Keele University

Sound design, composition and performance with interactive genetic algorithms

Ph.D. Music

Tom Prescott

06/2018

Contents

List of Figures

Abstract

A variety of work has been carried out investigating the suitability of interactive genetic algorithms (IGAs) for musical composition. There have been some promising results demonstrating that it is, in principle, an effective approach. Modern sound synthesis and processing techniques (SSPTs) are often very complex and difficult to use. They often consist of tens or hundreds of parameters and a large range of values can be assigned to each parameter. This results in an immense number of parameter combinations; listening to the result of each one is clearly not viable. Furthermore, the effect each parameter has on the audio output may not be immediately obvious. Effectively using these systems can require a considerable time commitment and a great deal of theoretical knowledge. This means that in many cases these techniques are not being used to their full potential.

IGAs offer a solution to this problem by providing a user with a simpler, more accessible interface to a range of SSPTs. This allows the user to navigate more effectively through the parameter space and explore the range of materials which can be generated by an SSPT.

This thesis presents compositions and software that investigate a range of approaches to the application of IGAs to sound design, composition and performance. While investigating these areas, the aim has been to overcome the limitations of previous IGA based systems and extend this approach into new areas. A number of IGA based systems have been developed which allow a user to develop varied compositions consisting of diverse and complex material with minimal training.

Acknowledgements

I would like to thank Professor Rajmil Fischman for his guidance, encouragement and infinite patience.

I would also like to thank the Keele Research Institute for the Humanities and Social Sciences for funding to purchase a computer with the specifications necessary to develop the software presented throughout this thesis.

Chapter 1 – Introduction

1.1 Overview

An interactive genetic algorithm (IGA) is a biologically inspired search procedure based on artificial selection, similar to the ways in which natural selection affects the evolution of animal and plant species. This automated selection capability has prompted the investigation of interactive genetic algorithms' suitability for sound design and musical composition. There have been some promising results demonstrating that it is, in principle, an effective approach (Dahlstedt 2007). However, there are also a number of limitations to current systems which the work carried out throughout this Ph.D. aims to address.

One of the primary motivations of IGA based systems has been to address the complexity of modern sound synthesis and processing techniques (SSPTs) which often require large amounts of data and may be difficult to use: SSPTs may consist of tens or hundreds of parameters and a large range of values can be assigned to each parameter, resulting in an immense number of parameter combinations. Therefore listening to the result of each one of these combinations is clearly not viable. Furthermore, the effect each parameter has on the audio output may not be immediately obvious. Effectively, using these systems can require a considerable time commitment and a great deal of theoretical knowledge, which suggests that in many cases these techniques might not be used to their full potential.

A number of systems have demonstrated the effectiveness of applying IGAs to sound design and composition by providing a user with a simpler, more accessible interface to a range of SSPTs (Dahlstedt, 2007). This allows the user to navigate more effectively through the parameter space and explore the range of materials which can be generated by an SSPT. However, these systems have been limited by making a trade-off between ease of use and flexibility. Systems which are simpler to use tend to be limited in the diversity and complexity of material which can be generated (Johnson, 1999). Conversely, systems which can generate a diverse range of materials tend to require extensive customisation, which demands a substantial amount of technical knowledge (Dahlstedt, 2007).

1.2 Thesis contents

The Ph.D. consists of compositions and software that investigate a range of approaches to the application of IGAs to sound design, composition and performance. These are grouped into three substantial components:

1.2.1 Sound Design and Fixed Media Composition

1.2.1.1 Software

The *Evolutionary Sound Design Environment* (ESDE) software provides an interface for a range of sound synthesis and processing techniques (SSPTs) which can be controlled with an IGA. The purpose of ESDE is to allow the user to navigate the search space of the SSPTs to generate material suitable for sound design and the composition of fixed media electroacoustic works. ESDE demonstrates one of the more important developments of earlier IGA based approaches to sound design and composition: the IGA evolves breakpoint functions rather than directly controlling the parameters, allowing a range of mapping strategies to be developed which increase the diversity and complexity of the material the software can generate.

1.2.1.2 ESDE Compositions

Two compositions were composed using the ESDE software in order to test and demonstrate its validity. The first, *The Blind Watchmaker* is a stereo fixed media work which establishes the core compositional strategies for IGA based composition. The second composition, *The Singing Forest* is a stereo fixed media work which uses a wider range of sonic material to further explore IGA based compositional strategies.

1.2.1.3 Large-scale composition

The Wind-Up Bird can be considered a culmination of the technical and compositional approaches developed throughout the Ph.D. The IGA based software and compositional strategies are applied to the composition of a multi-channel large-scale composition.

1.2.2 Live Performance

1.2.2.1 Software

The *MAES+IGA* software was developed to investigate approaches to live performance with IGAs. *MAES+IGA* extends Rajmil Fischman's Manual Actions Expressive System (Fischman, 2013) to allow the user to control IGAs driving SSPTs in real time with natural hand actions.

1.2.2.2 Compositions

Two compositions were composed using the MAES+IGA software. *Snowstorm* and *Terraform* are multi-channel works which demonstrate a range of approaches to constructing a composition performed with an IGA controlled by natural hand gestures. These compositions allow the user to interact with surreal landscapes, enabling improvisation within a structured environment.

1.2.3 Visualisation of Interactive Genetic Algorithms

An animated visualisation investigates further approaches to real-time control of IGAs. It consists of an interactive animation which represents the state of multiple IGAs in a customised version of the ESDE software. The visualisation is designed to be presented in the form of a customisable installation which cycles through three different states, functioning as sections of an interactive composition. The visualisation also serves as an educational tool by demonstrating the principles an IGA is based on in a way which can be quickly understood. It also demonstrates how an IGA can be applied to allow a simple interface to control multiple, sophisticated SSPTs.

References

Dahlstedt, P., (2001). A MutaSynth in Parameter Space: Interactive Composition through Evolution. *Organised Sound,* 6(2), pp. 121-124.

Dahlstedt, P., (2007). Evolution in Creative Sound Design. In: E. M. a. J. Biles, ed. *Evolutionary Computer Music.* London: Springer, pp. 79-99.

Fischman, R., (2013). A Manual Actions Expressive System (MAES). *Organised Sound,* 18(3), pp. 328-45.

Johnson, C.B. (1999). Exploring the sound space of synthesis algorithms using interactive genetic algorithms. In *Proc. of the AISB '99 Symposium on AI and Musical Creativity*. Brighton, UK. Society for the Study of Artificial Intelligence and Simulation of Behaviour, pp. 182-196.

Chapter 2 – Evolutionary Sound Design Environment

2.1 Introduction

This chapter describes the *Evolutionary Sound Design Environment* (ESDE), a system which implements an interactive genetic algorithm (IGA) to aid in sound design and the composition of acousmatic music. ESDE allows users to design complex and varied sounds using the IGA regardless of whether they have any formal musical training. This is possible as the users can experiment with a system that requires little direct interaction with the SSPT parameters and requires minimal configuration for initial experimentation.

The primary objective of developing the ESDE software was to create a system which makes use of an IGA to generate material of sufficient variety and complexity for composing acousmatic music. There have been a number of approaches to applying an IGA to sound design and composition which have demonstrated the benefits of this approach. However, they have limitations which prevent them from being used as fully functioning, practical tools for sound design and composition. This is generally due to serving primarily as a proof of concept or having a different area of focus; for example, composing instrumental music. The development of ESDE was informed by a review of these systems to determine which aspects have been successful and areas where further development is needed.

2.2 Background

2.2.1 Existing Systems

Johnson (1999) makes use of an IGA to construct an interface for the CSound FOF granular synthesis algorithm. This system applies a domain specific approach (see 2.2.2.1 below). The user assigns a numerical rating to each sound in a generation, the higher ranking sounds have a higher probability of reproducing. This system demonstrates an effective means of mapping binary strings onto synthesis parameters and allows the user to evolve usable sounds in relatively few generations. It also provides an example of an accessible interface. However, while serving as an effective proof of concept, the system has a number of drawbacks. Firstly, it requires a user to continuously rank sounds from 1 to 9: this is likely to become tedious fairly quickly. Also, as it only controls one sound synthesis algorithm, the range of material which can be generated is limited.

Palle Dahlstedt describes various approaches to implementing an IGA for controlling sound synthesis techniques and the properties such a system requires to be successful (Dahlstedt, 2001). He discusses MutaSynth which is a system designed to be compatible with a wide range of synthesisers, providing an example of an open ended approach (see 2.2.2.2 below). The generality of MutaSynth proved to be both its strength and weakness. While it could control almost any MIDI (Musical Instrument Digital Interface) synthesiser, extensive manual configuration was required.[1]

MutaSynth was then developed into Patch Mutator in collaboration with the synthesiser manufacturer Clavia (Dahlstedt, 2007). This is particularly notable as it is the first time that a professional sound synthesis tool has been provided with an interactive evolution mechanism. In contrast to MutaSynth, Patch Mutator is integrated into the synthesis environment. It is therefore a domain specific approach in which it is possible for the breeding engine to make fairly intelligent choices about which parameters should be included in the genotype and exclude those which are not likely to be useful. It appears that the domain specific knowledge of Patch Mutator is a significant advantage as the exclusion of parameters which are not useful decreases the probability of unsuitable solutions being evolved and speeds up the process.

[1] MIDI (Musical Instrument Digital Interface) is a protocol which allows communication between electronic musical instruments and computers.

Also, Patch Mutator arguably limits the drawbacks of the domain specific approach by controlling a modular synthesiser, which offers a vast range of sonic material.

Collins (2001) applies IGAs to reverberation, wavetable synthesis, synthesis of percussive sounds and an analytical solution to the stiff string. A notable aspect of this work is the comparison of a trial and error approach to exploring the search space with the IGA. The trial and error approach proved to be time consuming and unproductive. Also, selecting random points in the search space does not produce suitable solutions. This eliminates the possibility of a genetic algorithm being helpful in this context simply due to selecting a random number of points in the search space and happening to eventually find a solution. It was also discovered that the fitness bottle neck did not prove to be particularly problematic[2]

2.2.2 Review

Before proceeding with the development of the ESDE software, the existing systems were reviewed to determine which approaches should be carried forward and areas were further development is required. The review revealed that the systems can be loosely grouped into two categories, *domain specific* and *open ended*, each with their own strengths and weaknesses.

2.2.2.1 Domain Specific

The domain specific approach tailors the IGA to specific sound synthesis and processing techniques (SSPTs), allowing for the development of systems which are very easy to use and require no customisation from the user. Due to the tight integration with SSPTs, systems in this category can implement features which cannot be controlled solely by a numerical input. For

[2] A fitness bottleneck is a situation where too many new populations need to be created to find a usable solution.

example, evolving the shape of a filter. However, these systems also tend to support very few SSPTs which limits the variety of material they can generate.

2.2.2.2 Open Ended

The open ended approach allows an IGA to be customised for a vast range of sound synthesis and processing techniques providing endless possibilities. Given enough time and expertise on the part of the user, they can serve as sophisticated tools to aid in composition. The drawback of these systems is the complex and time consuming nature of the customisation process. This is arguably self-defeating as the primary benefits of an IGA are ease of use and quicker search space navigation. An example of a system in this category is Dahlstedt's MutaSynth (2007).

2.2.2.3 Shortcomings for a Practical System

The review also revealed that there are a number of features which are necessary for developing a practical composition tool which are not present in existing systems. These features are missing as some systems were intended to serve as a proof of concept or an investigation into the behaviour of IGAs rather than fully functioning tools. Also, there has not been a significant focus on composing acousmatic music so features required for this area have not been implemented.

An example of an essential feature for the composition of acousmatic music is time varying timbral parameters. Without these the variety and complexity of material which can be generated will be very limited and it is unlikely a composer would be able to use the system continuously over time for numerous compositions.

2.2.3 Summary

The review of existing systems showed that a middle ground between the domain specific and open ended approaches may be beneficial. The system should retain the ease of use of the domain specific approach while allowing some of the versatility present in the open ended approach. For a tool to move beyond the proof of concept stage, it is important for it to be able to integrate into a user's workflow by providing a preset system in which the state of the system can be saved at any point and easily recalled. This is an essential development as some earlier systems overwrite the parent generation and replace it with the child population, meaning the state of the system cannot be later recalled. Also, a large degree of flexibility should be provided to allow the user to customise the system as their skills with it increase.

2.3 Software Implementation

The software consists of an interface implemented in MAX, an IGA implemented in FTM [3] & Max, and the synthesis and processing modules (Figure 2.1).

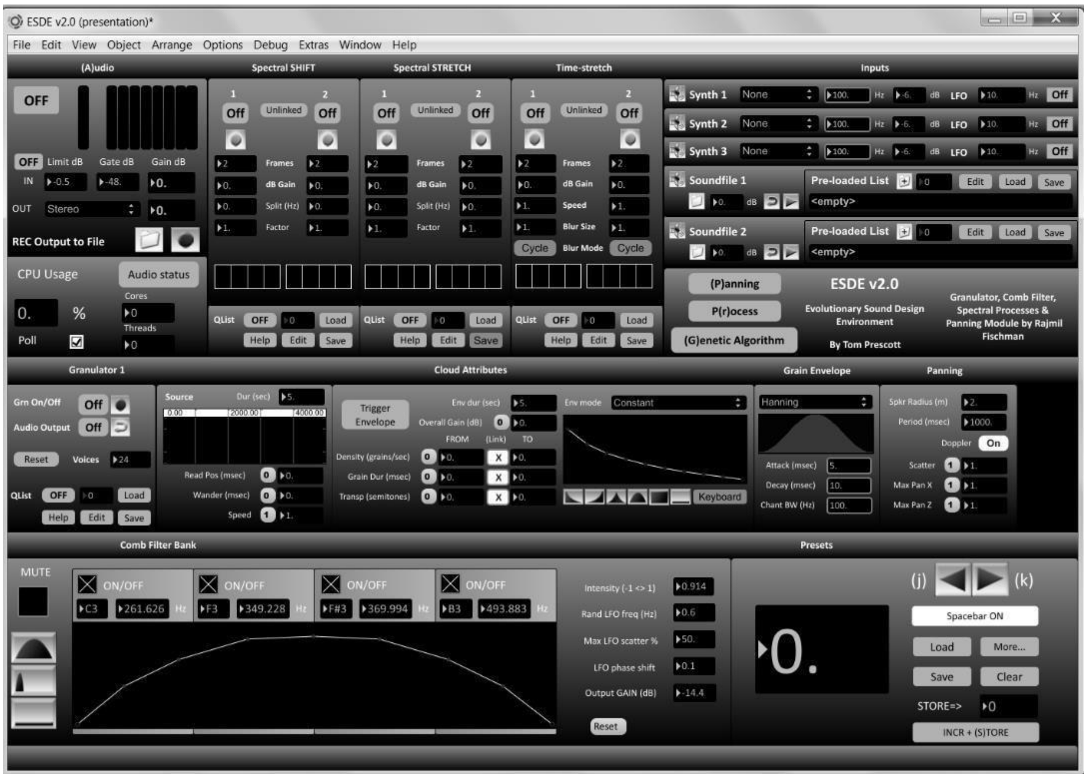

Figure 2.1: The ESDE Interface

2.3.1 Synthesis and Processing

The synthesis and processing modules are interconnectable by means of a patchbay emulating matrix. The synthesis and processing modules were provided by Rajmil Fischman[4] and consist of the following:

---

[3] FTM introduces a range of more complex data structures into MAX. For more information on FTM see (Schnell, et al., 2005)

[4] For more information on the synthesis and processing modules see (Fischman, 2015).

1. Sound sources

   1.1 Three synthesisers (including microphone capture).

   1.2 Two audio file players (1–8 channels).

2. Spectral processes

   2.1 Two spectral shifters (Fischman 1997: 134–5).

   2.2 Two spectral stretchers (Fischman 1997: 134–5).

   2.3 Two time stretchers (Fischman 1997: 134–5).

   2.4 Two spectral blur units (Charles 2008: 92-4).

3. Comb Filters[5]

- A bank of four comb filters.

4. Asynchronous granulation[6]

- This includes the control of sample read position, wander and speed (time-stretch); grain density, duration, transposition and spatial scatter; and cloud envelope.

5. QList Automation

- MAX QLists allow smooth variation of parameters in time according to breakpoint functions.

6. Spatialisation

- A proprietary algorithm developed implements spatialisation in stereo, surround 5.1 and two octophonic formats, including optional Doppler shift.

[5] For more information on comb filters see (Smith, 2010).

[6] De Poli, Piccialli and Roads (1991: 137–86) explain asynchronous granular synthesis.

> The matrix patchbay enables the connection of the outputs of any of the synthesisers, file players and processes to ten independent spatialisers. The granulator features six additional spatialisers and the file players can be routed directly to the audio outputs, which is useful in the case of multichannel files that are already distributed in space (Fischman 2015).

2.3.2 Interactive Genetic Algorithm Implementation

2.3.2.1 Mapping

The implementation of time varying parameters was identified as a requirement for a system which can generate material of sufficient variety and complexity for the composition of acousmatic music. This is achieved by implementing an IGA which can evolve up to 24 time-varying breakpoint functions independently of the SSPT parameters. The breakpoint functions are then mapped onto the SSPT parameters through a matrix which allows for one-to-one and one-to-many mappings (Figure 2.2). This differs from the previous approach of mapping values directly onto the parameters and provides a number of advantages. For example, a breakpoint function mapped onto the granulator's grain transposition parameter could be remapped onto the factor of the spectral pitch shifter. While creating a different effect, the pitch trajectory will remain the same and should be identifiable. This provides an example of how the IGA could be used as the basis for a composition strategy to generate and develop material.

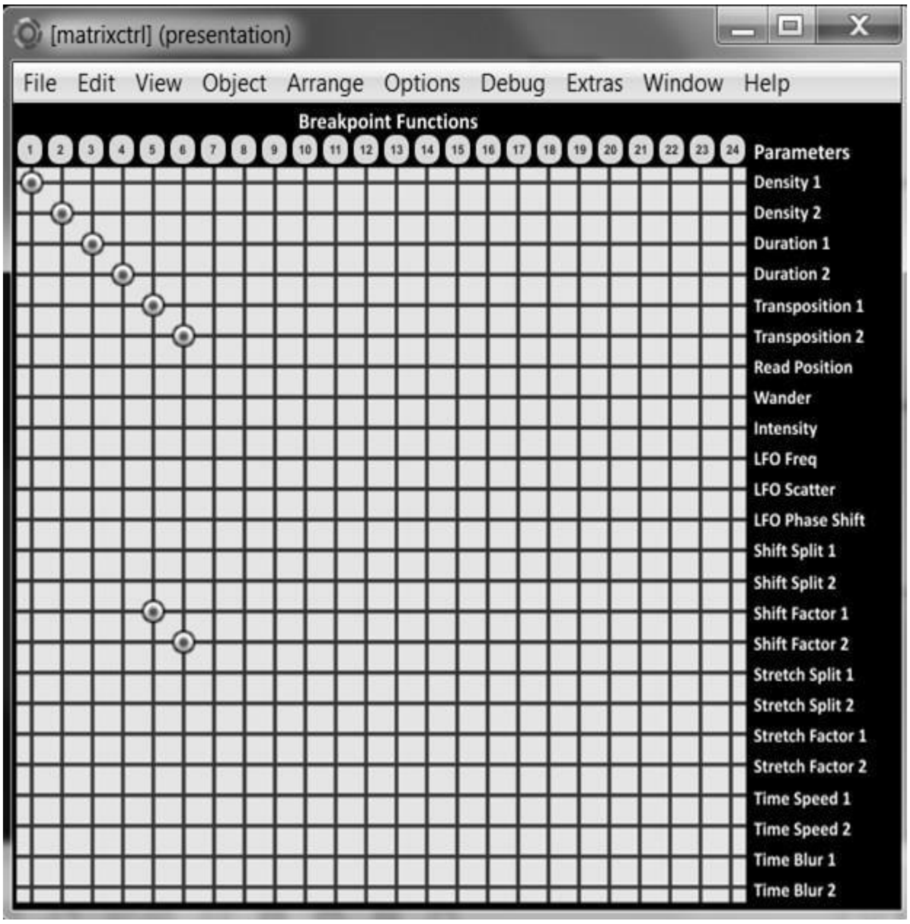

Figure 2.2: Mapping Matrix. Breakpoints mapped onto grain transposition and spectral shift factor.

The IGA is implemented within FTM, it allows new randomly generated populations to be created in one of 24 matrices. Each population consists of a number of binary strings, the number and length of these is determined by the user. Each binary string within the population is known as a genotype. The IGA uses these populations to generate breakpoint functions and then control the SSPTs through the following stages:

1. The genotype is divided into a number of substrings equal to the number of processing parameters under IGA control. This avoids redundancy in the genotype. If sections are not mapped onto any parameters this will cause the genetic operators (described in

2.3.3 below) to have unpredictable effects as they will be operating on parts of the genotype which are not mapped onto any of the breakpoint functions,

2. Each substring is then divided into a further set of substrings equal to the number of breakpoints in each function.

3. Each substring is then converted into a decimal number between 0 and 1 and placed at equally spaced intervals to create the time varying breakpoint function.

4. The values from the breakpoint function are sent to the inputs of the mapping matrix over time according to the duration of the breakpoint functions.

5. When the values are sent from the outputs of the mapping matrix they are rescaled from the range [0, 1] to a range suitable for the SSPT parameter they are being mapped onto. The range for each parameter can be controlled by the user, the menu for setting the ranges is shown in Figure 2.3. A range of defaults are provided which will be likely to provide usable material. They can then be customised as the user becomes more familiar with the SSPTs.

2.3.3 Genetic Operators

The user listens to the sound generated by each genotype, if none are suitable a new random population can be created. If a sound which may be useful for the user's purposes is discovered there are three options: create a child population through mutation, create a child population through crossover or save the genotype into temporary storage. *Mutation* uses a single parent genotype to create a child population, whereas *crossover* uses two parent genotypes to generate a child population. Three crossover techniques are implemented in ESDE: One point crossover, two point crossover and uniform crossover; as described below.

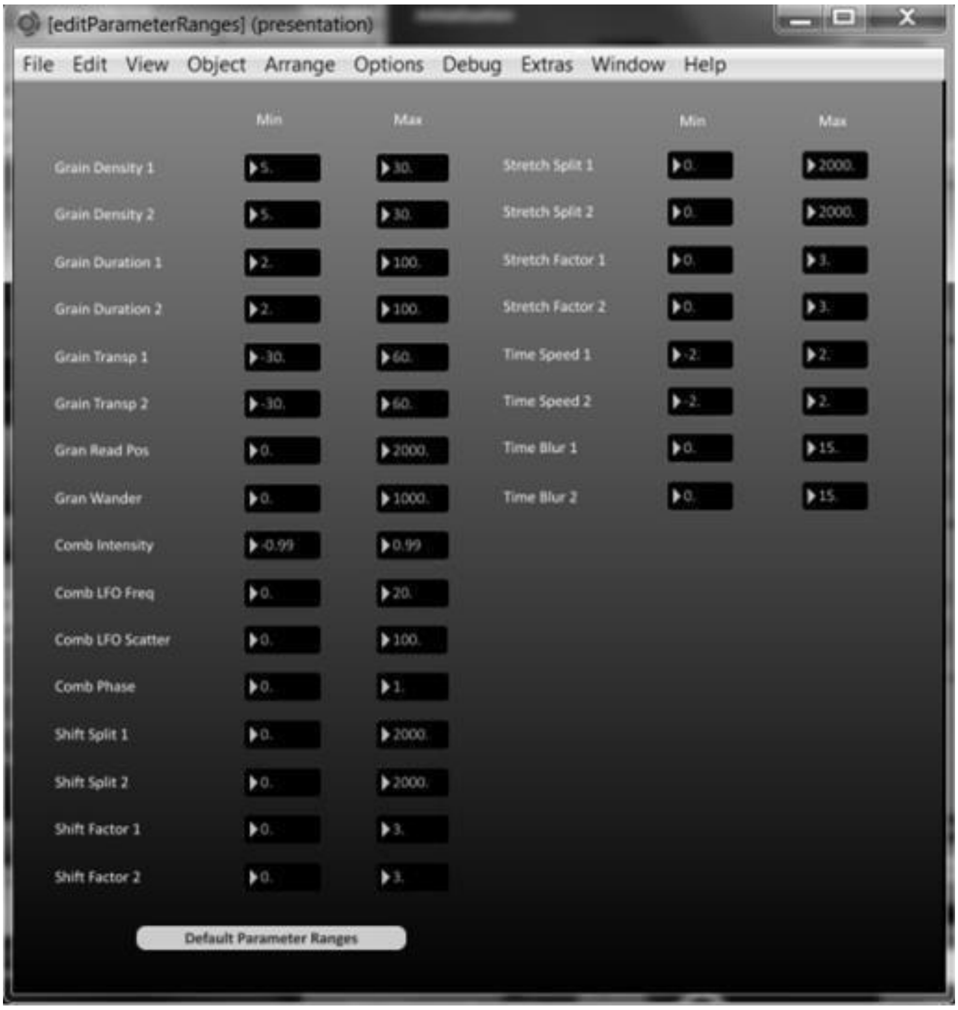

Figure 2.3: The parameter range window, this is used to select the range the values from the breakpoint function will be scaled to.

2.3.3.1 Mutation

The parent genotype is copied into a new matrix, then duplicated to create a child population containing the same number of genotypes as the parent population. Variation in the child genotypes is created by randomly selecting and flipping bits in the binary string. 1s are changed to 0s and 0s to 1s. The number of bits which are flipped is determined by the mutation rate, this allows the user to select the percentage of bits which will be flipped in each genotype. A low mutation rate of around 5 – 20% tends to result in minor variations from the parent, a rate of around 20 – 40% tends to create child genotypes which will be recognisable as connected to the parent but with some substantial differences. When the rate is above 40% the results become more unpredictable, there may be audible traces of the parent in the child genotypes or they may be entirely unrecognisable. The effect of a given mutation rate can vary somewhat depending on a number of factors such as the selected processing parameters, the range of values which can be assigned to each parameter and the number of breakpoints in each function.

2.3.3.2 One Point Crossover

One point crossover randomly selects a cut point, which is a randomly generated number between 0 and the length of the binary string. The section of the string to the left of the cut point from parent 1 is combined with section to the right of the cut point from parent 2 to create a new string.

2.3.3.3 Two Point Crossover

Two point crossover generates two cut points. The section to the left of cut point 1 and to the right of cut point 2 are selected from parent 1. They are combined with the section between cut points 1 and 2 from parent 2.

2.3.3.4 Uniform Crossover

Uniform crossover randomly selects half the bits from parent 1 and half from parent 2 and combines them to create a child genotype. This approach is less likely to preserve particular qualities which have been identified in the parents, it is intended to be used as a technique to further explore the search space.

After crossover has been applied to create a child genotype, the genotype is copied into a new matrix and duplicated to create a child population containing the same number of genotypes as the parent population. Mutation is then applied as described in 3.3.1. A range of crossover techniques have been implemented as the most suitable will vary depending on the parameters which are being controlled by the IGA.

2.3.3.5 Temporary Storage

A further matrix is provided as a temporary storage area if the user does not wish to immediately make use of a genotype but thinks it may be of some use lat er. Genotypes in the temporary storage matrix can be reintroduced into the population using crossover at any point, or mutation can be applied to create a new child population. ESDE places a lower degree of importance in temporary storage compared to some earlier IGA based systems as the creation of a child population does not overwrite the parent population. The primary purpose in ESDE is to provide a further means of exploring the search space.

2.3.4 Additional Features

There are also a number of additional features which are designed to make the system as practical as possible and allow it to be integrated into a composer's workflow.

2.3.4.1 Preset System

A preset system is included to allow for the storage and recall of process choice, mappings, spatialisation settings and discrete parameters, which cannot be controlled by a breakpoint function. There is also a system for saving and loading populations to and from .txt files. This means that the exact state of the system can be stored and recalled at any point. An approach to using this may be to have a preset for each section of a composition to allow the composer to retain a record of how all material for a composition was created and return to it for further development.

2.3.4.2 Population Management

Some additional population management features are provided so storage of the populations can remain organised. Populations can be copied, passed and swapped between matrices. For example, a user may decide a population stored in matrix 24 contains material suitable for the beginning of a composition. This can easily be relocated to matrix 1 so that the populations are stored in a logical order.

The IGA is customised to take advantage of the cloud envelope feature of the granulator, this is an example of a domain specific approach. If the trigger envelope option is selected and the breakpoint duration set to the same value as the cloud envelope duration; the cloud envelope will begin at the same time as the breakpoint functions, allowing for the generation of repeatable granular clouds. This approach can also be applied to the sound file and synthesiser inputs using the preset mechanism.

2.4 Summary

ESDE demonstrates a middle ground between open ended and domain specific IGA implementations which allows the user to control a range of sound synthesis and processing techniques and generate a diverse range of sonic material. The process control matrix allows for a level of customisation while not requiring complex and time consuming customisation procedures. As computer processing power increases over time, a greater range of sound synthesis and processing techniques could be included, providing further options for the user without increasing the complexity.

References

Charles, J.-F., (2008). A Tutorial on Spectral Sound Processing Using Max/MSP and Jitter. *Computer Music Journal,* 32(3), pp. 87-102.

Collins, N., (2001). Experiments with a New Customisable Interactive Evolution Framework. *Organised Sound,* 6(2), pp. 121-124.

Dahlstedt, P., (2001). A MutaSynth in Parameter Space: Interactive Composition through Evolution. *Organised Sound,* 6(2), pp. 121-124.

Dahlstedt, P., (2007). Evolution in Creative Sound Design. In: E. M. a. J. Biles, ed. *Evolutionary Computer Music.* London: Springer, pp. 79-99.

De Poli, G., Piccialli, A. and Roads, C. 1991. Representation of Musical Signals. Cambridge, MA: The MIT Press, pp. 137-86.

Fischman, R., (1997). The Phase Vocoder: Theory and Practice. *Organised Sound,* 2(2), pp. 127-145.

Fischman, R., (2013). A Manual Actions Expressive System (MAES). *Organised Sound,* 18(3), pp. 328-45.

Fischman, R., (2015). *Rajmil Fischman's Free Software.* [Online] Available at: http://www.keele.ac.uk/music/people/rajmilfischman [Accessed 12 May 2015].

Johnson, C.B. (1999). Exploring the sound space of synthesis algorithms using interactive genetic algorithms. In *Proc. of the AISB '99 Symposium on AI and Musical Creativity*. Brighton, UK. Society for the Study of Artificial Intelligence and Simulation of Behaviour, pp. 182-196.

Schnell, N., Borghesi, R., Schwarz, D., Bevilacqua, F. and Muller, R. (2005). *FTM - Complex Data Structures for Max.* Barcelona, ICMC.

Smith, J. (2010). *Comb Filters*. [ONLINE] Available at:

https://www.dsprelated.com/freebooks/pasp/Comb_Filters.html.

[Accessed 22 November 2017].

Chapter 3 – ESDE Compositions

3.1 The Blind Watchmaker

The Blind Watchmaker is a stereo fixed media work which was assembled from material generated with the ESDE software. A small amount of additional processing was used (e.g. reverberation) however this was kept to a minimum to draw attention to the material generated with the software and the interactive genetic algorithm (IGA) based compositional strategies.

The objectives of writing this composition were to:

1) Demonstrate the capability of the software to generate material suitable for electroacoustic composition, particularly through the use of time varying parameters.
2) Investigate strategies for generating sound material with IGAs.
3) Explore approaches to applying IGAs to compositional strategy.

3.1.2 Programme Notes

The Blind Watchmaker takes its inspiration from the theory of water being brought to the Earth by the eruption of the first volcanoes over 4.5 billion years ago. Steam was created and clouds were formed as the planet cooled. This started a rainstorm lasting for thousands of years, paving the way for the start of the evolutionary process and the development of life on Earth.

The title refers to the analogy used by Richard Dawkins to explain the development of complexity in nature from randomness coupled with cumulative selection. This process inspired the approach used to develop the sonic material for the composition. While the narrative themes and large scale structure of the composition were determined in advance, the micro and meso-level structure were unknown and generated by an open ended evolutionary process.

This was achieved through the use of an interactive genetic algorithm, a biologically inspired search procedure based on artificial selection, which was mapped onto a range of sound synthesis and processing techniques. This process naturally lends itself to musical composition as it has a great deal in common with aspects of the human creative process in which ideas are refined, combined, discarded and reintroduced in new forms.

3.1.3 Structure

The Blind Watchmaker consists of five main sections.

Section 1 - 00:00 – 00:41

Section 1 introduces the granular material which is repeated throughout the piece. This also introduces the core IGA compositional strategy of sequencing multiple genotypes from a population to develop material.

Section 2 – 00:42 – 02:18

Section 2 introduces high pitched note based material, which begin a movement away from the noise based granular material from the introduction. This movement is continued by processing the granular material with a comb filter bank.

Section 3 – 02:19 – 05:22

Section 3 begins with a return towards noise based granular material, which develops into layered textures. The section concludes with more prominent note based material in preparation for the climax of the composition.

Section 4 – 05:23 - 08:09

Section 4 provides the climax of the composition by applying a comb filter bank to the granular material to create a chord based sequence. This is accompanied by a return of the layered noise based material from section 3.

Section 5 – 08:10 – 09:34

The composition is concluded with a variation of the material introduced in sections 1 and 2 of the composition.

3.1.4 Compositional Strategies

This section describes the most frequently used strategies for generating sound material with the ESDE software. The core strategy for generating sound material for a composition is shown in figure
3.1, this will be referred to as strategy 1, further approaches are variations of this strategy.

3.1.4.1 Strategy 1

This strategy can be considered the core IGA based compositional strategy, and the starting point for the development of all other strategies.

1) Generate a random population.
2) Listen to sound material generated by genotypes in randomly generated population.
3) If a genotype generates material with some characteristics which may be appropriate for the composition, set mutation rate to ca. 40% and generate new child population. If material is not suitable, go back to 1.

4) Listen to sound material generated by child population to explore the search space. If material is an improvement on the parent then decrease mutation rate to 10% and generate new child population to explore minor variations. If material is not suitable go back to 3.

5) Listen to minor variations and select the genotype which generates material most suitable for the composition.

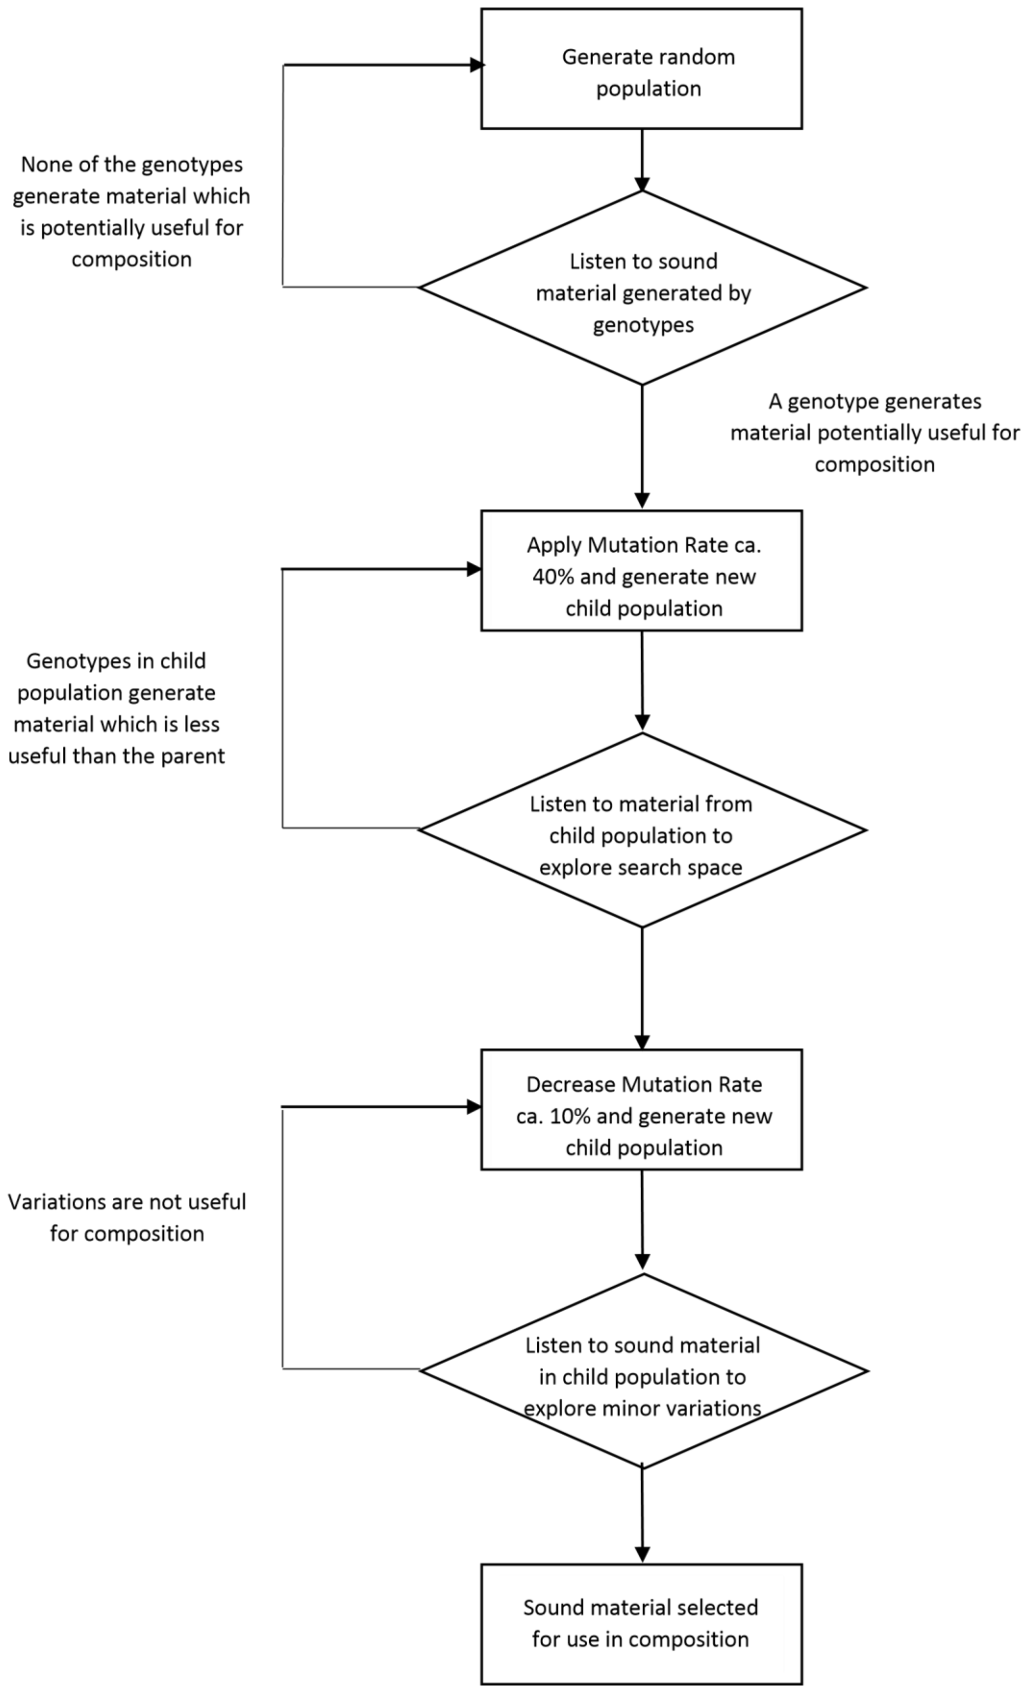

Figure 3.1: IGA strategy 1 workflow.

3.1.4.2 Strategy 2

Strategy 2 extends Strategy 1 by generating a further child population which enables the material to be developed. Also the duration of the breakpoint functions is altered for the child population to provide additional variety. It consists of the following stages:

1) Generate a random population.
2) Listen to sound material generated by genotypes in randomly generated population.
3) If a genotype generates material with some characteristics which may be appropriate for the composition, set mutation rate to ca. 40% and generate new child population. If material is not suitable, go back to 1.
4) Listen to sound material generated by child population to explore the search space. If material is an improvement on the parent then decrease mutation rate to 10% and generate new child population to explore minor variations. If material is not suitable go back to 3.
5) Listen to minor variations and select the genotype which generates material most suitable for the composition.
6) Generate new child populations with a low mutation rate with the genotype selected in 5 as the parent to provide developments of the material.
7) Adjust breakpoint function duration to provide additional variation.

The first example of this strategy can be heard at the start of the piece, with the low pitched, granular, bubbling material. This material can be heard at 00:01, 00:08 and 00:13. The first occurrence at 00:01 is the parent sound. A new generation was created based on this from which the second and third sounds (00:08 and 00:13) were selected. Additionally, the duration of the breakpoint function was reduced from 5 seconds to 3.5 seconds for the second occurrence and increased to 10 seconds for the third to provide some additional development of the material.

A second example of this process can be heard in the second major section of the piece beginning at 02:18. The genotype which generates the first noise based ‘rising water’ sound at 02:18 was used as the parent to create a new generation from which the following two sounds at 02:22 and 02:25 were selected. This sound was then used as a parent for the second time to create a new generation with a higher mutation rate which created a more significant variation heard at 02:28 and 02:30. This subsection was then completed by selecting two more sounds from the first child generation heard at 02:31 and 02:33.

3.1.4.3 Strategy 3

This strategy provides a significant development of material while retaining some connection to the original and providing a sense of cohesiveness throughout the composition[7].

1) Select a genotype used previously in the composition, developed through any other strategy.
2) Adjust the mapping in the mapping matrix to apply the same breakpoint function to a different parameter.

The subsection at (02:34) begins with the genotype which generates the ‘rising water’ sound at 02:18 acting as a parent for the third time: a higher mutation rate of 30% was selected and the breakpoint function duration extended to 15 seconds.

The variation at 02:41 was created by passing the previous sound from the granulator to the timestretcher. The same genotype was used for a third time at 02:52. This time the sound file

[7] There must be some similarities in the way the parameters affect the material for this strategy to be perceived. E.g. Granulation transposition to Spectral Shift factor.

used to create the low pitched granular material at the start of the piece was passed directly to the timestretcher, bypassing the granulator.

3.1.4.4 Strategy 4

This strategy applies the crossover procedure to develop material and reduces the importance of mutation.

1) Select two genotypes used previously in the composition, developed through any other strategy.
2) Create a new generation with the two genotypes using one point crossover and apply low mutation rate ca. 10%.

An example of Strategy 4 can be heard at 00:32, the granulated water material is generated from a crossover of the genotypes used to generate sounds at 00:18 and 00:20. Characteristics from both sounds have been preserved.

3.1.4.5 Strategy 5

This strategy is applied when the interactive genetic algorithm is not entirely suitable to the control of all parameters. For example, if applied to the Comb Filter Bank it will simply generate wildly fluctuating pitches which are unlikely to be of use in a composition. This is used to generate pitched material throughout the composition.

1) Provide fixed values for some parameters which are not suitable for control by the interactive genetic algorithm.
2) Provide variety in material by controlling parameters which are suitable for control.

The first example of the use of pitched material can be heard at 00:46, the pitch value for the comb filter is fixed, however the accompanying granulation is controlled by the interactive genetic algorithm.

The final section of the composition makes a further example of this approach. The values of the comb filter are fixed to generate a chord progression. However, the granular material allows for a significant level of development by the interactive genetic algorithm.

3.2.1 The Singing Forest

The Singing Forest is a stereo fixed media work which was assembled from material generated with the ESDE software. This composition uses a wider range of sonic material than *The Blind Watchmaker.* This provided an opportunity to further explore the compositional strategies developed for *The Blind Watchmaker* and also lead to the development of new strategies.

3.2.2 Programme Notes

The Singing Forest loosely tells the story of evolution across multiple time scales through a series of surreal scenes, beginning in the ocean and then transitioning to a forest. The composition explores the increasing complexity of organisms and ecosystems which emerge from the evolutionary process, rather than recreating a precise 'history' of how this happened on earth. The increasing complexity is expressed in two ways. Firstly, the mimetic material represents increasingly complex lifeforms. This begins with life in the ocean, followed by mammals on land, which then progress to more intelligent apes. Secondly, the increasing complexity of the ecosystems is expressed through the more frequent use of pitch throughout the composition and the harmonies generated by various 'singing' creatures.

The gradual changes which happen across many generations are explored by the development of the sonic material within each section which is generated by an interactive genetic algorithm. Populations of sounds are evolved, with child populations developing the musical material generated by their parents. The use of an interactive genetic algorithm as both a tool for generating material and for informing composition structure naturally lends itself to the themes explored in this composition.

3.2.3 Structure

The Singing Forest consists of four main sections.

Section 1 - 00:00 – 01:40

This first section introduces the nature/evolutionary theme with an ocean scene. This also presents some of the most prominent compositional strategies for the piece. For example, concurrently developing multiple populations.

Section 2 - 01:41 – 03:10

This section progresses the evolutionary theme with the introduction of the wolves, suggesting the progression to more complex organisms. It also continues the use of concurrently developing multiple populations, with the second serving as a counterpoint to the first.

Section 3 - 03:11 – 05:12

There is a shift towards more aural material (Emmerson, 1986) for the third section. However, some of the sounds are suggestive of birdsong which ensures this section remains cohesive with the rest of the composition. This section provides an opportunity to present the range of material which can be developed with the ESDE software and also allows for different compositional strategies to be applied

Section 4 - 05:13 – 08:00

The composition concludes with the final forest scene, returning to the more recognisable, mimetic material. The evolutionary theme is concluded with the monkeys, representing the most complex organisms portrayed in the composition.

3.2.4 Compositional Strategies

3.2.4.1 New Strategies

New strategies were developed while composing *The Singing Forest* beyond those used for *The Blind Watchmaker*.

3.2.4.1.1 Strategy 6

This strategy allows multiple lines of material to be concurrently developed.

1) Concurrently develop two populations
2) Interleave material generated by the populations to provide a counterpoint

An example of this strategy begins at 00:24 with the introduction of granulated seagull noises. This is developed through to 00:45 when it is joined by a second, substantially different variation of the seagull noises. These two populations are independently developed using mutation and crossover (Strategies 2 and 4).

A further example of this strategy can be heard beginning at 01:44 with two different types of granular material, the higher pitched crackling and low pitched scraping. The second serves as a counterpoint to the first. This approach is further developed in the third section beginning at 04:09 with the layered granular textures, three genotypes are layered over each other to increase the complexity of the material.

3.2.4.1.2 Strategy 7

Uniform crossover was applied as a further means of search space exploration instead of relying solely on mutation

1) Select two genotypes used previously in the composition, developed through any other strategy.
2) Create a new generation with the two genotypes using uniform crossover and apply low mutation rate ca. 10%.

The granular cloud at 03:59 was generated from two of the preceding clouds at 03:30, 03:41. The uniform crossover process introduced the more percussive quality which had not been discovered through mutation.

3.2.4.1.3 Strategy 8

This strategy allows for the development of more complex material.

1) Apply strategy 2 and 4.
2) Use crossfade or spectral morphing to combine genotypes to create longer gestures or sections and enhance range of material in the composition.

This strategy concludes the first section of the composition, there is a crossfade between three different genotypes to create the transition from splashing to a scraping effect.

This technique is also applied to create the transition into the final section of the composition at

05:09 with a crossfade from birdsong towards a sound suggestive of howling monkeys.

This approach is applied again immediately after to introduce the final section of the composition, beginning at 05:28. There are crossfades between seven different genotypes which were generated through mutation and crossover.

3.2.4.2 – Reapplied Strategies

These strategies first developed for *The Blind Watchmaker* were reapplied using new material for *The Singing Forest*.

3.2.4.2.1 Strategy 2

The wind sounds which open the composition were generated with strategy 2. A child population was created based on the genotype which provides the first wind sound. The duration of the breakpoint functions were altered for the child population to provide additional variety.

3.2.4.2.2 Strategy 5

This strategy was reapplied to generate the pitched material for the composition. However, it was further explored by increasing the complexity of the surrounding non-pitched material. For example, at 04:08 the genotype was allowed to generate values over a wider range of density, pitch and duration values than used in *The Blind Watchmaker*.

3.2.5 – Conclusion

The Blind Watchmaker and *The Singing Forest* show how the ESDE software can be used to create material suitable for composing acousmatic music; evolving time varying breakpoint functions in addition to the use of the process matrix allows for a range of complex and varied material to be generated. This is possible without any further customisation or configuration of the system, demonstrating its benefits over earlier approaches to IGA based composition systems. The latter tend to produce a limited range of material, in the case of domain specific systems or require complex and time consuming customisation procedures in the case of open ended systems. In comparison, *The Blind Watchmaker* and *The Singing Forest* feature a varied and diverse number of IGA approaches used to develop compositional strategies.

Chapter 4 - Live Performance with Interactive Genetic Algorithms

4.1 Introduction

This chapter describes a system for composing and performing music with an interactive genetic algorithm (IGA) which can be controlled in real time using natural hand actions[8] . The positive findings during the exploration of the use of IGAs for the generation of sonic material in the acousmatic medium in deferred time suggested that it is possible to use these algorithms for the generation and control of material in real time.

MAES+IGA allows users without formal musical training to design complex and varied sounds using the IGA and then perform them with gestures used in everyday life. One of the primary issues which has been encountered when developing an IGA system controlled by hand gestures is the lack of a clear causal connection between gestures and the resulting sound. This may be acceptable for composing music but presents some problems if the system is to be used for live performances. This is due to the level of randomness involved when generating material with an IGA.

MAES+IGA provides a solution to this issue by allowing the performer to work within a structured environment. Populations are evolved and stored, then loaded for a performance. This allows the performer a level of improvisation in choosing specific genotypes at specific moments in the performance, while being assured that the gestures will be suitable for the material which is generated: the gestures which have been developed use the metaphor [9] of interactions with organisms and natural processes, with the objective of creating a clear causal connection between gesture and sound despite the level of unpredictability present when

---

[8] This project was possible thanks to funding from the Keele Research Institute for the Humanities and Social Sciences which allowed me to purchase a computer with the specifications necessary to develop this software and perform *Snowstorm and Terraform*.

[9] The metaphor is the action which is being simulated with hand gestures

composing with an IGA. For example, grabbing ice particles from a blizzard or picking up and releasing a creature. This approach allows for complex and unpredictable sounds to be triggered by simple, logical, gestures.

4.2 Existing Systems

This section lists systems which have directly informed the development of the MEAS+IGA software. For a review of systems controlled by hand actions, see Fischman (2013).

Waschka (2007) describes a program called *GenDash* which has been used to compose a wide variety of works including compositions for solo human speaker, string quartets, operas and electronic computer music. The program is significantly influenced by evolutionary strategies in nature and reflections on human and animal behaviour such as individuals mating with more than one other individual or with the same individual more than once.

The *Sound Gallery Project* (Woolf and Yee-King 2003) is an interactive sonic art installation that uses a custom built ultrasonic sensing system to track the movements of participants in a gallery space. In the earlier implementations of the *Sound Gallery* this information was used to evolve sound distortion circuits implemented on reconfigurable analogue chips. In later versions, it was used as a physical interface to the *AudioServe* synthesis engine (described below). Four speakers are positioned in an open gallery setting, each one emitting the sound generated by one member of the current generation of genotypes. Participants are encouraged to interact with this gallery space by exploring the soundscape generated by the speakers within it, seeking out those areas they find the most aesthetically pleasing. The ultrasonic sensors track the movements of those present over time, and fitness values for each of the four speakers are derived from this information and passed to a genetic algorithm.

Audioserve (Woolf and Yee-King 2003) is an implementation of a collaborative IGA that allows multiple users to evolve and share audio synthesis circuits using a web based java interface. The audio synthesis circuits consist of several modules placed in a 2D grid. The system employs a genetic algorithm with a distributed population model where transient, small, local island populations evolve and share circuits via a persistent central population. The small island populations are the populations of circuits that the users develop on their machine via mutations driven by interactive evolution. This system provides an example of a means of composing a larger scale work with genetic algorithms, as opposed to designing sounds to be rearranged later (which is a much more common approach in this area).

The *Genophone* system (Mandelis 2001) allows a user to control an IGA using a data glove. However, as the mappings between glove and algorithm are evolved the causal connection between a gesture and the resulting sound is likely to be unclear to an audience. The system is more suitable for sound design and composition than performance. Nevertheless, this system demonstrates the feasibility of controlling an IGA through tracking hand movement and shows that users without formal musical training can rapidly create complex material with this approach.

The *IndagoSonus* (Gartland-Jones and Copley 2003) system uses blocks which are combined to make physical structures (In the form of a 3D graphical display). Each block has the ability to play and compose music: building a physical structure results in building a piece of music. Each block has a home musical phrase which is a phrase associated with that particular block. This phrase can be sent to another block, the second block will produce a new phrase which has a relationship between the phrase it has been passed and its own home musical phrase.

Eigenfeldt (2012) developed a composition system in which musical motives are treated as individuals in a population. The system also uses genetic algorithms to generate macro-level aspects that control how a population is presented. The user can adjust the probability of how operators evolve the population, both individually and collectively. These probabilities can change over successive generations. Once a series of generations have been created, which can be considered the history of the population, the populations are analysed to determine variation over the generations. This analysis is then used to create a trajectory through these generations by means of a genetic algorithm with a fitness function which varies depending on the results of the analysis. It is also possible to combine trajectories, referred to as braiding, for compositional instances, where it is deemed desirable to develop more than one idea during a section of music. The system has successfully generated complete compositions which are representative of the author’s style, yet produce results that are original, musically interesting and surprising. While the approach of using multiple genetic algorithms is applied to the generation of note based material there is no reason why it could not, in principle, be applied to the generation of electroacoustic works.

4.3 Implementation

MAES+IGA *(Manual Actions Expressive System* + *Interactive Genetic Algorithm)* is an extension to Rajmil Fischman's *Manual Actions Expressive System* which enables music creation and performance using natural hand actions such as hitting or shaking virtual objects. A brief overview of MAES is provided below, for a full description and contextualisation of this work see *A Manual Action Expressive System* (Fischman 2013).

4.3.1 Software Overview

MAES (Fischman 2013) consists of a user interface implemented in MAX/MSP[10]. The *P5 Glove* is used to capture the data necessary for the development of convincing gestures. It provides:

- tracking of three-dimensional translation and rotation, and finger bend,
- sufficient sensitivity and speed,
- detection within a wide spatial range calibrated by each user

4.3.1.1 Synthesis and Processing

The synthesis and processing modules are interconnectable by means of a patchbay emulating matrix. The modules were provided by Rajmil Fischman[11] and consist of the following:

1. Sound sources

    1.1 Three synthesisers (including microphone capture)

    1.2 Four audio file players (1–8 channels)

2. Spectral processes[12]

    2.1 Two spectral shifters (Fischman 1997: 134–5)

    2.2 Two spectral stretchers (Fischman 1997: 134–5)

    2.3 Two time stretchers (Fischman 1997: 134–5)

    2.4 Two spectral blur units (Charles 2008: 92–4)

    2.5 Two spectral morph units (Charles 2008: 92–4)

    2.6 A bank of four time varying formants [13]

---

[10] MAES+IGA is implemented using version 2.0 of MAES, which has not yet been released publically. Essentially it contains more processing modules than version 1 and the *P5 Glove* controls significantly more processing parameters (136 version 2 as opposed to 31 in version 1).

[11] For more information on the synthesis and processing modules see (Fischman, 2015).

[12] Implementation using jitter matrices based on Charles (2008)

3. Asynchronous granulation[14]

- This includes the control of sample read position, wander and speed (time-stretch); grain density, duration, transposition and spatial scatter; and cloud envelope for up to four different simultaneous sources.

4. QList Automation

- MAX QLists allow smooth variation of parameters in time according to breakpoint functions.

5. Spatialisation

- A proprietary algorithm implements spatialisation in stereo, surround 5.1 and two octophonic formats, including optional Doppler shift. The matrix patchbay enables the connection of the outputs of any of the synthesisers, file players and processes to ten independent spatialisers. The granulator features six additional spatialisers and the file players can be routed directly to the audio outputs, which is useful in the case of multichannel files that are already distributed in space (Fischman 2015).

---

[13] See http://en.wikipedia.org/wiki/Formant (accessed on 21 November 2017) for more information on formants.

[14] De Poli, Piccialli and Roads (1991: 137–86) explain asynchronous granular synthesis.

In addition to the modules described in the previous section, MAES+IGA has two additional modules which are a further development of those used in ESDE.

1. Interactive genetic algorithm

- An IGA which allows the control of genotype selection, population selection, mutation rate, mutation trigger, crossover trigger, temporary storage trigger, freeze duration and freeze trigger (these are explained below). All the IGA's parameters can be controlled in live performance.

2. Breakpoint functions

- There are sixteen breakpoint functions, which are generated and edited with the IGA's genetic operators (explained below), including individual control of duration, number of breakpoints, looping and trigger mode.

4.3.2 Mapping

MAES+IGA provides a configurable mapping matrix allowing 31 streams of data, 15 from the *P5 Glove* and 16 breakpoint functions from the IGA, to be mapped onto 54 continuous processing parameters, 10 parameters which control the IGA, a physical model for throwing/sowing particles, four audio file triggers, a preset increment and a preset jump. Mappings can be combined simultaneously to generate more complex metaphors using any combination of data from the *P5 Glove* and the IGA. Figures 4.1(a) and 4.1(b) show the mapping matrix windows. This combination of parameter control allows users of greatly varying skill to use the system. Beginners may choose for the parameters to be controlled entirely by the IGA and not interact with them at all and, as they gain a greater understanding of the effect the parameters have on the resulting sound, they may begin to design metaphors which also use a manual control of the parameters.

The interaction between the tracked data from the *P5 glove* and the IGA is handled through a deferred mapping; this is illustrated in figure 4.2. While the streams of tracked data from the *P5 Glove* are continuously mapped onto the parameters (assuming a mapping is activated) the breakpoint functions only provide data when they are triggered. They can be triggered by activating either the granulator's cloud envelope or the four sound file inputs: this is to ensure that it is possible to generate repeatable material. The breakpoint functions can also require conditions to be met before they are mapped onto parameters. Figure 4.3 shows a mapping between breakpoint function 1 and transposition 1 which is conditional on position X being greater than 0.5. The trigger mode for each breakpoint function can be assigned individually. For example, 8 breakpoint functions could be triggered by the granulator's envelope being activated and 8 by an audio file being activated. The first set of 8 functions could be used to generate a slowly evolving granular texture while the second set of 8 functions could generate shorter gestural material through the spectral shifters and stretchers.

Mappings can be stored using the preset mechanism, along with settings for sound synthesis and processing techniques (SSPTs), the IGA and breakpoint functions. Preset increments and jumps can be triggered when conditions are met. This allows for the construction of a composition which can be performed entirely through hand actions.

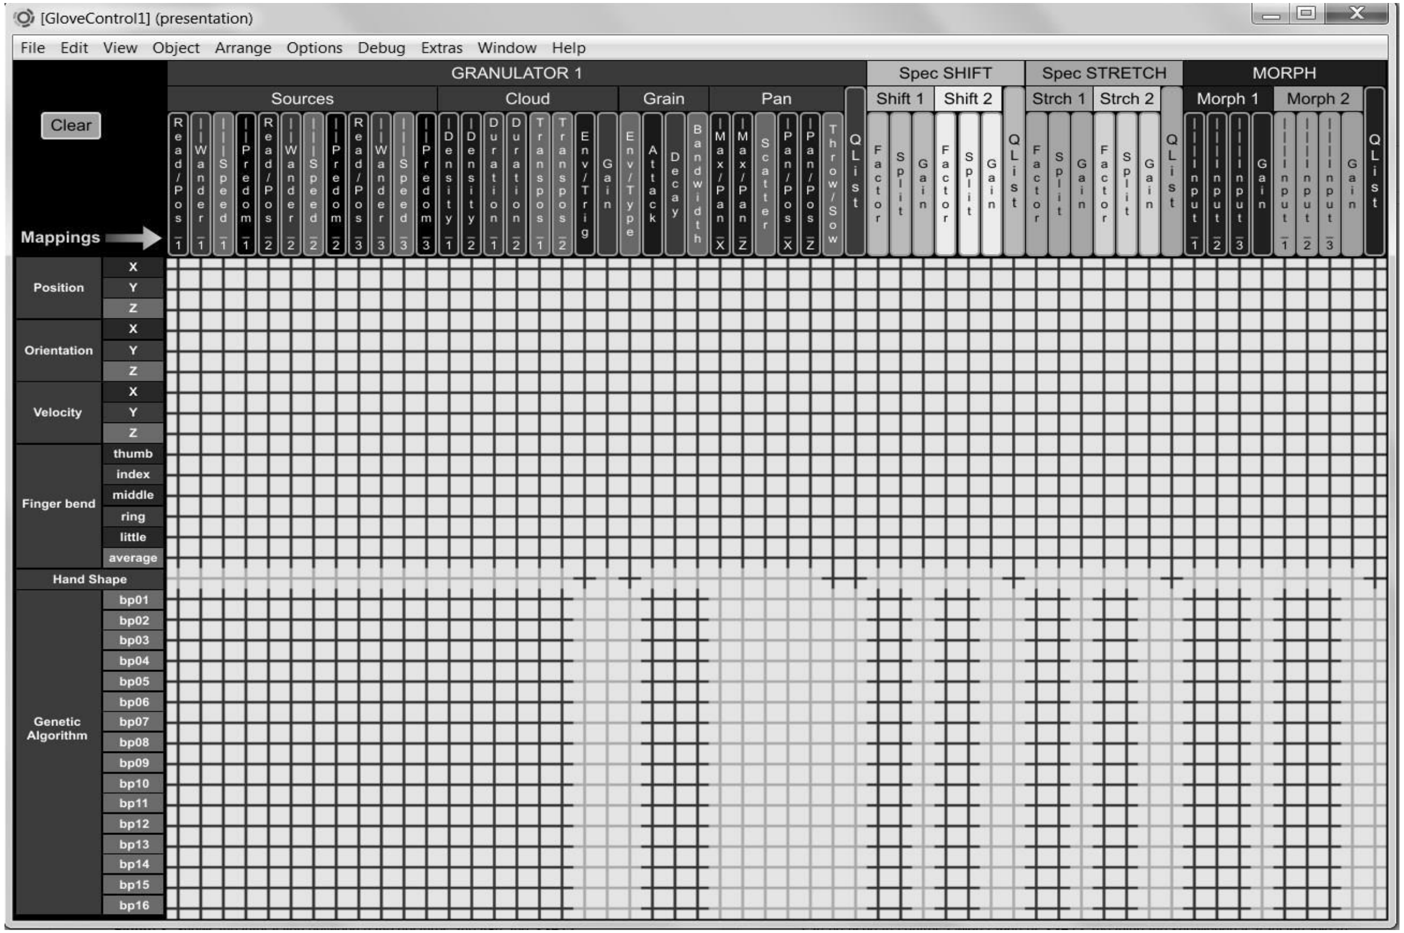

Figure 4.1a: Tracked parameters from the *P5 Glove* and breakpoint functions generated by the IGA appear in the left column. Parameters controlled by the *P5 Glove* and IGA appear along the top row. Fuzzy entries indicate mappings which do not make sense (E.g. the IGA cannot be used to control itself).

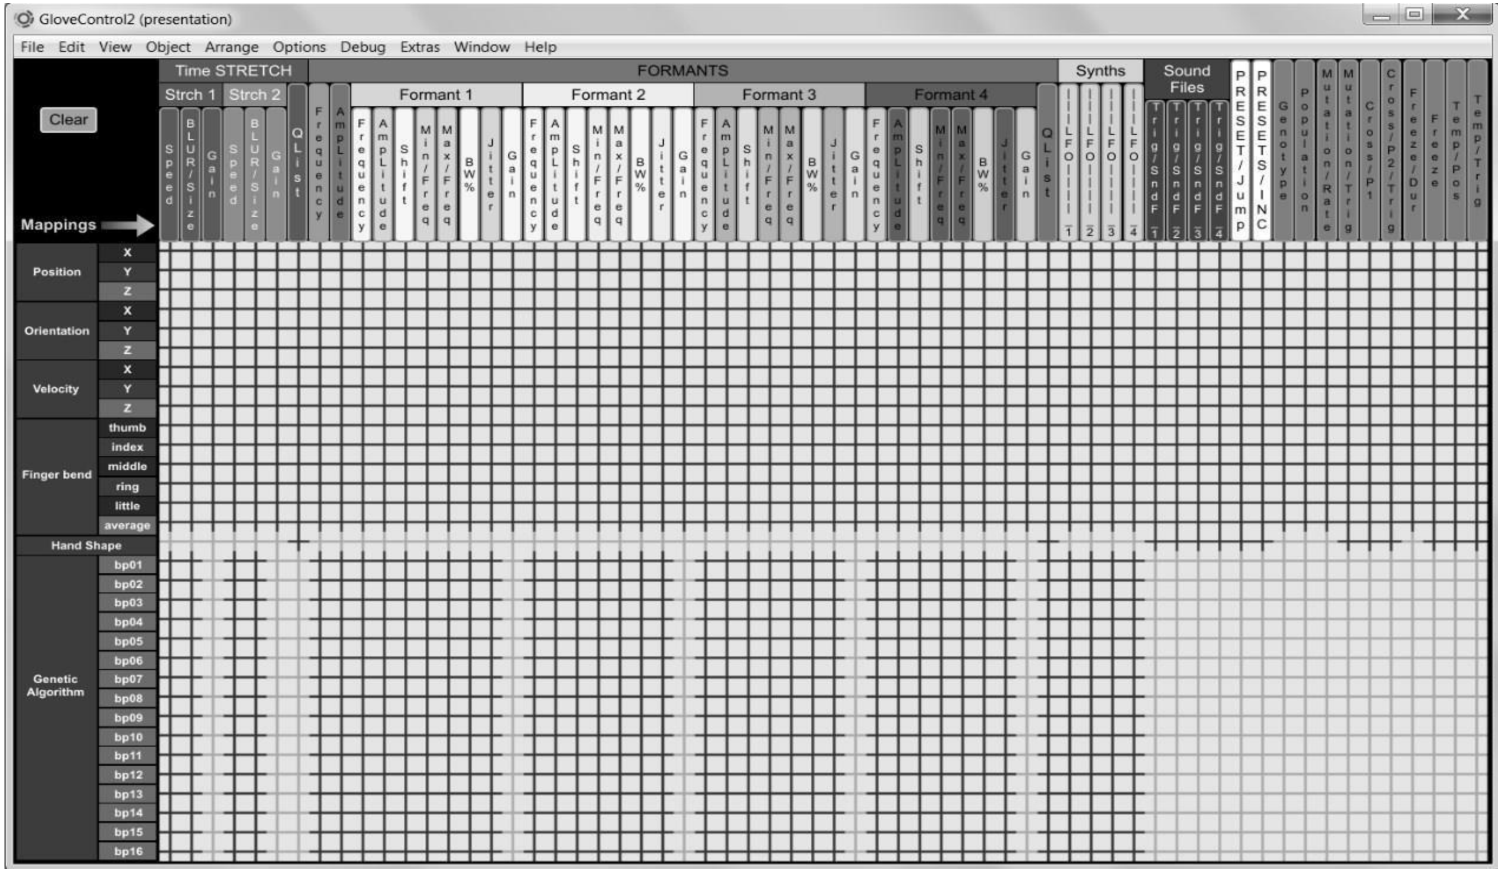

Figure 4.1b: Shows the second mapping window

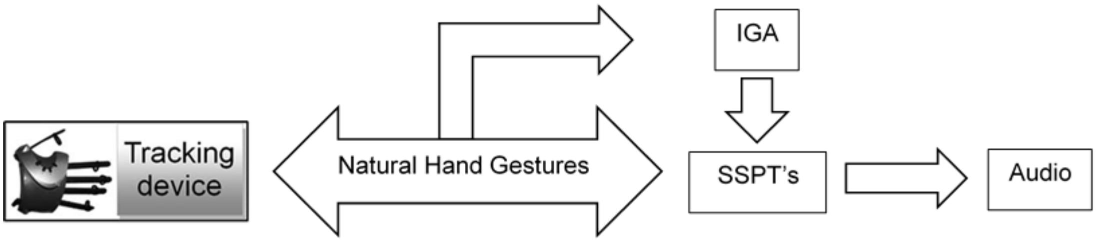

Figure 4.2: Shows the interaction between hand gestures, the IGA and SSPT's

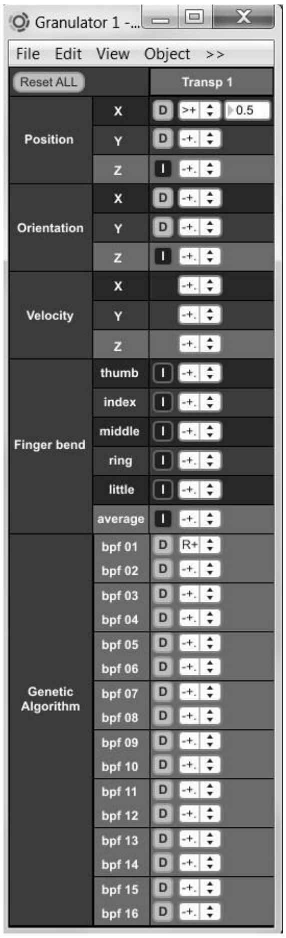

Figure 4.3: Mapping conditions for correspondence between breakpoint function 1 and transposition 1. Data from the breakpoint function will only be mapped when position X is greater than 0.5.

4.3.3 IGA Mapping

As the processing parameters are controlled by the IGA, the user does not need to understand how the SSPTs operate. This means that a beginner can design complex sounds: while it is necessary to acquire a basic knowledge of how an IGA operates, this can be done much more quickly than obtaining a full understanding of the SSPTs. Additionally, the IGA can be used to control a wide range of SSPTs; therefore, knowledge of IGA is applicable to other contexts[15]. The IGA also provides a starting point for a simple compositional structure: material can be developed by creating a child population based on a genotype and new sections in a composition can begin by creating an entirely new population.

The evolutionary process is started by selecting a location to store a new randomly generated population of genotypes which can be stored in one of 32 matrices. This allows multiple child populations to be created while preserving the parent population. Each genotype consists of a bit string[16] which is used to generate a sound through four stages, based on the following approach proposed by Johnson (1999):

1. The bit string is divided into a number of substrings equal to the number of active mappings of breakpoint functions onto processing parameters. For example, if 8 breakpoint functions were to be generated as shown in figure 4.4 the bit string would be divided into 8 substrings.

2. Each substring is then divided into a further set of substrings equal to the number of breakpoints each function is to consist of. For example, if each functions consists of 5

[15] For example, the composition of note based music.
[16] A bit string is a list of 0's and 1's which represent each solution, in this case each one represents a number of breakpoints.

breakpoints, as shown in figure 4.4, then each substring would be further divided into 5 substrings.

3. Each of these substrings is then converted to a decimal number between -1 and 1 and is placed at equally spaced intervals in the breakpoint function. For example, if the breakpoint duration is set to 8 seconds, this would result in a breakpoint being placed at the following intervals:

   0:00, 0:02, 0:04, 0:06 and 0:08.

4. When a genotype is triggered the interpolated values from the breakpoint functions are sent to parameters according to the configuration of the mapping matrix. For example, if the grain transposition mapping was active, the values from the breakpoint function would be applied to the grain transposition parameter.

The number of parameters to evolve is selected before beginning the evolutionary process, according to the configuration of the mapping matrix. This is to avoid redundancy in the genotype. If parts of the latter are not being mapped onto any parameters this will cause unpredictable behaviour in the mutation and crossover procedures since these would be operating on parts of the genotype which have no bearing on the parameters. Figure 4.4 shows an example of evolved breakpoint functions.

4.3.4 Genetic Operators

Genetic operators are used to develop the initial randomly generated material or to generate variations on existing material. This section describes the use of the genetic operators at the composition, as opposed to performance, stage. Figure 4.5 shows the window used to control the IGA. The genetic operators implemented in MAES + IGA are mutation, crossover and temporary storage. The same implementation is used as in the ESDE software, described in chapter 3.

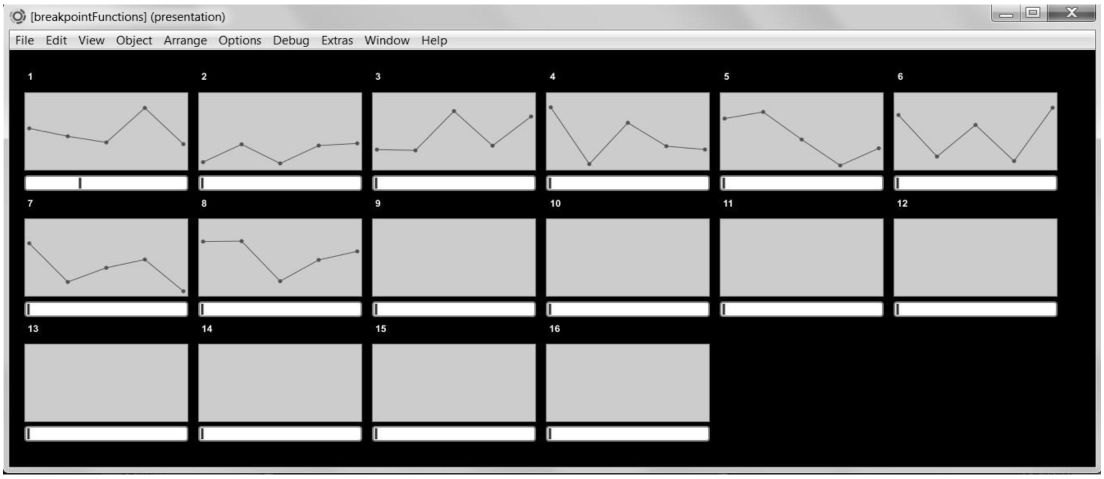

Figure 4.4: An example of breakpoint functions: in this case breakpoints have been evolved for 8 parameters. Each function consists of 5 breakpoints.

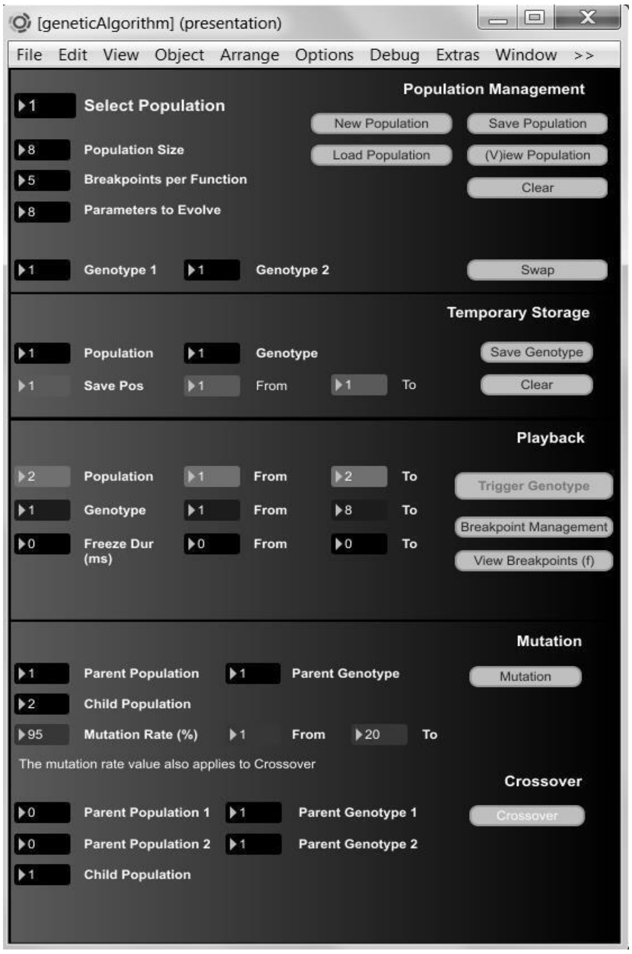

Figure 4.5: The window used to control the IGA.

4.3.5 Live Performance

An approach to using the IGA in live performance is to map position X, Y or Z onto the genotype number, associate the breakpoint triggers with either the granulator envelope or an audio file, and activate the granulator or trigger an audio file when a hand shape is detected. This allows the performer to reach out and pick up a sound as if it were a physical entity. A metaphor which may be helpful to understand this process is to visualise a group of flying creatures in the air in front of the performer which are of the same species but each one varies slightly from the others. Therefore, a different sound will be produced depending on which creature is selected: reaching out in a different position will result in picking up a different creature. Variations on this approach were used to construct the majority of *Snowstorm* and *Terraform*, which are described in chapter 5.

The genetic operators can also be activated in performance, this can be used for constructing sections of a composition which allow for extended improvisations: in such cases, it would usually be necessary to evolve a population as a suitable starting point beforehand. Generating a completely random population in live performance is likely to be too unpredictable[17] and causal connections between the gesture and the resulting sound would be lost. One approach to constructing an improvisation is to evolve multiple populations which consist of genotypes which generate contrasting sounds. For example, one population which generates low grain transposition values and another which generates high values. This can be combined with a preset jump to allow the composition to fork out at various stages.

4.3.6 Causality and Expressiveness

The objective in this system is to retain the level of transparency which can be achieved in MAES while using material generated with the IGA; there should be a clear causal connection between a gesture and the resulting sound, which is more challenging when using the IGA due to the level of randomness involved. For example, imagine a sound which is defined by its pitch rising over its duration. If this were triggered by moving the hand upwards this action would be logical to an audience. However, if a child population was created based on the genotype which generates this sound, it is possible that the pitch would instead fall over its duration for some members of the child population which would mean that the gesture used to trigger it becomes illogical. One way to overcome this is by shifting the focus of the metaphors from sound sculpture to selection: the genotypes are treated as life forms or natural processes and their spectromorphology is assumed to be fixed[18] , the performer may pick them up and move them around but not fundamentally change them. For example, imagine attempting to contain a

[17] It is possible to use a randomly generated population if tight constraints on parameter ranges are selected, although this will limit the variety of material which can be generated.
[18] With the exception of spatial location, and gain in some cases.

buzzing insect in your hand, it can be moved around and released in different directions, but the sound resulting from its movements cannot be affected in any substantial way[19].

Due to the control of the parameters being deferred to the IGA there is the risk of losing the expressive aspect present in the original system. This is overcome by controlling some parameters through hand movements. For example, if the density, duration, and transposition of grains is fixed, the performer may choose how to locate them in space and control their amplitude. This combination of parameter control and IGA independent evolution is logical in practice as hand gestures can be used to control those parameters which an IGA is not suited to. For example, an IGA is unsuitable for controlling gain and spatialisation, as it will result in what is essentially a random fluctuation of these parameters which will in most cases be undesirable.

When using the IGA in live performance an additional parameter is provided, genotype freeze. When this is activated the breakpoint functions are held at their current values for a time equal to their freeze durations. This is analogous to placing the hand under running water: it will temporarily be held back but will rapidly find its way round and continue to flow.

4.4 Conclusion

This chapter presented the MAES+IGA system which allows a user without formal musical training to compose and perform music with the aid of an IGA. The use of time varying breakpoint functions in conjunction with the mapping matrix allows complex and varied material to be developed.

[19] This is a personal preference that proved to be effective during performance; other users of the system are free to develop their own approaches to the construction of gestures.

The creation of *Snowstorm* and *Terraform* (see Chapter 5) demonstrates that the use of IGA's is a viable approach to composition and live performance. As *Snowstorm* and *Terraform* can be performed using simple hand gestures it is hoped that this system will allow a greater number of people to engage in the process of composition and performance, particularly those who may have previously considered it inaccessible.

References

Charles, J.F. (2008). A Tutorial on Spectral Sound Processing Using Max/MSP and Jitter. *Computer Music Journal* 32(3), pp. 87–102.

Dahlstedt, P. (2007). Evolution in creative sound design. In E.R. Miranda and J.A. Biles (eds.) *Evolutionary Computer Music*, London: Springer, pp. 79–99.

De Poli, G., Piccialli, A. and Roads, C. 1991. Representation of Musical Signals. Cambridge, MA: The MIT Press, pp. 137-86.

Eigenfeldt, A. and Pasquier, P. (2012). Populations of Populations: Composing with Multiple Evolutionary Algorithms. In: P. Machado., J. Romero. And A. Carballal (eds.) *Evolutionary and Biologically Inspired Music, Sound, Art and Design*, Heidelberg: Springer, pp. 72–83.

Fischman, R. (1997). The Phase Vocoder: Theory and Practice. *Organised Sound* 2(2), pp. 127–45.

Fischman, R. (2013). A Manual Actions Expressive System. *Organised Sound* 18(3): 328-345.

Gartland-Jones, A. and Copley, P. (2003). The Suitability of Genetic Algorithms for Musical Composition. *Contemporary Music Review,* 22(3), pp. 43-55.

Johnson, C.B. (1999). Exploring the sound space of synthesis algorithms using interactive genetic algorithms. In A. Patrizio., A.G., Wiggins. And P. Pain (eds.) *Proceeding of the AISB 99 Symposium on Musical Creativity.* Brighton, UK: Society for the Study of Artificial Intelligence and Simulation of Behaviour, pp. 182-196.

Mandelis, J. (2001). Genophone: an evolutionary approach to sound synthesis and performance. In: E. Bilotta., E. R. Miranda., P. Pantano., P. Todd (eds.) *Proceeding ALMMA 2001: Artificial Life Models for Musical Applications Workshop*. Consenza, Italy: Editoriale Bios, pp. 37-50.

Waschka, R., (2007). Composing with Genetic Algorithms: GenDash. In E.R. Miranda and J.A. Biles (eds.) *Evolutionary Computer Music*, London: Springer. pp. 117-136.

Yee-King, M. and Woolf, S. (2003). Virtual and Physical Interfaces for Collaborative Evolution of Sound. *Contemporary Music Review,* 22(3), pp. 31-41.

Chapter 5 – MAES + IGA Compositions

5.1 Introduction

Snowstorm and *Terraform* demonstrate approaches to constructing a composition using an interactive genetic algorithm (IGA) controlled by hand gestures with the MAES+IGA system. The compositions allow the performer to interact with the organisms and natural processes found in a journey through surreal landscapes. The metaphor used for the construction of the compositions is evolution in a changing environment, this is based on the principle of the properties of an environment affecting the organisms which will evolve within it. Each section defines the current conditions of the environment and the performer interacts with the organisms and processes present within this environment. The compositions are pre-composed and repeatable; they are recognisable as the same composition with each performance although they allow scope for improvisation.

5.2 Generating Material

Three mechanisms are used to perform the compositions:

1) Triggering breakpoint functions pre-evolved by the IGA through the performer gestures. 2) Controlling continuous parameters through performer gestures to retain expressiveness in the performance.
3) Triggering pre-composed material

5.2.1 Triggering evolved breakpoint functions

Triggering breakpoint functions evolved with the IGA provides the majority of material directly controlled by the performer. The populations are constructed in such a way to allow for some improvisation. For example, the section of *Snowstorm* beginning at 02:05 [20] which uses the metaphor of gathering and throwing 'ice cubes' consists of 8 genotypes with a range of grain density, grain duration and grain transposition values. These options, alongside the manual control of gain and spatialisation, allow this section to be performed in a variety of ways. For example, it can be performed in a subdued or more intense manner.

5.2.2 Pre-composed material

Pre-composed material is also triggered at various points throughout the composition. This constitutes the environment: it is possible to interact with organisms and the processes present in an environment but not change the environment itself. There is some overlap between triggering pre-composed material and triggering of genotypes which represent natural processes. For example, at 02:50 in *Snowstorm* the performer reaches down and triggers a rising bubbling process, which can be seen as analogous to water rising from a hot spring.

5.2.3 Fully Automated Processes

Fully automated processes caused by initiating MAX QLists are less frequently necessary than in the original version of MAES as some of their functionality is replaced by the IGA. They are primarily used to add an additional layer to material generated by the IGA. For instance, prior to the downbeats. They also provide the fade out at the end of the composition.

---

[20] As this is a live composition, the timings will vary with every performance. The timecodes in this chapter refer only to those on the recorded performance.

5.3 Compositional Strategies

The strategies used to compose the fixed media works[21] were used as a starting point for the development of IGA based compositional strategies suitable for live works.

5.4 Snowstorm

5.4.1 Programme Notes

Snowstorm takes a performer on a journey through a surreal winter landscape filled with strange organisms and natural processes. The sonic material alludes to blizzards, ice, hot springs and otherworldly creatures. Throughout the journey a snowstorm begins and grows in intensity then gradually fades away.

The material consists of two main types. The first is fixed material representing the environment which cannot be changed by the performer. The second represents the organisms found within the landscape which respond to the performer's actions causing sounds to move, evolve and scatter across the landscape. Each section defines rules for the evolution of sounds and processes, which represent the impact of environmental conditions on the evolutionary process.

The performer interacts with the landscape through natural hand gestures, which act as an interface to an interactive genetic algorithm controlling a range of sound synthesis and processing techniques. The composition has been designed to provide a high level of causality between gestures and the resulting sonic material which allow interactive genetic algorithm based strategies to be used in live performance.

[21] The compositional strategies used to compose the fixed media works (*The Blind Watchmaker* and *The Singing Forest*) are described in chapter 3.

5.4.2 Structure

Snowstorm consists of four main sections

Section 1 – 00:00 – 02:00

Section 1 introduces the winter landscape, the performer can interact with the ice particles in a blizzard and the surreal creatures within it. The section is concluded with the metaphor of water rising from a hot spring.

Section 2 – 02:01 – 04:14

Section 2 allows the performer to further interact with ice particles, they can be scattered across the landscape, flicked and shattered.

Section 3 – 04:15 – 06:23

The performer's gestures create a layered wind texture. The section is brought to a climax with a return of the ice scattering effect and the rising water metaphor.

Section 4 – 06-24 – 08:45

The composition concludes with a return to the blizzard scene introduced in section 1.

5.4.3 Compositional Strategies

5.4.3.1 Strategy 1

The performer reaches out and grabs a granular cloud representing an organism or natural process. The performer can hold the particles together and move them around the room or disperse the particles in space. Depending on the X position, a different genotype will be triggered. In this case, the variation in the genotypes affects the grain density, grain duration and grain transposition values. This demonstrates an approach to controlling processing parameters with a combination of the IGA and hand gestures.

An example of this strategy can be seen at 00:19 where the performer simulates grabbing ice particles in a blizzard. The mappings used to create this effect are shown in figures 5.1(a) and 5.1(b).

The mappings are as follows:

1. Breakpoint functions are mapped onto: grain density 1 & 2, grain duration 1 & 2, grain transposition 1 & 2, spectral shift 1 factor and spectral stretch 1 factor.
2. X Position is mapped divergently onto the genotype selection and panning of the sound on the X axis
3. Y Position is mapped onto gain.
4. Z position is mapped onto the panning of the sound on the Z axis.
5. The cloud envelope is activated on recognition of a hand shape (closed fist).
6. The breakpoint functions are triggered when the cloud envelope is activated.

A second example of this effect can be seen at 01:00 where the performer simulates picking up and releasing a creature found in the landscape, the same mapping as above is used. However, a different sample has been loaded into the granulator’s buffer and the genotypes generate a different range of values for the grain density, grain duration and grain transposition parameters.

5.4.3.2 Strategy 2

The performer flicks a particle causing it to bounce.

An example of this strategy can be seen at 02:37 where the performer flicks an ice particle. Figures 5.2(a) and 5.2(b) show the mappings for this example. Figure 5.2(c) shows the conditions required for this mapping to be activated. The mappings are as follows:

1. Breakpoint functions are mapped onto: grain density, grain duration, grain transposition, formant amplitude, formant frequency.
2. X Position is mapped onto genotype selection.
3. The cloud envelope is activated when finger bend is below 0.3.
4. The breakpoint functions are triggered when the cloud envelope is activated.

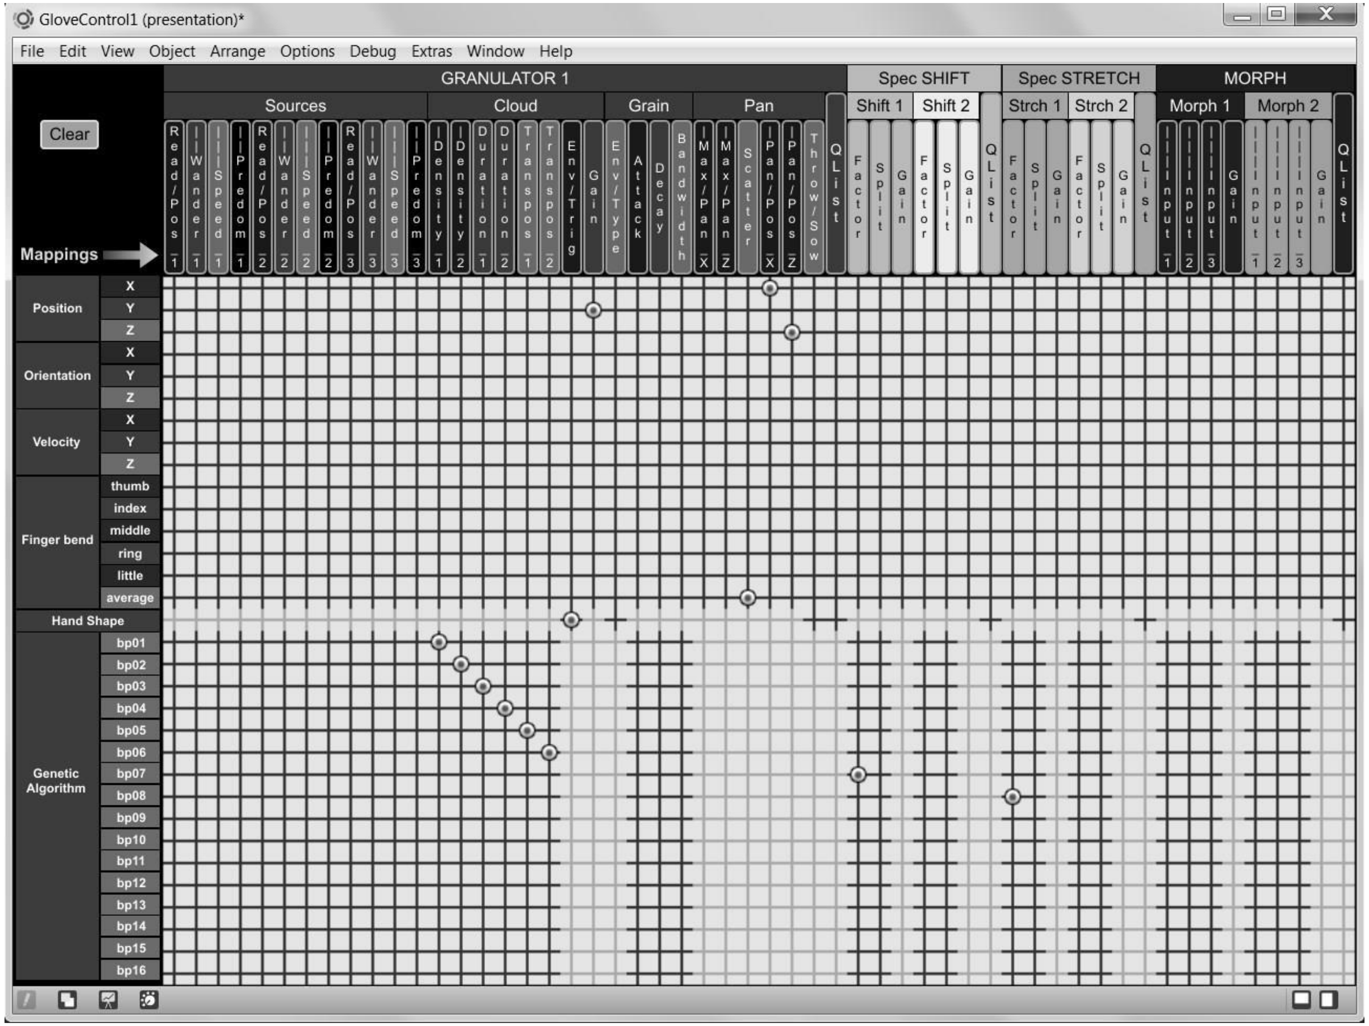

Figure 5.1a: Mapping used to create the blizzard effect.

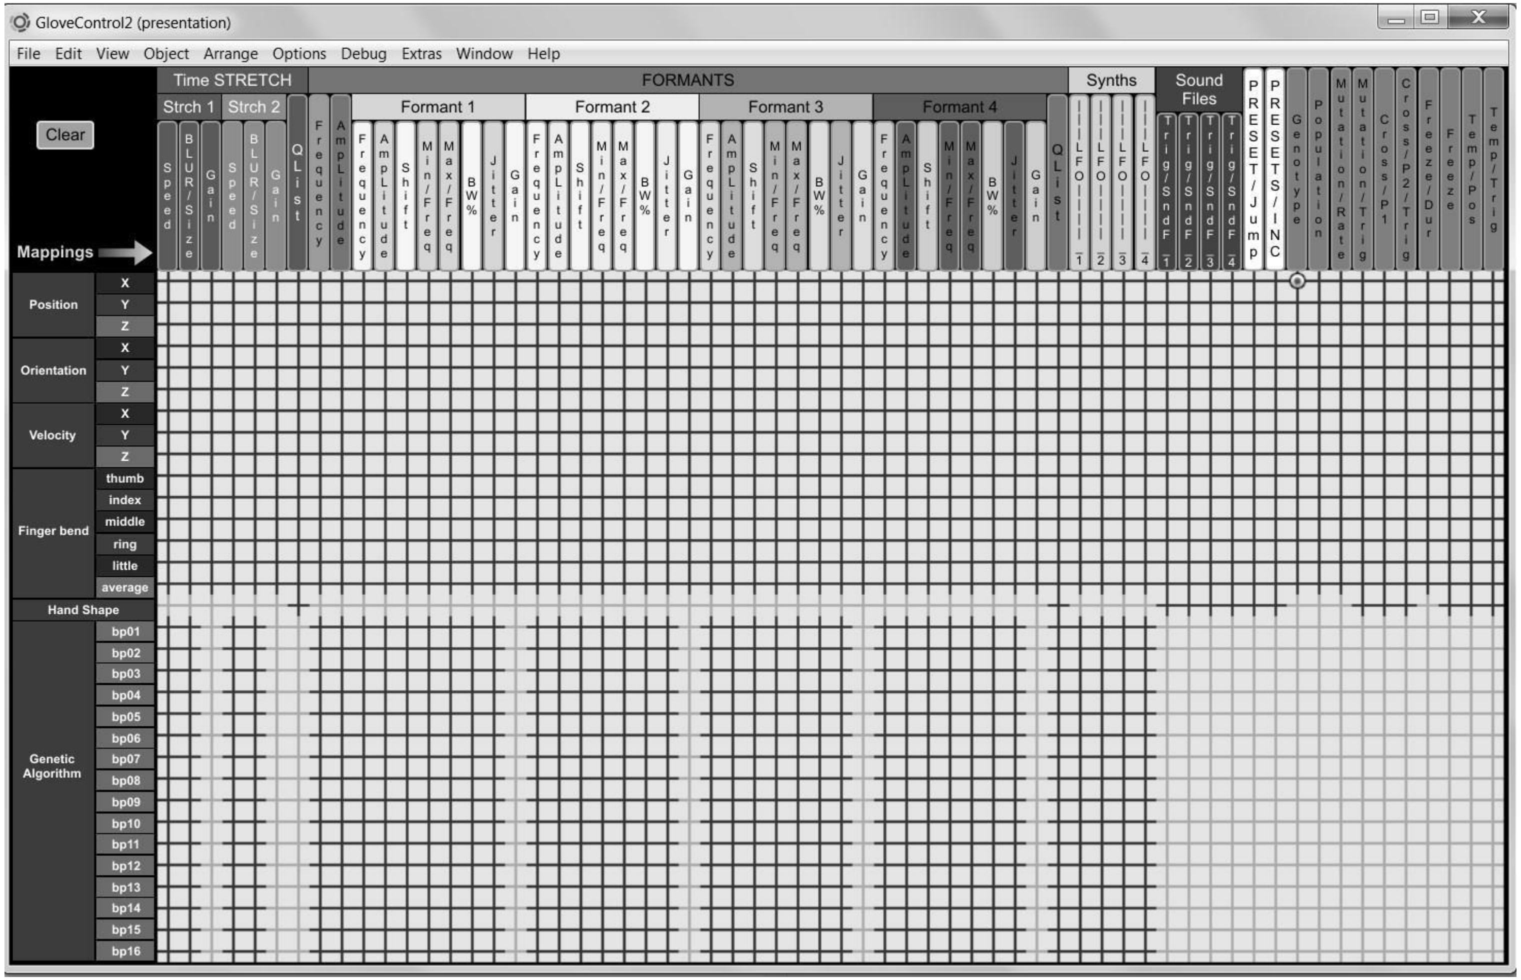

Figure 5.1b: Mapping used to create the blizzard effect.

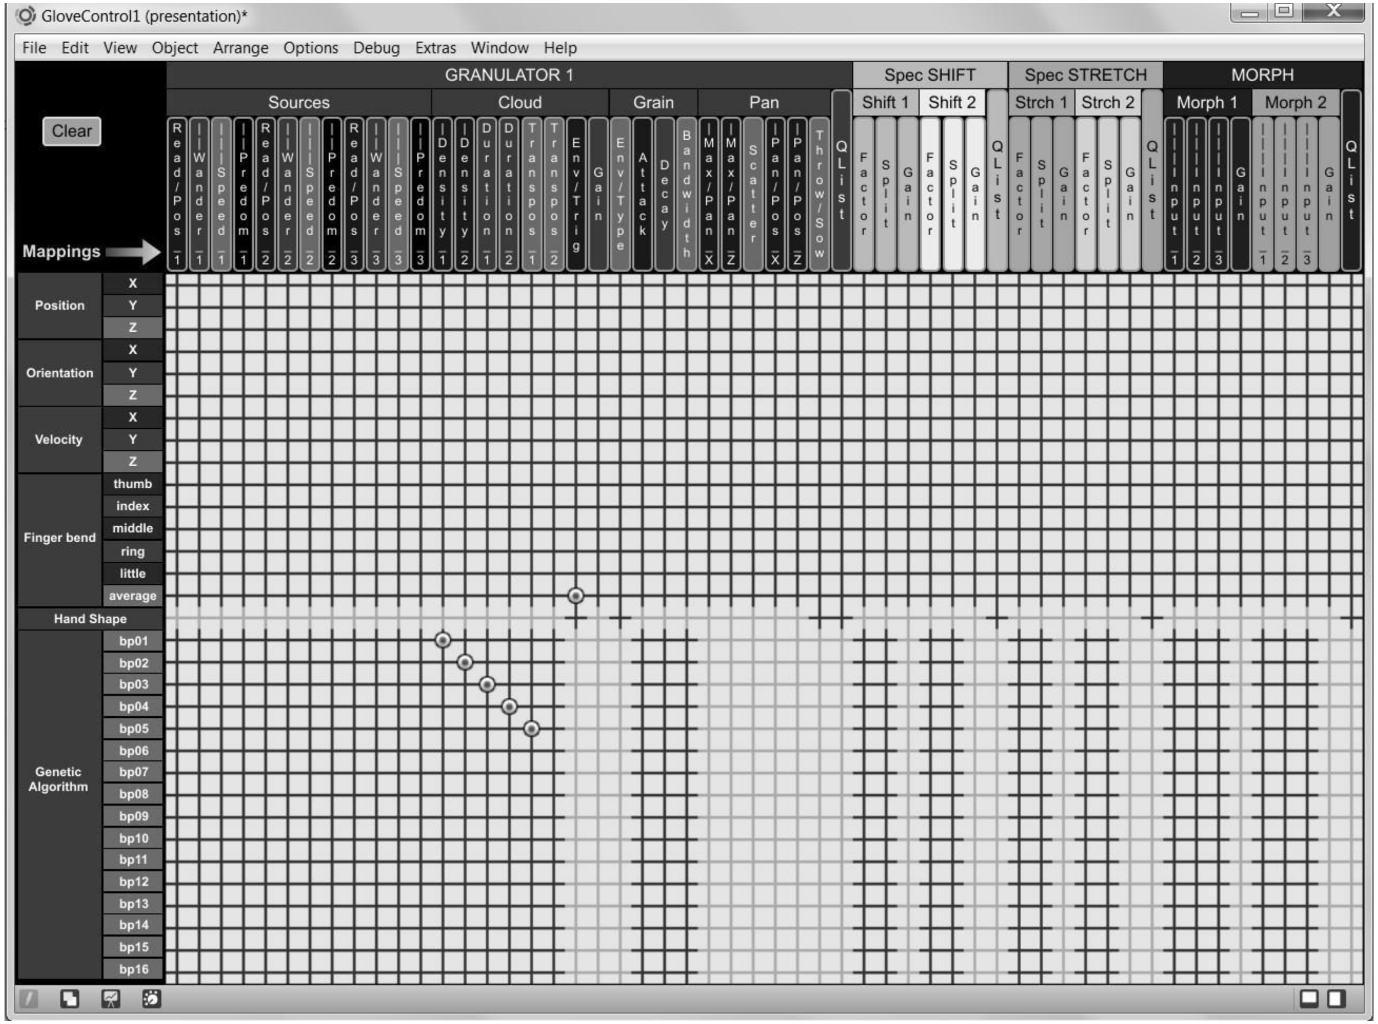

Figure 5.2a: Mapping used to create flicking ice effect

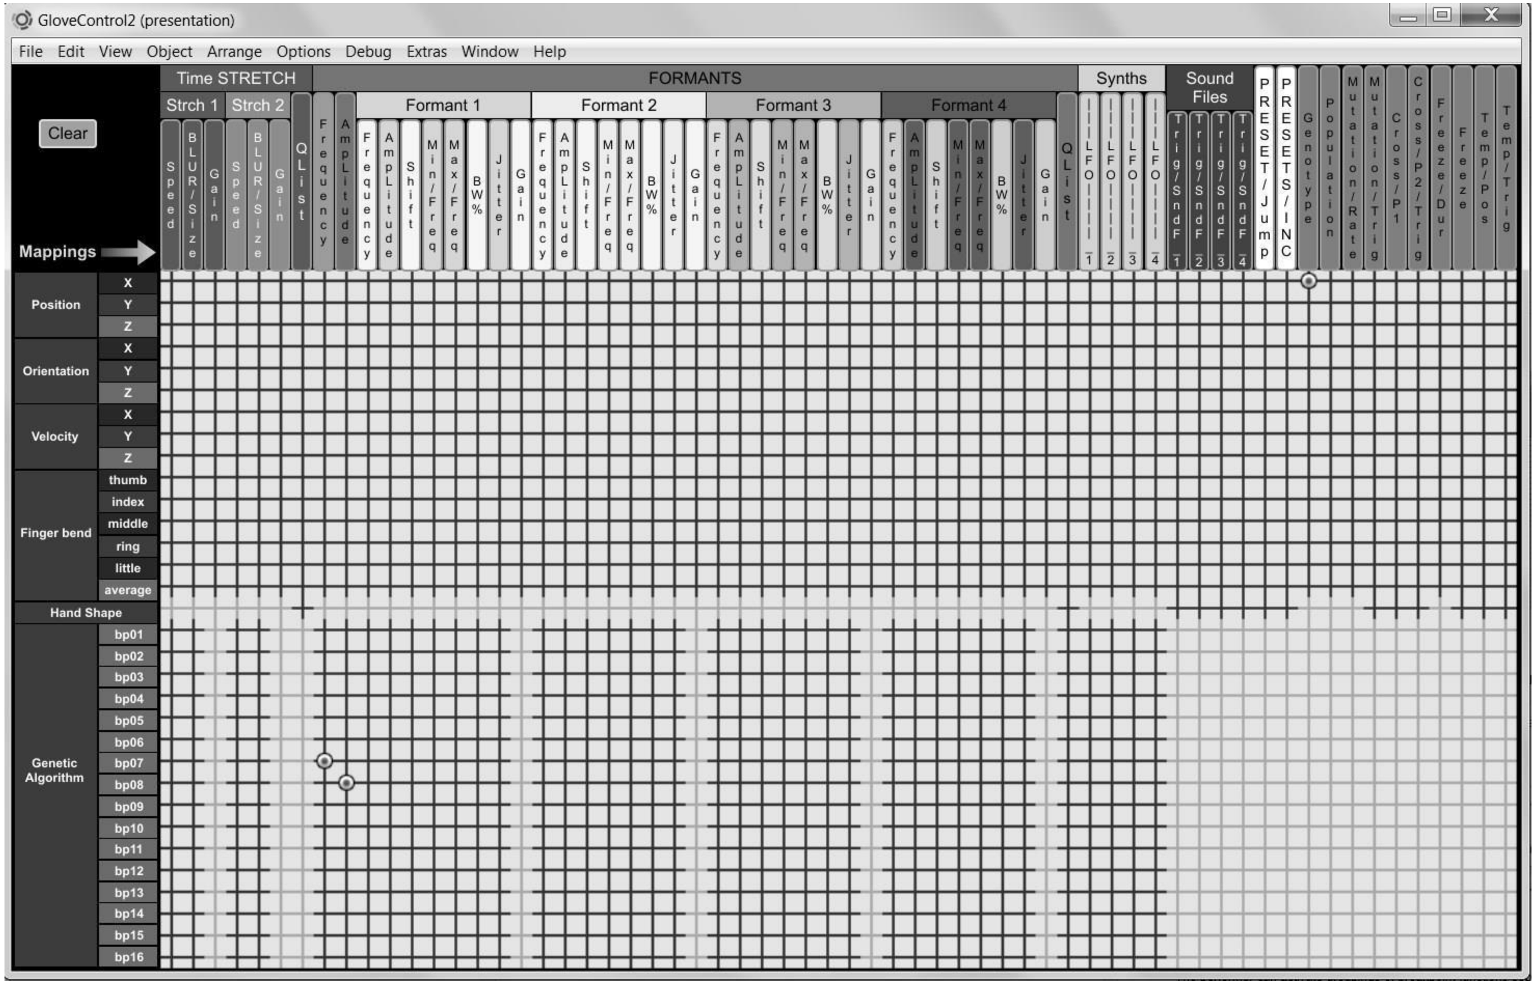

Figure 5.2b: Mapping used to create flicking ice effect.

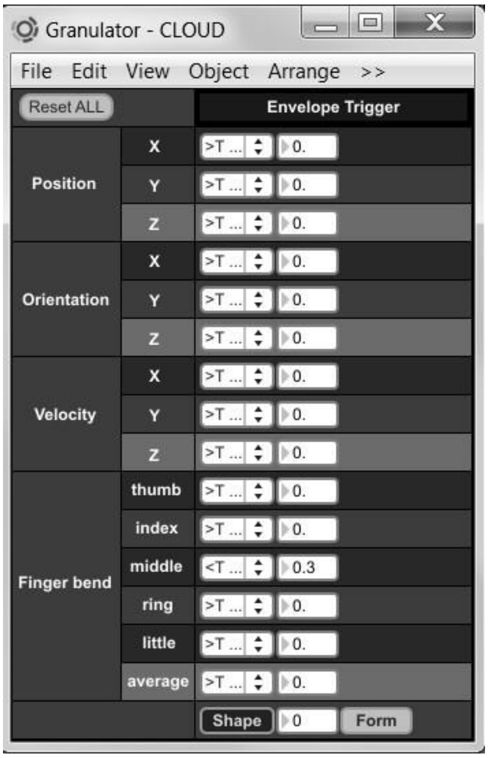

Figure 5.2c: Conditions for ice flicking mapping to activate

5.4.3.3 Strategy 3

The performer's hand movements create layered textures.

An example of this strategy can be heard at 04:15 where the performer's creates a layered wind texture. This strategy demonstrates the importance of ancillary gestures (Fischman 2013: 340). The mappings are as follows.

1. Breakpoint functions are mapped onto: spectral shift 1 & 2 factors, spectral shift 1 & 2 split frequencies, spectral stretch 1 & 2 factors, spectral stretch 1 & 2 split frequencies. 2. X Position is mapped onto genotype selection

3. Audio file 3 is activated on recognition of a hand shape (closed fist).

4. Audio file 4 is activated on recognition of a hand shape (closed fist, index finger extended).
5. The breakpoint functions are triggered when audio file 3 is activated.

5.5 Terraform

5.5.1 Programme Notes

Terraform allows a performer to use natural hand gestures to influence the formation of a number surreal landscapes. The composition begins with more abstract sound material to provide opportunities for the performer to improvise and work with a diverse range of hand gestures and metaphors for controlling sound processes. Despite the wide range of gestures available to the performer, a strong sense of causality to the resulting material is retained. As the piece progresses, there is a shift towards more recognisable material and the central metaphor of landscape construction. The performer progresses through a series of landscapes with varying conditions. The landscapes are covered in earth and rocks, filled with wind and water, and populated with a range of lifeforms. *Terraform* can be seen as a loose telling of the story of a number of evolutionary processes. Possibly processes which diverges from those seen on Earth, yet which all follow a journey from simplicity to increasing complexity and intelligence.

5.5.2 Structure

Section 1 - 00:00 – 01:33

The first section places the performer in a surreal landscape, they can interact with a range of aural material[22]. The section is brought to a conclusion by triggering multiple granular clouds with an increasing grain density.

Section 2 – 01:34 – 04:06

This section begins a transition towards mimetic material, the performer interacts with wind chimes and controls wind like textures. The section concludes with a brief transition back towards unrecognisable, aural material.

Section 3 – 04:07 – 06:41

This section initially continues with the move back towards unrecognisable, aural material then quickly reverses this trend. The performer gradually constructs a landscape by scattering rocks and water across it, then introducing life forms into it.

Section 4 – 06:42 – 10:11

This section concludes the composition by returning to material from the previous to sections and combining the aural and mimetic material heard throughout them

5.5.3 Compositional Strategies

5.5.3.1 New Strategies

New strategies were developed while composing *Terraform* beyond those used in *Snowstorm*.

[22] See Emerson (1986).

5.5.3.1.1 Strategy 1

The performer can activate groupings of breakpoint functions separately to control multiple sound streams independently.

An example of this strategy can be seen at 08:34. A texture is generated by controlling the granulator and formants with the first breakpoint function group. Audio file 1 (wind chimes) can then be independently triggered and controlled with the second breakpoint function group. Figure

5.3 shows the settings of the breakpoint management window used to create this example. The mappings are as follows:

1. Breakpoint functions 1 to 6 are mapped onto: grain density 1 & 2, grain duration 1 & 2, grain transposition 1 & 2.

2. Breakpoint functions 9 to 12 are mapped onto: formant 1 – 4 amplitude.

3. Breakpoint functions 13 to 16 are mapped onto: choruser speed, choruser depth, harmoniser transposition 1 and harmoniser transposition 1.

4. The duration of breakpoint functions 1 to 12 is 60 seconds.

5. The duration of breakpoint functions 13 to 16 is 5 seconds.

6. The granulator's cloud envelope is activated on recognition of a hand shape (closed fist, index finger extended).

7. Audio file 1 is activated on recognition of a hand shape (closed fist).

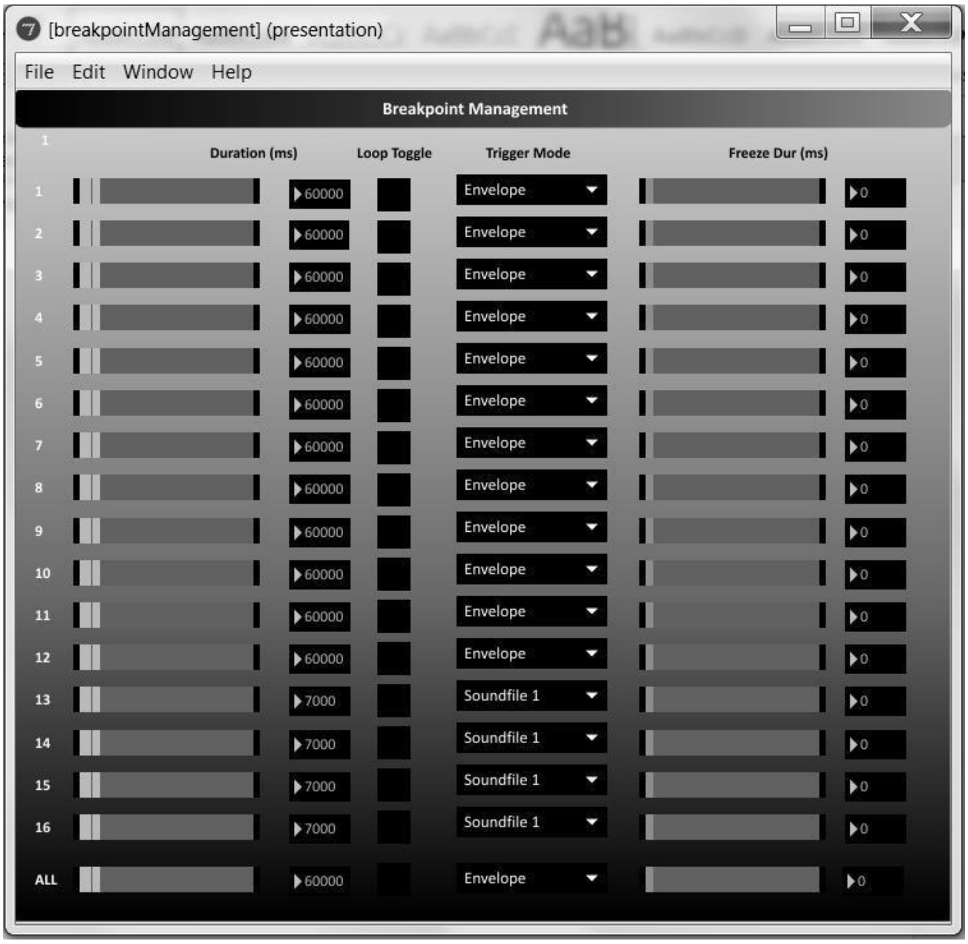

Figure 5.3: Breakpoint groupings

5.5.3.1.2 Strategy 2

The performer's hand movements simulate sweeping through wind chimes.

An example of this strategy can be heard at 02:00. The granulator is activated when the performer opens their hand and accelerates it along the X axis. The mappings are as follows:

1. X Velocity is mapped onto grain density 1.

2. Average finger bend is mapped onto grain density 2.

3. Breakpoint functions 1 and 2 are mapped onto grain duration 1 and grain duration 2.

4. The cloud envelope is activated when the average finger bend is below 0.3.

5.5.3.2 Reapplied Strategies

Strategies applied during the composition of *Snowstorm* were reapplied to *Terraform,* with some variations.

5.5.3.2.1 Strategy 1

This strategy was used to simulate the performer grabbing ice particles in a blizzard. It has been reapplied in *Terraform* to simulate grabbing electrical particles, which can be seen at 00:18. The process was created by controlling 2 spectral shifters and the harmoniser instead of the granulator.

The mappings for this are as follows:

1. Breakpoint functions 1 - 6 are mapped onto: spectral shift 1 & 2 factor, spectral shift 1 & 2 split and harmoniser transposition 1 & 2.

2. X Position is mapped onto the genotype selection and panning of the sound on the X axis

3. Audio file 1 is activated on recognition of a hand shape (closed fist).

6. The breakpoint functions are triggered when audio file 1 is activated.

5.5.3.2.2 Strategy 3

Strategy 3 is applied in *Snowstorm* to creating a layered wind texture, this process is reapplied in *Terraform* using the formant bank. This can be seen at 02:31. The mappings are as follows**:**

1. Breakpoint functions 1 – 4 are mapped onto formant shift 1 – 4.

2. X Position is mapped onto formant frequency.
3. Y Position is mapped onto formant amplitude.
4. The formant input is provided by the granulator.
5. The breakpoint functions are when the grain envelope is activated.
6. The grain envelope is activated on activated on recognition of a hand shape (closed fist).

5.6 Conclusion

Snowstorm and *Terraform* demonstrate how the MAES+IGA system can be used to compose and perform live compositions. MAES+IGA builds on the initial developments made in the ESDE software by allowing the IGA to serve as an interface between hand gestures and SSPTs. It has allowed multiple live compositions to be created which consist of diverse and complex material which can be controlled in real time. The compositional strategies which have been developed have overcome the issue of the lack of causality between gesture and resulting material when performing with an IGA. This has been achieved through the development of new metaphors in which the performer simulates interacting with organisms and processes in a structured environment.

References

Emmerson, S. (1986). The Relation of Language to Materials. In S. Emmerson (ed.) *The Language of Electroacoustic Music*. Basingstoke: Macmillan, pp. 17-39

Fischman, R. (2013). A Manual Actions Expressive System. Organised Sound 18(3), pp. 328-345.

Chapter 6 – Visualising Interactive Genetic Algorithms

6.1 Introduction

This chapter describes a system for generating an Interactive Genetic Algorithm (IGA) visualisation. IGAs allow for the construction of a simple, intuitive interface to SSPTs which is more accessible to a complete beginner or a composer who has not worked extensively with technology before (Johnson, 1999). This is possible because it is not necessary to directly control all the parameters of the SSPT, only a small number of the IGA parameters need to be controlled. This principle has been applied to the development of an interactive installation in the form of a visualisation which represents the state of an IGA. A user can also interact with the visualisation to change the latter and control the generation of musical material. The visualisation also functions as an educational

Gakki-Mon Planet (Dahlstedt and Berry 2003) is a more developed version of *Feeping Creatures*. It features more complex graphics and allows a greater deal of user interaction. The user can travel around the world using a joystick, and interact with and listen to the creatures. The creatures can be placed near a food source which provides energy and increases the probability of their survival, or caused to mate (crossover). This level of interaction can enhance the users' understanding of the process as they can experiment with various creatures and use the camera to track their movements. They can also move the camera to track the effect of an individual in a population.

The genotypes determine the appearance and behaviour of the creatures, both visual and musical.

The *IndagoSonus* system (Gartland-Jones and Copley 2003) uses blocks which are combined to make physical structures (in the form of a 3D graphical display). Each block has the ability to play and 'compose' music; building a physical structure results in building a piece of music. Each block has a home musical phrase which can be sent to another block, the second block will produce a new phrase which has a relationship between the phrase it has been passed and its own home musical phrase, and so on.

A number of visualisations have also been constructed in the field of computer science. However, they tend to presume the user already possesses a relatively high level of technical knowledge which is less likely to be the case among musicians. Also, these visualisations tend to serve primarily analytical purposes and are more suited for genetic algorithms which operate over thousands or millions of generations whereas IGAs are often run for less than ten (Ying-Hong Liao 2001).

6.3 - Objectives

The visualisation should allow the user to rapidly gain an understanding of the principles the IGA operates on. While demonstrating these principles, the system should produce varied, high quality sound material so as to be suitable for an installation at a public event.

To ensure a user gains a sufficient understanding of the operation of an IGA, the animations convey the following information:

- The currently selected population and genotype
- The relationship between the genotypes in a population and the level of variation from the parent population
- The creation of a new population

6.4 Implementation

The visualisation system consists of two components. The first component is a modified version of the *Evolutionary Sound Design Environment* (ESDE) software, described in Chapter 2. The second component is the Processing [23] (Processing, 2001) visualisation. The visualisation shows an underwater scene inhabited by different species of surreal creatures.

6.4.1 ESDE Extension

The ESDE software has been modified in the following ways. The modified version of ESDE is shown in figure 6.1.

6.4.1.1 SSPT changes

The Cycling 74 FM Synthesis module has been added to allow for the generation of sequences of notes. Also, the Comb Filter bank has been replaced by the Formant Filter Bank used in MAES+IGA (described in Chapter 4), which generates material with a closer correspondence to the animations used in the visualisation.

[23] Processing is an open source programming language developed primarily for graphics programming

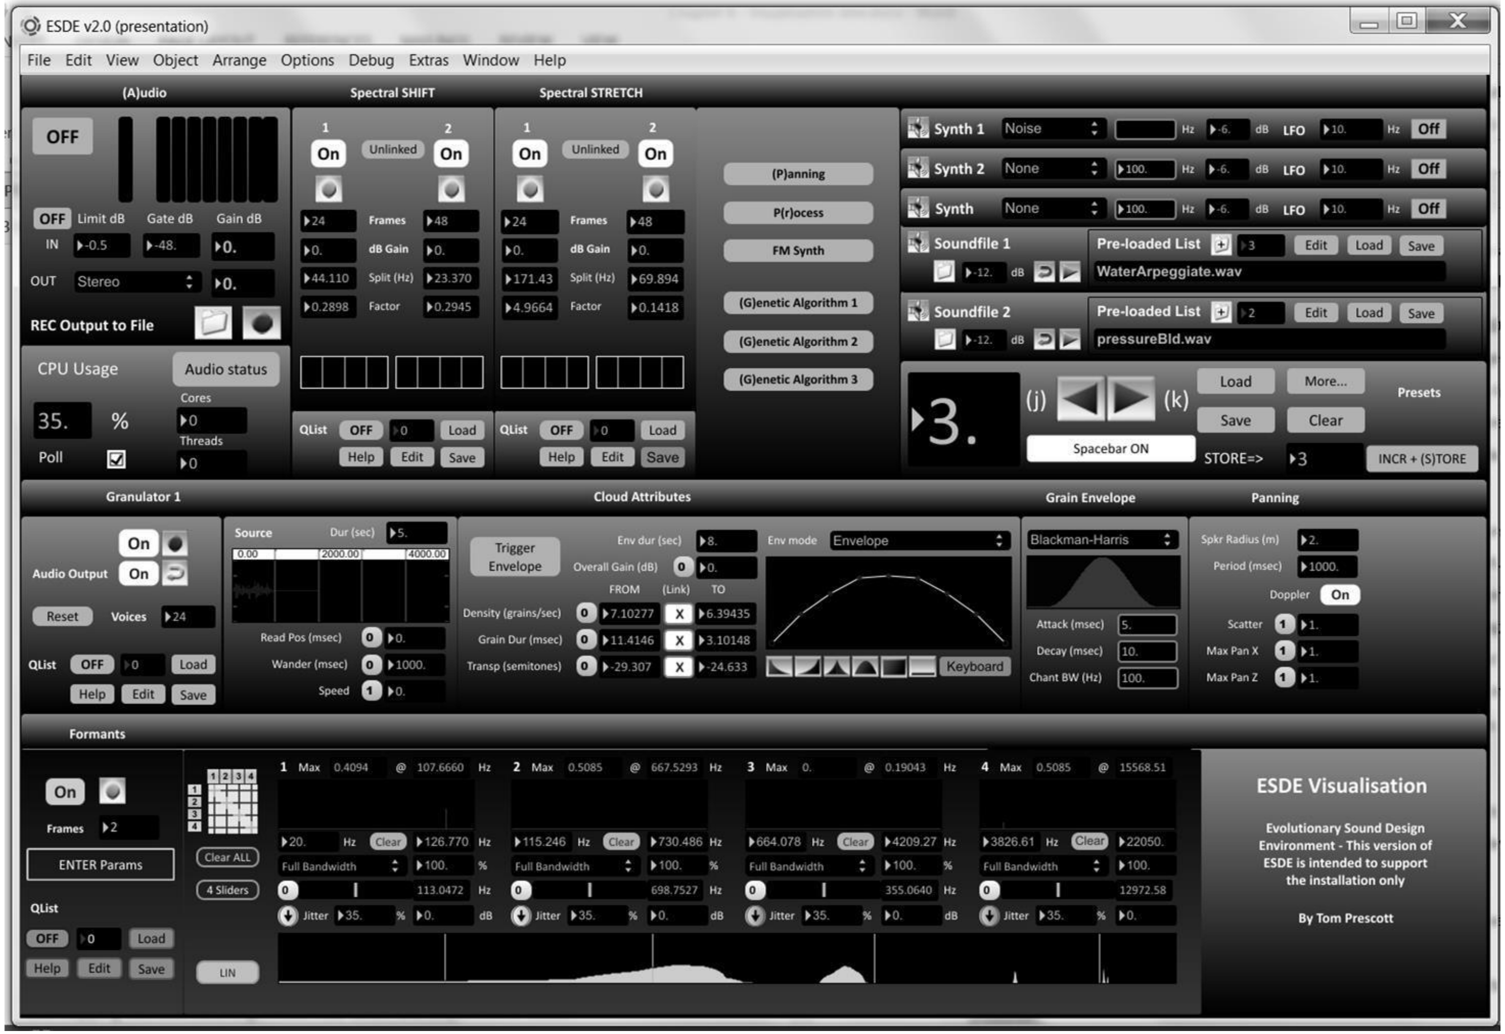

Figure 6.1: The modified ESDE software

6.4.1.2 Multiple Interactive Genetic Algorithms

This version of *ESDE* contains three IGAs which can be independently controlled, as opposed to the single one present in the original software.

6.4.1.3 RGB Mapping

The mapping was extended to generate RGB values and sequences of notes. The same values are mapped onto both SSPT parameters and the animations. This creates the correspondence between the creatures and the sounds they generate. Creatures which look similar, will generate similar sound material.

6.4.1.4 FM Synthesis Mapping

The MIDI data for the FM synthesis is generated by mapping the IGA output onto pitch, velocity and duration.

6.4.1.5 MAX/Processing Communication Module

A module was added to allow for communication between MAX and the Processing Visualisation. RGB values are sent from MAX to processing which determine the creature's appearance. Values are sent from Processing to MAX to control the state of the IGA and to activate the SSPTs. This communication is implemented with the Open Sound Control (OSC) protocol which was designed for sending data between musical instruments and computers (Wright and Freed 1997).

6.4.2 Visualisation

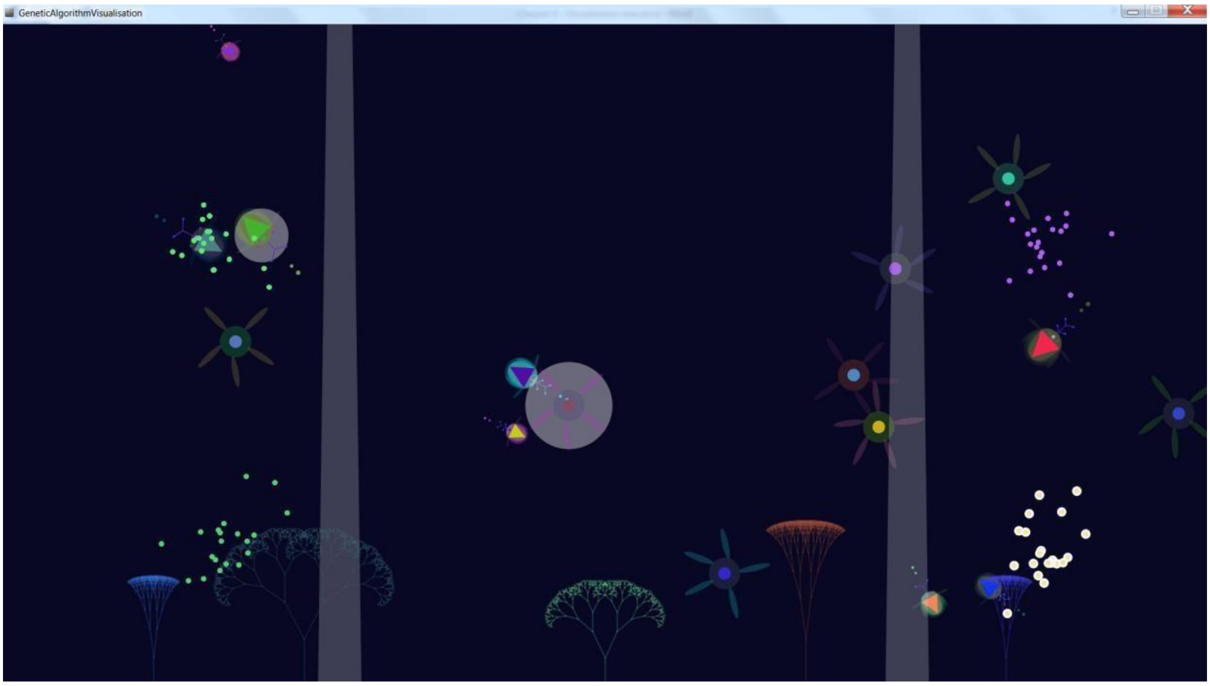

Figure 6.2: The visualisation

The visualisation is implemented in Processing. Genotypes are represented as creatures which move on a 2D plane. The level of genetic diversity between members of a population and their parents is conveyed using an approach inspired by Gakki-Mon Planet (Dahlstedt and Berry 2003); in which the genotype determines the physical appearance of the creature as well as the sound it will trigger. Creatures which look similar, will trigger material which sounds similar. Animated creatures were chosen for the visualisation since IGAs are modelled on the evolution of living organisms. This should allow a user to rapidly understand the principles which an IGA operates on.

There are three types of creatures: 'insects', 'starfish' and 'plankton'. The shared behaviour of the 'insects' and 'starfish' are defined in the Creature Processing class[24]. This defines the following behaviour of a creature:

[24] See https://en.wikipedia.org/wiki/Object-oriented_programming (accessed 8 January 2018) for more information on object oriented programming and its terminology.

- Position
- Movement speed
- Rotation
- Fade in and fade out animation
- Animation when creature is playing associated sound
- Growth rate when new child population is created
- Colour changes when new child population is created

The Creature class is then extended to implement the SeaCreature and Starfish classes which provide the customised animations required for each creature. The 'plankton' are implemented through the ParticleSystem class. Each instance of a class draws a single creature. When the program is started, 2D arrays are created for each creature type which stores all the instances of each creature. The array rows represent populations and the columns store the individual creatures within each population. The main draw loop increments through each of these arrays to animate the creatures.

Each type of creature is mapped to one of the three IGAs in ESDE. The visualisation is shown in Figure 6.2. The fractal plants and light beams were added to provide some additional visual interest to the installation. They are not IGA controlled and instead trigger audio files when clicked on. This can be seen as an extension of the evolution in a structured environment metaphor introduced in chapter 5.

The visualisation implements sends and receives OSC messages from the ESDE application. The mousePressed and keyPressed functions allow the user to interact with the visualisation by listening for mouse and keyboard input. They allow the user to trigger the sounds associated with each creature, trigger the sounds associated with the environment and trigger the mutation process in ESDE. When a creature is clicked on, the genotype is triggered in ESDE. When a user has not interacted with the system for 30 seconds, genotypes are randomly selected to be automatically triggered to allow the system to function as an installation.

The oscEvent function listens for OSC messages passed from ESDE which generates the RGB values for the creatures. A new message is passed when a new child population is created by triggering the mutation process.

6.4.2.1 Visualising IGA state changes

The initial random population generates creatures with a wide range of colours. The user can select one of the creatures and create a new child population based on it. This child population inherits some visual similarities to their parent, but with some random variation provided by the mutation process to prevent the population converging. Mutation is represented by creatures from the new population emerging from the current location of the parent followed by the parent population fading out. The new generation is initially 50% smaller then gradually grows in size to demonstrate the parent-child relationship. Figure 6. 3 shows the new generation animation.

Figure 6.3: New population animation. Arrow 1 points to the parent, arrow 2 points to the child population fading in, at this point the fade is in progress so the colours are only partially developed.

6.5 Installation

The installation operates on a 9 minute loop. Every 3 minutes, the preset number is incremented and new parameters and audio files are loaded. This represents the changes over time in an ecosystem due to environmental changes or natural disasters. On a practical level, it retains the interest of the user, prevents the installation from becoming repetitive and avoids the problem of populations converging to an extent that the material is too similar.

The default mapping configuration is as follows:

- The 'insects' trigger genotypes in IGA 1. This is mapped to MIDI note generation, and the parameters of the FM synthesis module.
- The 'starfish' trigger genotypes in IGA 2. This is mapped to the Spectral Shifter and Spectral Stretchers.

- The ‘plankton’ trigger genotypes in IGA 3. This is mapped onto the granulator.
- The ‘fractal trees’ trigger audio file 1.
- The ‘light rays’ trigger audio file 2.

The duration of the loop, time between presets, audio files and FM synthesis scales can all be easily changed so that the installation can be customised to be suitable for different venues and occasions.

The user interacts with the system by clicking on a creature, this will map the IGA parameters for the creature’s corresponding genotype onto an SSPT. Users can also trigger a mutation process to create a new generation of genotypes. Additional audio file triggers allow the users more scope for improvisation.

As a number of the settings are fixed by the presets, users are provided with a structured environment in which they can exercise a level of creativity while being able to rapidly produce satisfactory results.

6.6 Conclusion

This chapter presents an approach to visualising an IGA in the form of an interactive installation.

The visualisation should rapidly provide users with an understanding of core principles which an IGA operates on and allow them to generate a wide range of varied material. The system also demonstrates how, with the aid of an IGA, a range of sophisticated SSPTs can be controlled by a simple interface. The simple interface means the system is accessible to non-expert users, such as children and adults who are familiar with gaming. This also provides an opportunity to present electroacoustic music in an accessible form to an audience who may otherwise not be exposed to it. The system is also highly customisable, which means users can make adjustments to the former as they become more familiar with the IGA.

References

Berry, R., Dahlstedt, P. (2003) Artificial Life: Why Should Musicians Bother? *Contemporary Music Review* 22(3), pp. 57 – 67.

Dahlstedt, P. (2001). A MutaSynth in Parameter Space: Interactive Composition through Evolution. *Organised Sound* 6(2), pp. 121-124.

Gartland-Jones, A., Copley, P. (2003). The Suitability of Genetic Algorithms for Musical Composition.

Contemporary Music Review 22(3), pp. 43-55.

Johnson, C.B. (1999). Exploring the sound space of synthesis algorithms using interactive genetic algorithms. In *Proc. of the AISB '99 Symposium on AI and Musical Creativity*. Brighton, UK. Society for the Study of Artificial Intelligence and Simulation of Behaviour, pp. 182-196.

Processing. (2017). *Processing*. [ONLINE] Available at: https://processing.org/. [Accessed 16 May 2016].

ScaleMaker. (2011). *ScaleMaker*. [ONLINE] Available at: https://cycling74.com/tools/harmony-maker. [Accessed 16 May 2016].

Wrigt, M., Freed, A. (1997) Open Sound Control: A New Protocol for Communicating with Sound Synthesizers: International Computer Music Conference. Thessaloniki, Greece.

Ying-Hong Liao., Sun. (2001). An Educational Genetic Algorithm Learning Tool. *IEEE Transactions on Education* 44(2) pp. 20-40.

.

Chapter 7 – Multi-channel Large-Scale Composition

7.1 The Wind-Up Bird

This chapter discusses *The Wind-Up Bird*, a large scale fixed media composition with both a stereo and octophonic version[25]. *The Wind-Up Bird* can be seen as a culmination of the work carried out previously. The lessons learned by returning to the fixed medium, and composing a large-scale work, are also applicable to live composition and performance.

The objectives of this composition were to:

1) Apply the more sophisticated sound synthesis and processing techniques (SSPTs) developed for the MAES+IGA software[26] to a fixed media composition. These include the addition of a bank of four IGA controlled time varying formants, and an IGA controlled choruser and harmoniser.
2) Investigate the versatility of IGA as a generator of sonic material of sufficient variety for the composition of a large scale work on Interactive Genetic Algorithm (IGA) based compositional strategies.

7.2 Programme Notes

The Wind-Up Bird is an investigation into composing a large-scale work with material generated by an interactive genetic algorithm and arranged with interactive genetic algorithm based compositional strategies. It is the culmination of investigations into the application of interactive genetic algorithms to sound design, fixed media composition, live performance and installations.

The composition explores the simultaneous evolution of multiple populations and the coexistence and divergence of populations throughout this process. This can be seen as analogous to the

---

[25] There are some slight variations between the stereo and octophonic versions of *The Wind-Up Bird*. The timecodes in this chapter refer to the stereo version. [26] The MAES+IGA software is discussed in chapter 4.

development of natural ecosystems. The landscapes which support the development of these systems are also explored. The landscapes and processes are represented with a more minimal recourse to mimetic material. Instead, the evolutionary process is explored through the development of musical material and the relationships between these materials.

7.3 Structure

Section 1 – 0:00 – 01:02

Section 1 begins the composition with the introduction of abstract granular material then begins a gradual transition to more recognisable bird song. It makes use of a larger number of IGA populations than present in *The Blind Watchmaker* and *The Singing Forest* due to the increased opportunity to develop the material allowed by the large-scale structure. The section concludes with crossfades between aural [26] (non-recognisable) materials generated by multiple populations to create the transition towards recognisable mimetic material.

Section 2 – 01:03 – 05-17

Section 2 continues to introduce material which will be developed throughout the composition. Firstly, a more pronounced filtered birdsong completes the transition suggested in section 1. A low pitched gesture which serves as a counterpoint to the birdsong is also introduced. The section is developed by layering variations of the filtered birdsong and concludes with a return towards abstract granular material introduced in section 1.

Section 3 – 05:18 – 09-40

Section 3 makes a departure from the previous events by moving away from the granular material and introducing multiple noise based textures. This section is brought to a climax by layering the noise based textures.

---

[26] See Emmerson 1986.

Section 4 – 09:41 – 14:40

Section 4 further develops the noise based material introduced in section 3. The low pitched gesture introduced in section 2 returns to provide counterpoint to new granular material and prepares a transition back to the material from sections 1, 2 and 3 which conclude the composition.

Section 5 – 14:41 – 18:11

Section 5 provides the climax of the composition. It reintroduces the abstract granular material from Section 1 and repeats the transition towards recognisable material. Material presented throughout the composition is combined, including the filtered birdsong and the noise based textures introduced in section 2. The layering strategy from section 2 is also repeated to conclude this section. However, additional layers were added to achieve higher event density to emphasise the climax.

Section 6 – 18:11 – 20:05

Section 6 concludes the composition with a combination of the filtered bird song and noise based textures which lead to a fade out.

7.4 Compositional strategies

7.4.1 New Strategies

New compositional strategies were developed while composing *The Wind-Up Bird* beyond those used for *The Blind Watchmaker* and *The Singing Forest*. The new strategies were possible due to the advancements in IGA based SSPT control developed for the MAES+IGA software and because of the opportunities provided when composing a large scale work.

7.4.1.1 Strategy 1

Multiple populations can be concurrently developed throughout a section as a large scale composition allows time for the development of a wider range of material, as follows:

1) Generate multiple populations.

2) Create child populations based on each parent to develop the material generated by each population.

3) The child populations provide multiple lines of development throughout a section.

An example of this strategy can be heard throughout section 2, beginning at 01:02. The section begins by developing the granular material. A low pitched gesture which is generated by a second population is introduced at 0:34. This is followed by the introduction of a third population which generates material suggestive of birdsong at 02:02. Each of these three populations serves as a parent to create a new child population. This allows the three lines of material to be developed concurrently. The development of the low pitched gesture can be heard at 02:22, the birdsong at

02:43 and the granular material at 04:15

A second example can be heard beginning at 11:50: the low pitched gesture is again developed alongside granular material. A second occurrence of the low pitched gesture can heard at 12:21 and the developed granular material can be heard at 12:40.

7.4.1.2 Strategy 2

To take advantage of the opportunities provided by the large scale structure, populations were developed across multiple sections of the composition, as follows:

1) Introduce a population and develop it by creating child populations based on the parent.

2) Reintroduce the population in a later section and create further child populations based on the parent for further development

An example of this approach can be heard in the granular material developed throughout sections 1 and 2. It is first introduced at 00:12 until 01:33. It is then reintroduced and further developed at 04:15.

A second example can be heard in section 2 and section 5. The noise based gesture introduced in section 2 at 05:16 is reintroduced and developed further in section 3 from 10:40 onwards.

7.4.1.3 Strategy 3

This is a variation of a strategy applied while composing *The Singing Forest*[27] consisting of crossfades between material generated by multiple genotypes from a population. However, it differs from the latter because the crossfades are between materials generated by genotypes from different populations as opposed to between genotypes generated by developments of a single population. The strategy is implemented as follows:

1) Generate multiple populations and develop them by creating child populations based on the parent.

[27] *The Singing Forest* is discussed in chapter 3.

2) Use crossfade or spectral morphing to combine genotypes from different populations to create longer, more complex, gestures.

An example of this strategy can be heard at 00:42 where this approach is used to create the transition from abstract granular material towards more recognisable birdsong. Crossfades between genotypes from four different populations were applied to create this transformation. This strategy is also applied to the conclusion of section 3 and section 5 which begin at 07:42 and 16:43, respectively. It complements the layering of noise based textures from multiple populations by crossfading between them.

7.4.1.4 Strategy 4

This strategy is used to increase the complexity of the layered material which is used to conclude section 2.

1) Generate and develop multiple populations.

2) Select genotypes from different populations and use them to process the same audio file.

3) Layer the output audio files.

This approach is applied towards the end of section 2 at 03:36: the same audio file is processed by both the formant filter bank and the harmoniser. Each SSPT is controlled by a genotype from a different population.

7.4.2 Variations on Previous Strategies

In addition to the new strategies developed, some strategies developed while composing *The Blind Watchmaker* and *The Singing Forest*[28] were reapplied and, in some cases, further developed to take advantage of new features in the MAES+IGA software.

7.4.2.1 Strategy 5

This strategy was originally used to provide significant developments in material while retaining some connection to the original. It was reapplied in *The Wind-Up Bird* to generate more complex layering of material which is used to conclude sections 2, 3 and 5 of this composition, as follows:

1) Select a genotype used previously in the composition, developed through any other strategy.

2) Adjust the mapping in the mapping matrix to apply the same breakpoint function to a different parameter.

An example of this approach can be heard at the conclusion of section 3 and section 5, from 08:21 onwards and 16:48 onwards. The same breakpoint was used to control both a spectral shifter and spectral stretcher to create the layered noise based textures.

7.4.2.2 Strategy 6

Multiple forms of crossover were applied to generate the granular material in Section 4. This was combined with crossfading between genotypes, as follows:

[28] *The Blind Watchmaker* and *The Singing Forest* are discussed in chapter 3.

1) Select two genotypes used previously in the composition, and developed through any other strategy.

2) Create a new child population with the two genotypes using one point crossover and apply a low mutation rate ca. 10%.

3) Use two genotypes generated in the previous step as a parent and create a new child population with uniform crossover with a low mutation rate ca. 10%.

An example of Strategy 6 can be heard at 10:27. The granular material was generated by applying one point crossover to the genotypes which generate the previous granular texture heard at 10:06 and 10:15. The genotype which generates the material at 10:27 was combined with another genotype from the same population to generate the granular texture heard at 11:23. A crossfade between these two genotypes was then applied to develop this material.

7.5 Conclusion

The Wind-Up Bird demonstrates that IGAs can be used to control a wide range of SSPTs. This approach has allowed material of sufficient diversity and complexity to be generated to complete a large scale composition. In addition, it has provided an opportunity to develop a wide range of IGA based compositional strategies suitable for different scenarios such as fixed media compositions, live performance and installations. Along with *The Blind Watchmaker* and *The Singing*

Forest, The Wind-Up Bird demonstrates that IGA based techniques can be repeatedly applied to meet a range of compositional requirements. This is possible due to the flexibility of the IGA to SSPT mapping systems implemented in the software. As breakpoint functions are evolved independently from the SSPT parameters, the software can be customised depending on the current requirements. The replacement of the comb filter bank used In the composition of *The Blind Watchmaker* and *The*

Singing Forest with the time varying formant bank, the harmoniser and the choruser, allowed for a composition with different spectromorphological characteristics and avoided the risk of the compositions becoming repetitive.

References

Emmerson, S. (1986). The Relation of Language to Materials. In S. Emmerson (ed.)

The Language of Electroacoustic Music, Basingstoke: Macmillan, pp. 17-39.

.

Chapter 8 – Conclusion

8.1. Overview

During the course of this Ph.D. approaches have been developed which apply IGA based strategies to sound design, composition and performance. The aim has been to overcome the limitations of existing systems and to extend IGA based techniques into new areas. The developments made can be categorised according to the grouping of the three substantial items which make up the Ph.D.

8.2 ESDE Software and Fixed Media Compositions

The *ESDE* software overcomes the trade-off which is traditionally made between the diversity of material which can be generated by a system and ease of use. This limitation was overcome by evolving time varying breakpoint functions independently of the SSPT parameters. The combination of the processing and mapping matrix allow for a wide range of complex material to be generated and controlled by an IGA without the need for customisation. However, if the user should wish to customise the software, the independently evolved breakpoint functions make this process possible by simply enabling the routing of breakpoint functions into a selection of SSPTs and assigning custom ranges which the breakpoint values should be mapped onto. For instance, this was demonstrated by exchanging the comb filter bank used in the composition of *The Blind Watchmaker* and *The Singing Forest* with the chorus, harmoniser and time varying formants used in the composition of *The Wind-Up Bird*.

The Blind Watchmaker, The Singing Forest and *The Wind-Up Bird* demonstrate how the *ESDE* software is capable of generating a range of material of sufficient diversity and complexity to create multiple, varied, fixed media compositions. These compositions also demonstrate how IGA based strategies can be extended into the development of compositional strategies, as opposed to just the generation of material. The core strategies were developed during the composition of *The Blind Watchmaker* and *The Singing Forest*. These strategies culminated in the composition of *The Wind-Up Bird,* a multichannel, large-scale work.

8.2.1 Further Developments

The number of SSPTs which can be controlled in real time is limited by current CPU speeds: as these increase a larger number of SSPTs can be included in the software which will respectively increase the diversity of material which can be generated and further reduce the need for customisation.

8.3 MAES+IGA Software and Live Compositions

The MAES+IGA software applies IGA based techniques to live performance by mapping hand gestures onto SSPT parameters through an IGA interface. A limitation of previous gesture controlled IGA systems concerned the lack of causality between a gesture and the resulting sound, due to the randomness involved in IGAs. This limitation of the usefulness of IGA based systems for live performance was overcome by applying a metaphor of sound selection in which genotypes are treated as lifeforms or natural processes which the performer may interact with.

Snowstorm and *Terraform* demonstrate approaches to retaining causality between hand gestures and sound material to allow IGA based strategies to be used in live performance. The compositions allow the performer to interact with life forms and natural processes in a journey through surreal landscapes. The compositions provide a structured environment in which some elements are precomposed and others are generated in real time with hand gestures through the IGA interface.

This allows for compositions which are repeatable and recognisable as the same composition in each performance, while allowing scope for the performer to achieve individual expression and/or improvise.

8.3.1 Further Developments

As with the *ESDE* software, the number of SSPTs which can be controlled in MAES+IGA is limited by current CPU speeds. This issue is further compounded by the additional overheads from gesture tracking and IGA mapping. As CPU speeds increase and the range of material which can be generated widens it may be possible to develop new gestures. Additionally, it will be possible to extend gesture tracking and IGA interface to control audiovisual material. The visualisation provides a first step in this direction. Due to the flexibility of the mapping strategies these approaches could be extended to full body tracking and may be usable by people with limited mobility.

The IGA could also be developed further to offer suggestions for material based on the performer's choices, with the objective of moving towards a collaborative process between performer and computer.

8.4 Interactive Visualisation of Genetic Algorithms

The visualisation consists of an animation which represents the state of an IGA and allows the user to change the state by interacting with the animation. The visualisation is presented as a customisable installation which cycles through a range of presets. This system offers a development over previous systems by serving both as an educational tool as well as a musical work. It achieves this by demonstrating the principles an IGA is based on through the animation in a way that the user can quickly understand.

8.4.1 Further Developments

Future developments in visualisation may include developing a 3D version of the animation. This approach could be extended further through the use of video game technology to provide an environment in which the user can explore a landscape and interact with the organisms within it.

This approach may also be integrated with the gesture tracking capabilities of the MAES+IGA software.

Bibliography

Berry, R., Dahlstedt, P. (2003) Artificial Life: Why Should Musicians Bother? *Contemporary Music Review* 22(3), pp. 57 – 67.

Charles, J.-F., (2008). A Tutorial on Spectral Sound Processing Using Max/MSP and Jitter. *Computer Music Journal,* 32(3), pp. 87-102.

Collins, N., (2001). Experiments with a New Customisable Interactive Evolution Framework. *Organised Sound,* 6(2), pp. 121-124.

Dahlstedt, P., (2001). A MutaSynth in Parameter Space: Interactive Composition through Evolution. *Organised Sound,* 6(2), pp. 121-124.

Dahlstedt, P., (2007). Evolution in Creative Sound Design. In: E. M. a. J. Biles, ed. *Evolutionary Computer Music.* London: Springer, pp. 79-99.

De Poli, G., Piccialli, A. and Roads, C. 1991. Representation of Musical Signals. Cambridge, MA: The MIT Press, pp. 137-86.

Eigenfeldt, A. and Pasquier, P. (2012). Populations of Populations: Composing with Multiple Evolutionary Algorithms. In: P. Machado., J. Romero. And A. Carballal (eds.) *Evolutionary and Biologically Inspired Music, Sound, Art and Design*, Heidelberg: Springer, pp. 72–83.

Fischman, R., (1997). The Phase Vocoder: Theory and Practice. *Organised Sound,* 2(2), pp. 127145.

Fischman, R., (2013). A Manual Actions Expressive System (MAES). *Organised Sound,* 18(3), pp. 328-45.

Fischman, R., (2015). *Rajmil Fischman's Free Software.* [Online] Available at: http://www.keele.ac.uk/music/people/rajmilfischman [Accessed 12 May 2015].

Gartland-Jones, A., Copley, P. (2003). The Suitability of Genetic Algorithms for Musical Composition.

Contemporary Music Review 22(3), pp. 43-55.

Johnson, C.B. (1999). Exploring the sound space of synthesis algorithms using interactive genetic algorithms. In A. Patrizio., A.G., Wiggins. And P. Pain (eds.) *Proceeding of the AISB '99 Symposium on Musical Creativity.* Brighton, UK: Society for the Study of Artificial Intelligence and Simulation of Behaviour, pp. 182-196.

Mandelis, J. (2001). Genophone: an evolutionary approach to sound synthesis and performance. In: E. Bilotta., E. R. Miranda., P. Pantano., P. Todd (eds.) *Proceeding ALMMA 2001: Artificial Life Models for Musical Applications Workshop*. Consenza, Italy: Editoriale Bios, pp. 37-50.

Processing. (2017). *Processing*. [ONLINE] Available at: https://processing.org/. [Accessed 16 May 2016].

ScaleMaker. (2011). *ScaleMaker*. [ONLINE] Available at: https://cycling74.com/tools/harmony-maker. [Accessed 16 May 2016].

Schnell, N. et al., (2005). *FTM - Complex Data Structures for Max.* Barcelona, ICMC.

Waschka, R., (2007). Composing with Genetic Algorithms: GenDash. In E.R. Miranda and J.A. Biles (eds*.) Evolutionary Computer Music*, London: Springer. pp. 117-136.

Wrigt, M., Freed, A. (1997) Open Sound Control: A New Protocol for Communicating with Sound Synthesizers: International Computer Music Conference. Thessaloniki, Greece.

Yee-King, M. and Woolf, S. (2003). Virtual and Physical Interfaces for Collaborative Evolution of Sound. *Contemporary Music Review,* 22(3), pp. 31-41.

Ying-Hong Liao., Sun. (2001). An Educational Genetic Algorithm Learning Tool. *IEEE Transactions on Education* 44(2) pp. 20-40.

Appendix – DVD Contents

The thesis is accompanied with 2 data DVDs which contain the compositions and software presented throughout.

The contents, and the corresponding thesis chapters, are as follows:

Data DVD 1:

ESDE Compositions

The Blind Watchmaker – Stereo 2.0 (Chapter 3)

The Singing Forest – Stereo 2.0 (Chapter 3)

MAES+IGA Compositions

Snowstorm (Chapter 5)

Terraform (Chapter 5)

Large-Scale Multi-Channel Composition

The Wind-Up Bird – Stereo 2.0 and Octophonic (Chapter 7)

The Wind-Up Bird Octophonic Configuration Guide.

Visualisation Recording

Sample Visualisation Recording (Chapter 6)

Data DVD 2:

Software + Live Composition Presets

ESDE (Chapter 2)

MAES+IGA (Chapter 4)

ESDE v2.0 Visualisation (Chapter 6)

Processing Sketch (Chapter 6)

ESDE User Guide

MAES+IGA User Gide

Visualisation Configuration Guide.

Composition Presets and Instructions

Snowstorm & Terraform Preset Files (Chapter

5) Visualisation MAX Preset Files (Chapter 6).

Documentation to support the configuration process for the compositions.