piece to deviate from pleasant harmonic relationships and comfortably return to them in a controlled way.

Chess

The interface

“Chess” is a performance piece intended to be played live by two performers. It is unconventional in several respects. Both performers have equal control over the sound world. This means that neither performer is ever fully in control of the sonic output and in fact, the only way to gain predominant control is to win the game. The piece collects information throughout the game and interprets chess moves into control data for the sound systems in the piece.

The piece is performed on a normal chessboard by two performers. Simultaneously, a third person inputs the moves as they happen into the Max 6 environment. Other ways of interfacing include the ability to play the game directly on the computer screen. This format of control is far less visually dramatic and is therefore not used in the final videos submitted within this project. One potential format of control would be a purpose-built hardware interface. This would incorporate the visual aspect of using a real chess board, whilst at the same time retaining the time-accuracy benefits of the computer version. The main challenge for a performer when faced with “Chess” is accepting that the control of the sonic output will not be completely theirs. Apart from the other performer interfacing with the same system, they will also be met with the unpredictability of the system itself. A performer has to realise that their movements will not directly influence specific parameters at any given time, but rather a whole host of multiple parameters simultaneously.

“Chess” can be seen as a piece that takes raw data - with no obvious links to music, and converts it to a musical process. Non computer-based versions of this concept include Joe Hamlen’s “Shipping” (Hamlen, 2014). In this piece, the performer interprets the live shipping forecast based on an existing connection between geographical coordinates and the chords used. The piece therefore, is different every time depending on the weather. In “Chess” the performers are fully aware that their actions are directly resulting in sonic output; however, they must still play the game to win. The fact that both players are instructed to win is the first hint towards the fact that this will piece will comment on a conflict. Both performers are instructed to “perform” better than the other. Inevitably, one will always outdo the other. This conflict is emphasised sonically throughout the piece. The expression in the piece is not generated from direct user interaction (like a guitar and a whammy bar) but rather it is inferred from the data provided. This tension comes in the form of sharp amplitude modulation and heavy compression that “glitches” due to its extremely fast attack.

Intent

The intent of this piece is to create an atmosphere of conflict and suspense. Uncomfortable sounds are juxtaposed with clear and sharp melodies. These melodies are soon masked by the much louder tense atmospheric drones that eventually takeover in volume. The melodies heard at the beginning of the piece are put through more and more sonically intrusive processing until the tense atmosphere is undeniably obvious. This signals that the players are nearing the end of the game. As the game develops, the listener will experience a much harsher and tense sonic experience. This loose structure is generally followed by most games; however, the open interpretation of “Chess” makes every new performance a strikingly new experience. The aim is to create a piece that whilst distinct in its sonic output, also contains subtle variations that are experienced within conventional instrument. These subtleties are used to create drama and illustrate the progression of the game. Distinct moments of disorder and chaos can be heard at times, and it is the intention that these reflect the complexity of the game. At the same time, one will also find moments of rest and clarity, in which the pace of the music slows down or stops to reflect the pace of the game.

As an instrument

To consider this interface an instrument would not fit a lot of the criteria already proposed above. The main conflict is the lack of human interaction with any aspect of the sound world. The interface is merely a format for collecting data and a means to interpret this data into structured sound. It is therefore, more interesting to discuss the potential this system has as an instrument. A discussion of this can be found within the conclusion section.

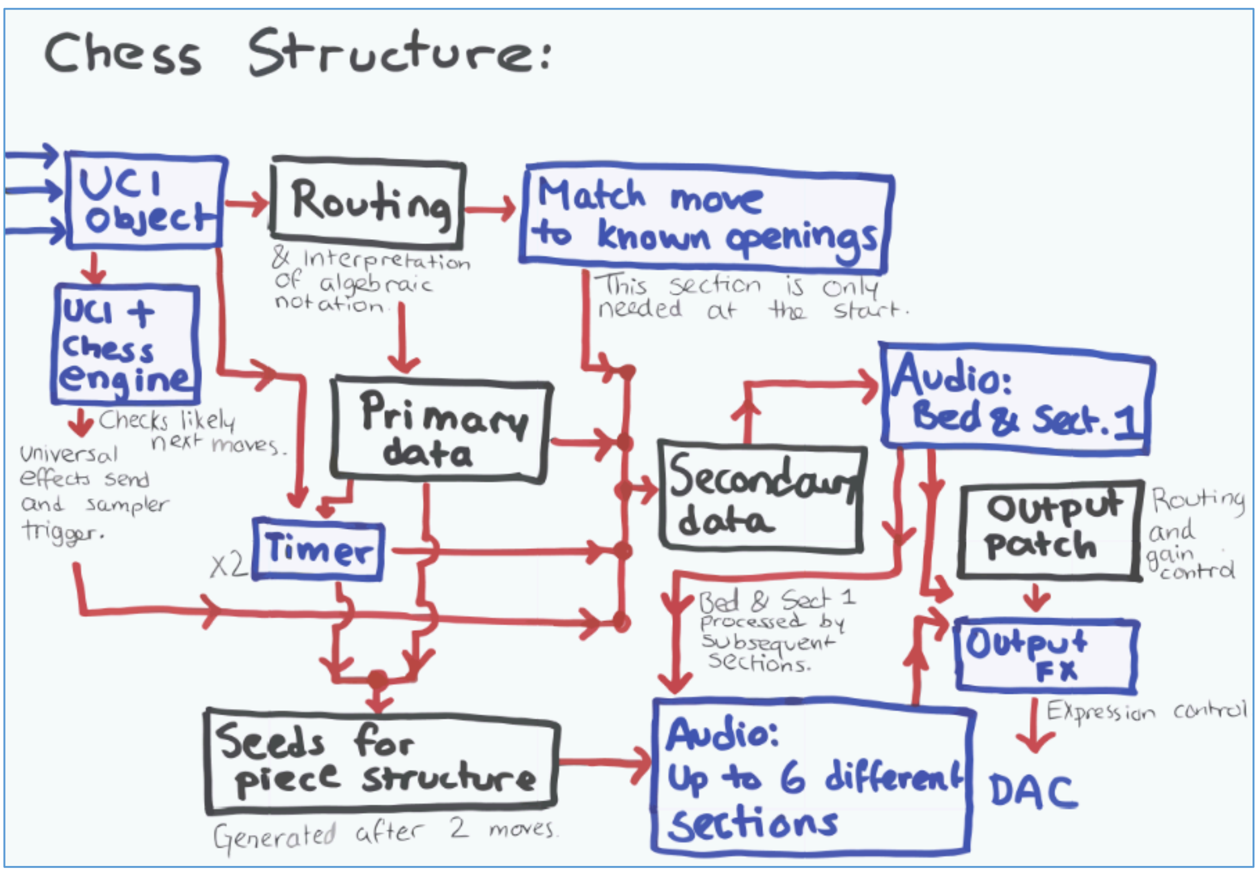

Figure 15: This shows a basic structure of the programming behind chess.

Allowing variance and interaction within the piece.

The patching for “Chess” features a somewhat linear methodology in the way sounds are produced. The piece is organised into eight sections which are introduced one after another with a maximum overlap of three at any given time. This patching is the only fixed form of structure within the piece but even so, these eight sections can come in any order, depending on how the game is played. As can be seen in the submitted videos, the variance between performances is highly evident. This variance is down to two main factors:

- Instrument intricacies as discussed for “Leaps”.
- Gameplay differences

Ordinary acoustic instruments feature subtle intricacies that define them. Detailed acoustic qualities such as the effect temperature in a room has on instruments, or the sympathetic resonance of the strings on a piano, all contribute towards a unique sonic experience with every use of an instrument. “Chess” emulates these intricacies with subtle changes in its sonic output. Variances with every rendition will always be found regardless of the gameplay. These variances can be subtle, but also can completely redefine the opening of the game. It is the intent of the piece to steer the performance in the direction of conflict, regardless of the initial sounds that occur.

The structure of the performance is taken in a predefined direction at the pace of the gameplay. The interaction between the players and the derived intent behind their moves is analysed and mapped to characteristic changes in the piece. Landmark actions such as "the first piece taken" are marked by significant events. In the case of "first piece taken", the playback of a player theme is generated and triggered. These developments are not instantaneous in the way that a

conventional interface would react to input, but rather quite progressive and slow in its playback. This way of introducing new elements differs to the methods previously mentioned in "Leaps", where the mapping is directly and instantaneously correlated to the actions being performed.
In terms of performance, this way of deriving structure gives the individual performers almost no certainty over their sonic output. The intent behind introducing this unpredictability is to emphasise the lack of individual control the performers are given. Indeed, whilst the piece would not exist without the performers, their involvement merely drives the composition forwards. The indeterminate nature places the performers in a similar perspective to that of the audience.
"Chess" presents no continuous expressive control in the way a traditional instrument does. It does however, uses derived data to interpret the intent of both performers to create this expression. To an extent these aspects are simulated, as there is no direct human interaction creating said expression. Instead, the system derives control data for expression from the data collected throughout the game. These considerations contribute the idea that “Chess” is not currently definable as an instrument.

Real time and offline models

"Chess" could theoretically function as an offline model, in which the performance is fed into the system as a seed for an instantaneous sonic output. The model would almost push the concept of "Chess" from being a performance piece to a playback piece or even a standalone installation with numerous algebraic notation format games loaded in its memory. Whilst of course, this drastically changes the piece, the more prominent reason using a real time model is the aspect of time. The time parameter adds a lot to the quality of data there is to interpret, and subsequently adds to the performance aspect. The system’s ability to portray tension between moves and evolve over time is invaluable to the desired effect.

Performance

The main conflict facing a performer is the idea that they are not indeed performing a piece but playing a game of chess. The intent of the piece is made clear to them in the score (shown below in figure 16) where the performers are presented with the following instructions: “Play chess. Make sure you win.”. This not only reinforces the intent of the piece (which is to create an atmosphere of conflict) but also clarifies what the performer is there to do. For example, if the performer playing as black sees a way of defeating the opponent playing as white, he or she must follow through the move they think is best, even if they believe it would be sonically beneficial for the piece for them to allow the opponent playing as white to have a chance. This removes the ambiguity for the performers but also enables the piece to have a clear direction which will always depend on the individual ability of the players.

A consideration to be made is the duration of the game of chess. It is very difficult to determine the length of a game of chess as there are many variables such as the experience of individual players or external pressures such as spectators. For some chess games, clocks are used to keep the game under a certain length of time. Each player’s clock is only active during their turn and there Is a finite limit on the time they can spend thinking before they move. This system is not enforced within this piece. This is potentially a choice that was faced in the past, when John Cage and Teeny Duchamp played as part of John Cage’s performance utilising the Chess set. It is said that the game they played went long into the night until players decided to call it a draw. This was not before plenty of the audience had left. In this situation would have been beneficial for a time limit to have been established. This would have perhaps maintained interest of the audience ("Marcel Duchamp & John Cage Play Musical Chess", 2016). Within this piece, nothing has been done to prevent this very same issue from coming up. Whilst that never has come up, it is definitely noteworthy of future development. Beyond the scope of this project it would be ideal to incorporate the clock system mentioned prior. This system would be useful not only to keep the game limited to reasonable time limit, but also to trigger events if a dedicated chess interface is developed.

Figure 16: This shows the only score/performance instructions that are provided for the performers.

Technical overview

Collecting data.

“Chess” collects all the data it requires from one object by Jeremy Bernstein (Bernstein, 2016) called “UCI”. This object presents chess moves in the form of algebraic notation. Algebraic notation is the standard for human-side notation of chess moves. It consists of piece data followed by location data.

For example:

Nb1c3 = Knight has moved from b1 to c3.

This algebraic notation is followed by the turn number. From this data, I can interpret the precise location of the piece, and the colour of it. As the game goes on, this data builds up a statistical understanding of what is happening within the game. It is this data that supports the entire infrastructure of mapping that follows it. From this information, the following data can be extracted:

- Coordinates for all pieces
- Current active piece
- Turn data
- Taken pieces
- King's check status
- Current count for all individual pieces
- Game score
- Check-mate

This data is what will be referred to as "primary data", as it is the most direct data that can be gathered directly from the object with the use of simple mapping. Whilst this data is useful, it does not provide an overall picture of the game being played. For this, more complex, "secondary data" must be derived. In addition to the "UCI" object, "Chess" also collects time data. Strategically-placed timers will start and stop depending on when a move is made, in order to collect the "think time" parameter. This "think time" parameter is used alongside the primary data taken from "UCI", to create the following secondary dataset:

- Player's average thinking time
- Player's forward thought
- Player's aggressiveness
- Player's intent
- Accumulative "taken" count
- Accumulative "check" count

Opening

“Chess” uses the first two moves to build the structure for a given performance. It is often said that the opening moves are the most important to a game of chess, and the piece mimics this concept by determining a structure based directly on the first two moves. The player’s opening moves are repeatedly compared against eleven of the best known chess openings. The percentage of similarity is noted after every move. This running total forms the “openings” parameter data. “Chess” also incorporates a game engine that reacts to every move ("Stockfish", 2016). This game engine will predict the best move at every stage of the game. If this “best move” is played, then the parameter will output an “on” value for a recorder patch. This recorder will sample the current output and play it back with a slight delay. This signifies that the player’s intent is known and from an outsider’s perspective, they are formulating a structured plan.

Audio systems

Introduction and section A

The piece begins with an additive synthesiser playing tones which consists of varying harmonics. These are automated up in amplitude at varied intervals to create an ever-changing texture. The parameters of this processing can drastically change the way the synthesiser comes across sonically. This processed sound forms a bed which will be present throughout the piece. “Section A” uses a theme-generating patch to create a player theme that follows the player throughout the performance. This theme will be different each performance; however, the notes that form it will always be within a (procedurally selected) musical mode. This harmonic relationship is not necessarily the same for both white and black, as this emphasises the conflict found between the two players.

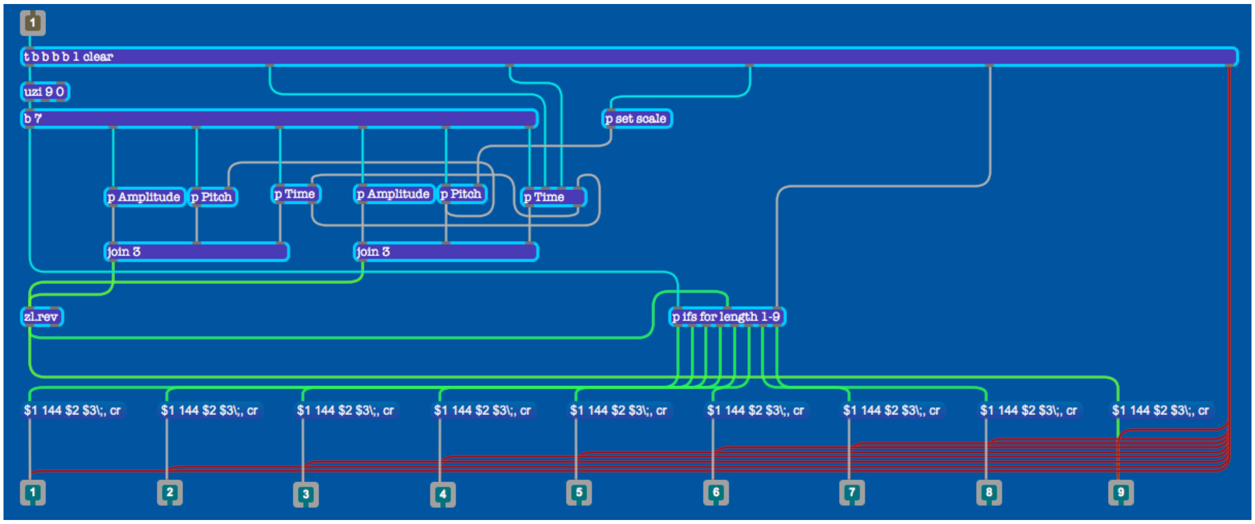

Figure 17: Shows the melody generation patch which will output different MIDI sequences every time it is played based on variations in parameters determined by early chess moves.

Subsequent sections

After section A, the game's structure is followed based on the opening of the game. The "seed moves" create a format for the rest of the game, and outline the order of appearance of the following sections:

Section B:

This section samples the output of section A and plays it back in a "glitch" fashion. This is done for landmark negative events for the player such as a piece taken or repetitive check. The intensity of this is increased by the waiting time of a player.

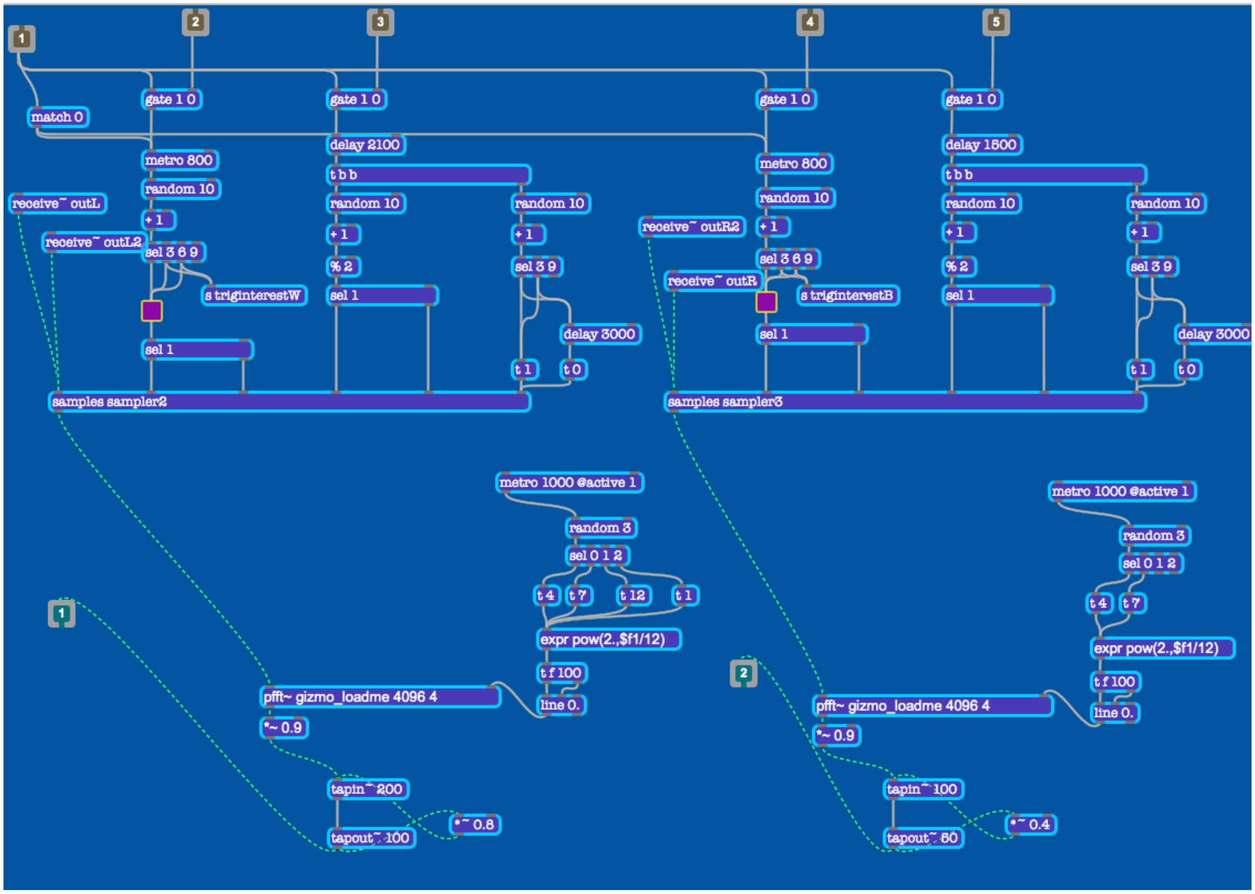

Figure 18: This shows the patcher that is within section B.

Section C introduces a continuous tone filtered to remain un-intrusive. This adds to the existing bed and is manipulated by any subsequent sections along with the opening bed sounds and Section A. This section is triggered as soon as the take count reaches a certain point. Because this point varies depending on the opening seed, it cannot be definitively determined.

Section D marks the introduction of an evolving drum machine. This drum machine will always continue to change. The pace of this change, as well as the pace of the playback, is directly correlated to the tension of the game. That is to say, the more tension, the faster the drum machine will change. It is capped to never go too fast as this would end up sounding far too ridiculous if left unchecked.

Section M8

This section creates gaps in an otherwise continuous bed of sound. These gaps can be rhythmic in nature. This section derives its data from percentage correlation to famous chess moves.

Section CH emulates the sound of a “side chain” in which the low-frequency kick sound also tigers a momentary dip in the master audio. This dip is subtle at first but increases with intensity along-side the game to the point where the sound can even momentarily cut out completely.

CUT is the last section and as the name suggests, it will “cut” out the sound for a short period of time. Generally, this occurs near the end of the piece, as it will always be placed as the last section regardless of what order of sections is generated. If it does occur, it still won’t activate unless the threshold is passed in the number of pieces that have been taken within the game.

These sections are independent sound events that will trigger at predetermined intervals based on the progress of the game. At the point of check-mate, the piece fades from whatever state the game is in, to a state of silence. This silence signifies the end of the piece and indeed, the end of the game. The output system is a segregated patching section for a two reasons. Firstly, it enables easy control of all spatialisation aspects of the piece, and secondly, because it enables global parameter adjustments and effects to be done. Within this output section, processing of existing sound content is done, such as the expression and amplitude modulation mentioned earlier.

Conclusion

This project was about creating pieces in which the interfacing style was different and innovative. I have attempted to create controls for music that can be universally interacted with by anyone. This was done through the metaphorical representation of input data as sound. These new ways to interact have posed a variety of challenges in their programming, mapping, compositionally, and in practice. The aim was to create a performance that not only showcases a coherent composition within these new interfaces, but also creates an atmosphere through expression and musicality. These interfaces aimed to be unpredictable and varied in their detailed features but to follow the control of the performer at every step. Finally, these pieces aimed to be versatile enough to present the opportunity to be developed upon in the future. “Leaps” is primarily striking because of the lack of tactile interaction. This way of control opens opportunities and creates drawbacks simultaneously. The piece works straight away for a performer without the same level of practice as would be required form a conventional instrument. Whilst it is not strictly true to say that no musicality is required in order to perform “Leaps”, little to no understanding of musical theory is explicitly necessary. Indeed, it has been preferable to see the piece in the hands of non-musicians. “Leaps” continues to deliver unexpected sonic output not just because of the mapping but because of the individuals performing it.

The mapping on “Leaps” aims to be intuitive to anyone by utilising recognisable metaphorical actions. Whilst I think this has been done well, it is also the case that there are times where a performer’s actions will not always be seen to be making sonic output. This is because there are some times when the performer will be using these gestures to navigate through menus or set up sounds. Whilst this system is preferable to the alternative action of physically interacting with the computer, it can be confusing for a spectator to see actions that previously triggered sounds, not trigger them. Overall, I think there is ample scope to alter these navigation control gestures to be different to the sound control gestures.

“Chess” creates an instantly unfamiliar atmosphere by utilising a known g ame as an interface for music. The piece plays on the tension of the game and aims to make an increasingly uncomfortable sonic atmosphere. The mapping is not directly linked to the physical movement of pieces but rather it is linked to the meaning these moves hold within the context of the game. This concept is sometimes difficult to contextualise within a game, as it is not always instantly obvious why a sound has been played back, without thinking about the context of the game. This mapping can potentially exclude some spectators from understanding the piece. Whilst everyone’s chess ability is varied, if the performers are drastically more experienced at playing chess than the audience, their intent may not always be clear.

The obstacles faced within “Chess” have included the interpretation of chess information and the collection of data itself. The submitted performance is performed on a standard chess board and the data is collected live, and processed live. “Chess” has the scope to develop upon its GUI interface and utilise hardware components within a physical chess set to track player movements in real time. This would eliminate the need for the data to be input into the system manually. Furthermore, it would be extremely convenient to place these interfaces in a standalone system utilising modern existing technologies such as the Raspberry PI ("Raspberry Pi", 2016) or the Arduino (“Arduino", 2016).

Throughout this project, I have developed my ability in creative programming and composition. I have gone through many iterations of the submitted systems and developed the structure of the programming. This development in my understanding has been immense and the quality of output has drastically evolved. Whilst there is still plenty of scope for development within both of these interfaces, I am confident that they represent the concepts described within this document. Both pieces created contain a lot of attention to detail when it comes to shaping the desired sonic output. The patching within both of these pieces goes into a lot of detail to make sure there is a varied performance within a coherent structure. In the future, I hope to continue learning about I hope to continue to develop upon these pieces and indeed the interfaces that form them.

References

Agnelli, J. (2009). Birds on the Wires. Retrieved from https://www.youtube.com/watch?v=LoM4ZZJ2UrM

Arduino - Home. (2016). Arduino.cc. Retrieved 11 March 2016, from https://www.arduino.cc/

Bernstein, J. (2016). jeremy bernstein :: bootsquadresearch. Bootsquad.com. Retrieved 11 March 2016, from http://www.bootsquad.com/

Fels, S., Gadd, A., & Mulder, A. (2002). Mapping transparency through metaphor: towards more expressive musical instruments. Organised Sound, 7(02). http://dx.doi.org/10.1017/s1355771802002042

Françoise, J. (2014). Leap Motion skeletal tracking in Max. Jules Françoise. Retrieved 11 March 2016, from http://julesfrancoise.com/leapmotion/

Hamlen, J. (2014). About. Joe Hamlen - Composer. Retrieved 27 September 2016, from http://www.joehamlen.co.uk/about/

Harker, A., & Tremblay, P. (2016). The HISSTools Impulse Response Toolbox: Convolution for the Masses. In The International Computer Music Association (pp. 148-155). ICMC 2012: Non-cochlear Sound.

Kinect - Windows app development. (2016). Developer.microsoft.com. Retrieved 26 September 2016, from https://developer.microsoft.com/en-us/windows/kinect

Krefeld, V. & Waisvisz, M. (1990). The Hand in the Web: An Interview with Michel Waisvisz. *Computer Music Journal*, *14*(2), 28. http://dx.doi.org/10.2307/3679709

Leon Theremin playing his own instrument. (2008). Retrieved from https://www.youtube.com/watch?v=w5qf9O6c20o

Marcel Duchamp & John Cage Play Musical Chess. (2016). *The Toronto Dreams Project Historical Ephemera Blog*. Retrieved from http://torontodreamsproject.blogspot.co.uk/2015/03/marcel-duchamp-john-cage-play-musical.html

Max Msp (Version 6). (2016). Retrieved from https://cycling74.com/

Motion, L. (2016). Leap Motion. Leapmotion.com. Retrieved 11 March 2016, from https://www.leapmotion.com/

Paine, G. (2013). New Musical Instrument Design Considerations. IEEE Multimedia, 20(4), 76-84. http://dx.doi.org/10.1109/mmul.2013.60

Propellerhead. (2016, August 08). Retrieved September 26, 2016, from Propellerhead, Reason, https://www.propellerheads.se/reason/instruments/player-devices#scales-chords

Raspberry Pi - Teach, Learn, and Make with Raspberry Pi. (2016). Raspberry Pi. Retrieved 11 March 2016, from http://www.raspberrypi.org

Reactable | Music Technology Group. (2016). Mtg.upf.edu. Retrieved 26 September 2016, from http://mtg.upf.edu/project/reactable

Steinberg,. (2013). Cubase iC Air gestures control for Cubase. Retrieved from https://www.youtube.com/watch?v=ObIAFP7E5_Q

Stockfish - Open Source Chess Engine. (2016). Stockfishchess.org. Retrieved 11 March 2016, from https://stockfishchess.org/

the Gloves. (2016). Imogenheap.co.uk. Retrieved 11 March 2016, from http://www.imogenheap.co.uk/thegloves/

Traubeck, B. (2011). Years. Schmiede Hallein. Retrieved from http://traubeck.com/years/

Yoo, M.-J., Beak, J.-W., & Lee, I.-K. (2011, June 1). Proceedings of the International Conference on New Interfaces for Musical Expression. Retrieved September 26, 2016, from NIME.org, http://www.nime.org/proceedings/2011/nime2011_324.pdf