



MONASH University

***A Portfolio of Compositions based on Trajectories through Musical
State Spaces***

Michael John Spicer

*BA Hons (LaTrobe University), Dip Com Sci (LaTrobe University), MSc (National
University of Singapore)*

A thesis submitted for the degree of Doctor of Philosophy at
Monash University in 2015
Theatre, Music & Performance HDR Program

© The author 2015. Except as provided in the Copyright Act 1968, this thesis may not be reproduced in any form without the written permission of the author.

I certify that I have made all reasonable efforts to secure copyright permissions for third-party content included in this thesis and have not knowingly added copyright content to my work without the owner's permission.

Abstract

This project attempts to develop a compositional approach where electro-acoustic music can be produced in a way that can easily integrate techniques and processes from different types of musical practices. The method adopted is inspired by an approach developed in Computer Science, to integrate the various fields of Artificial Intelligence. In this portfolio, music composition is viewed as creating the conditions so that the resultant music is an emergent phenomenon of the interactions of a collection of autonomous musical agents. To create a composition, a composer assembles a collection of musical agents and then delegates the task of creating the low-level musical material to each agent, specifying how each agent behaves. Every agent can exist in a number of states and their current state is reflected in their current contribution to the musical texture. By determining the ways an agent can travel through its state space, the composer can shape the composition. This state transition strategy can be implemented in such a way as to emulate the various idiomatic strategies a musician might apply in different musical contexts. Encapsulating compositional processes inside agents provides a way to potentially combine disparate compositional approaches.


To create the works in this portfolio a general compositional framework was devised based on this idea, and three composition models were developed: the Painting Approach, the Pruning Approach, and the Goal-Led Approach. Several compositions were made with each model, varying the details of the different implementations. The portfolio includes interactive and fixed form works for computer-based instruments, analog synthesizers and acoustic instruments. Almost all of the works have some improvisational component.

The approach of encapsulating generative processes in agents, so as to guide them as they explore their state spaces proved to be a successful means of producing music in a way that enables free mixing strategies for generating musical material. The ability to choose each agent's state space traversal strategy gives a composer high-level tools to effectively determine the overall shape of the music. Most of the portfolio works, created with this approach, have been performed either at curated international conferences or festivals.

Declaration

This thesis contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

Signature:

A solid black rectangular box used to redact the signature.

Print Name: ...Michael John Spicer.....

Date: ...8 August 2016.....

Acknowledgements

I would like to thank everyone who has helped me on this project. I have met many people at conferences, on collaborative projects and in more formal educational settings, who have shaped the way the project has evolved. I would like to thank these people, and the many other people from whom I have learnt, either through reading about or listening to their work. I would particularly like to thank my supervisor Thomas Reiner for all his insights and guidance, and my co-supervisor Paul Watt for his input. I would like to thank Katrina Spicer for proofreading the exegesis and Dr Timothy O'Dwyer for recording *Pandan Meditations*. The various members of Sonic Escapade deserve mention, particularly Deborah Tan, Chris Clark, Irfan Rais and Ng Zheng Jie. Particularly helpful have been my colleagues at DMIT, Singapore Polytechnic, especially the DMAT staff, who have willingly accommodated my divided attention whilst working on this project. Lastly, I would like to thank my family, Linda and Eryn, for allowing me the time to do this work, and my Mother for accommodating me during those times when I was in Australia.

A Portfolio of Compositions based on Trajectories through Musical State Spaces

By Michael Spicer	Chapter 1: Statement of the Research Question	8
The Compositional Framework at a Macro Level		13
<i>The Painting Approach</i>		14
<i>The Pruning Approach</i>		15
<i>The Goal-Led Approach</i>		16
The role of Strategies/Heuristics		16
Summary of the Composition Framework		17
A personal perspective of the project		19
Chapter 2: Outline of some Theoretical Concepts		25
Overview of artificial agent concepts from Artificial Intelligence		25
<i>Performance Measure</i>		25
<i>Rationality and Autonomy</i>		26
<i>The Task Environment</i>		27
<i>Basic Artificial Agent Design Approaches</i>		28
Interactive Performance Systems		30
Strategies/heuristics from popular music, with particular reference to the production process in Redgum		32
Application of ideas from Spectromorphology		38
The portfolio works in the relation to other agent-based approaches to composition. ..		42
Chapter 3: The Portfolio Works		49
Aesthetic placement and stylistic predecessors		49
<i>Tape Music</i>		49
<i>Live Signal Processing</i>		50
<i>Fixed Form Computer Music</i>		51
<i>Synthesizer Technique</i>		52
<i>Stochastic Processes</i>		55
<i>Interactive Composing Systems</i>		56
<i>Motivic Transformation</i>		59
<i>Popular Music Structure</i>		60
<i>Summary of influences</i>		61
Compositions that use the Painting Approach		62
<i>Description</i>		62
<i>Compositional approach</i>		64
<i>Personal reflections</i>		78
Compositions that control the ensemble by pruning		84
<i>Description</i>		84
<i>Compositional Approach</i>		86
<i>Personal reflections</i>		91
Compositions that lead the ensemble by setting goals		96
<i>Description</i>		96
<i>Compositional approach</i>		99
<i>Agent Design</i>		101

<i>Performer Agents</i>	101
<i>Controller Agents</i>	104
<i>Restrictions and possibilities</i>	105
<i>Personal reflections</i>	107
Chapter 4: Concluding Summary and Future Work	113
Time, Form and Improvisation	114
Conclusion	118
Section 5: References	123
Bibliography	123
Appendix	127

Composition Portfolio				
	Title	Dur mins	Performance Venue	Type
The Painting Approach				
1	<i>Along the Corridor</i>	5	ICMC 2008	Fixed Media
2	<i>In Transit</i>	5	ACMC 2010	Fixed Media
3	<i>Flute Trio in the Diminished Scale</i>	9	Mosaic Festival 2014	Fully notated score.
4	<i>A Painting in Sound</i>	7	ICMC 2015, SI15	Fixed Media
The Pruning Approach				
5	<i>Sky Castles</i>	2	ICMC 2010 (60x60) Keele Uni 2012 Mosaic Festival 2014	Score + real-time interactive DSP agents
7	<i>Momentary Diversions</i>	12	Modular1 2014	Real-time analog synth
9	<i>Sonic Escapade</i>	6	Various	Live ensemble improvisation
The Goal-Led Approach				
10	<i>Pandan Musings Examples</i>			Score + Real-time interactive DSP agents
	<i>C Lydian Dominant</i>	4	Mosaic Festival 2014	
	<i>D Phrygian</i>	4	Mosaic Festival 2014	
	<i>D Oriental</i>	4	Mosaic Festival 2014	
11	<i>Pandan Meditations Performed by Dr Timothy O'Dwyer</i>	5	Various	Real-time interactive DSP agents

Chapter 1: Statement of the Research Question

The compositions in this portfolio explore an approach to creating (mainly) electroacoustic music with an agent based, bottom up composition approach. The working hypothesis is that musical works exhibiting a comprehensible sense of organisation (coherence) and possessing a distinctive sonic identity, can be created with a compositional approach centred on the creation of a collection of autonomous agents which independently work to create musical material that are combined to make up the piece. In contrast to other applications of the autonomous agent concept to music creation, the focus of this work is on how a multi-agent approach can benefit by embedding a variety of compositional formalisms, or generative processes, within different types of agents so that each agent makes a complimentary contribution to the

resultant musical texture. The embedded formalisms serve as navigational tools for each individual agent, enabling it to find its unique path through the domain of all of the possible states (this set of states is referred to as their state space). The musical output produced by each agent reflects the particular path the agent traverses through its state space.

The use of the term bottom up implies that the process of producing a composition begins at the level of fine detail and works up to the higher conceptual levels. The implication of this is that in this approach, the act of composing is addressed by creating a collection of ways to generate individual musical layers that will be combined to form the final musical texture. Brian Eno eloquently uses the analogy of planting seeds and growing the composition to describe this type of process (Eno, 2004). This implies that some sort of generative strategy (a formal generative process combined with an appropriate mapping function) will be implemented which will determine the activity of each agent as it produces its layer of the musical texture. These layers are combined to form a larger structure. This parallels an approach that some instrumental ensembles have historically used when creating a group arrangement, or whilst engaging in collective improvisation. Individual performers apply their improvisational skills and draw on their knowledge of relevant musical idioms to create an effective, spontaneous ensemble listening experience for the audience. Each player creates their part independently whilst listening and responding to each other in real-time. When successful, the players create a synergistic situation, creating an emergent sonic entity that is greater than the sum of its individual parts.

In this portfolio the composition framework is adapted in different ways for each work. In every composition, each layer of the musical texture is created by an agent implemented either as electronic hardware, software or in the thought processes of a human. The musical knowledge of each agent is manifest in a formal process/generative strategy guiding the unfolding of each part. This modular design provides a framework to plug in arbitrary strategies when creating the various textural layers. Agents may or may not have access to information about the actions of other agents in the ensemble, and agents may or may not have the ability to individually adjust their behavior to collectively achieve some global musical goal.

In the epilogue of his book *The Rest is Noise* Alex Ross states, “The impulse to pit classical music against pop culture no longer makes intellectual or emotional sense” (Ross, 2007, p. 589). He continues and talks of young composers “seeking the middle ground between the life of the mind and the noise of the street” (Ross, 2007, p. 589). The compositions in this folio have an element of attempting to draw on a range of cultures, specifically to integrate various aspects of my musical background which includes working in classical music, rock music, analog electronic music and computer music. In this portfolio there is a conscious attempt to incorporate particular generative processes for creating musical material found in the work of various twentieth century composers with the processes that form the foundation of the work flow I have adopted over a thirty-year period when producing popular music and electronic/computer music.

The composition framework adopted in this project is founded on two concepts from computer science. The first is the concept of the autonomous

artificial agent and the second is a general problem-solving model known as state space search. I was directly inspired to apply these ideas to the area of electro-acoustic music composition when I saw the way Russel & Norvig applied the agent concept as an abstraction to encapsulate and unify the various types of state space search in their book *"Artificial Intelligence: A Modern Approach"* (Russel & Norvig , 2003). The composition framework applied in this project parallels their approach. The agent concept is an analogy that helps me to create the bottom up compositional systems that realise each piece. The agent concept and the state space search model provide a degree of abstraction and modularity that allow a composer to plug in a diverse set of generative techniques and processes drawn from many different knowledge domains, and enable these to work together.

The compositional framework assumes the following:

Assumption 1

A bottom up compositional approach where high-level musical structures are created by the selection and combination of musical material that was generated by an arbitrary collection of autonomous composition agents, is capable of producing music that exhibits a comprehensible sense of organisation and a distinct sonic identity.

Assumption 2

Many compositional strategies can be implemented as formalisms, meaning that they can be described in formal mathematical or logical terms. Some combinations of strategies facilitate the creation of complementary idiomatic musical material. Compositional strategies implemented as formalisms

can be encapsulated within autonomous agents to direct the generation of musical material.

Assumption 3

The first two assumptions allow a composer to combine disparate musical generative processes so that a composer can freely draw upon incongruous sets of musical practices in order to create new works.

This portfolio attempts to establish the validity of these assumptions by exploring the application of these assumptions in several musical settings. To do this, each work in the portfolio is created with a multi-agent based, bottom-up composition framework that is based on these assumptions. The framework was implemented in three different ways (models), and several works were created with each model. The details of the different implementations of the same model are quite varied in each instance. In some works, the agents are implemented in software, in others as configurations of electronic hardware, or they could be affected by people consciously acting out an agent design either as a composer of a fixed form piece, or as an improviser working in real-time. The focus of this portfolio means that the works are primarily electro-acoustic, and draw on a variety of timbres, formalised compositional processes and performance paradigms from my professional work creating electronic and popular music. This portfolio is not intended to be an exhaustive investigation, but attempts to establish that multi-agent based compositional systems provide a practical framework to create music that can coherently draw on disparate approaches to generate musical material, and create a musical work exhibiting a distinct identity.

The Compositional Framework at a Macro Level

In this composition portfolio, the process of creating a musical work is analogous to running a collection of state space search problems in parallel, where each search is embedded inside an agent. A state space, in this context, is the set of all the possible states that the agent could be in. When dealing with multi-agent systems, the state space of the system can be considered at the low level, focusing on the internal state of each individual agent, or at a higher level, looking at the collective state of the whole ensemble. Low-level agents individually produce musical material that reflects their exploration of their state space. The resulting collection of material from all of the low-level agents then becomes the state space of another high-level agent, whose task is to present this material to the audience in a coherent way.

While the form of an agent can vary considerably, all agents share some basic characteristics. Each can perceive their operating environment in some way. In a software or hardware agent, this corresponds to whatever control inputs are present. Agents also have some ability to change or interact in that environment, which could include their audio output, the ability to generate control signals and contribute data for analysis. Embedded inside every agent is a mechanism that determines how it will interact with the environment through these inputs and outputs. The choice of the strategy implemented in this mechanism to guide each agent as it traverses its state space has a fundamental bearing on the agent's contribution to the musical texture. The details of an agent's operational environment varies from work to work, but it consists of the state of all the agents, plus any control devices used.

The act of composing with this framework involves the assembly of a collection of agent performers that each control some sort of sound generator or processor. The musical output of an agent is dependent upon its internal state, so a systematic mapping the agent's internal state to particular musical parameters needs to be implemented. Once the state space search strategies for each agent are determined, the piece progresses as each agent traverses its state space from a specified initial state, and the musical output of each agent reflects the changes of the state along the path it takes. Agents may be designed to either move through a successive sequence of goal states, or to simply explore the state space. A large part of the effort of creating music with this framework is to ensure that the agents in the ensemble contain complimentary search strategies, collectively producing a musical output that exhibits some synergistic emergent behaviour. Three basic composition models are derived from this multi-agent composition framework to realise the works in this portfolio. Some works are created in real-time, some have partial in real-time elements, and some are completely pre-composed.

The Painting Approach

The first composition model is analogous to a simple approach to creating pictures. In this approach, pictures are considered to be made up of a number of discrete shapes that are arranged to create the overall scene. Similarly, a composition can be considered to be made up of a number of discrete musical objects, such as patterns of sounds or notes, which are arranged into larger scale structures. This implies that the composition process can be divided into two

district phases: the generation of the basic musical material and the assembly of that material into a larger structure.

A similar analogy is that of building a plastic model kit, such as a model sailing ship or a plane. The first phase involves creating the basic components, which is done by the model manufacturer, and the second phase is assembly, done by the hobbyist. The significant difference between the plastic model building process and the works made using this approach, is that the final form of the plastic model is known from the very beginning while, when composing, I prefer to start with very few, if any preconceptions and allow the final form to emerge during the assembly process. I view this sort of composition process as an opportunity for musical exploration that produces a result that is not quite predictable, reflecting my view that if everything turned out exactly as I had planned, I must have missed a lot of opportunities that occurred along the way.

The Pruning Approach

Another way to implement a multi-agent composition framework is to create an ensemble of agent performers, allow them to explore their state space and develop a high-level agent whose role is to determine which performer agent's musical output is heard by the audience. The high-level agent prunes the tree of possible sounds that are heard. Dynamically altering which agents are heard provides the mechanism to shape the entire piece.

The Pruning Approach is commonly used in modern popular music, especially in the dance related genres. The difference between the works in this portfolio and common popular music practice is that in the popular context, the musical material usually consists of loops rather than the more variable musical

material generated by the multi-agent systems in this project. This increase of complexity adds an element of uncertainty that can provoke a more spontaneous reaction from a human performer.

The Goal-Led Approach

In the third implementation model of the multi-agent composition framework, the composer works at the highest level and shapes the collective state of the system during performance. This model is most suitable for real-time application, where the composer/performer influences the musical output by specifying high-level musical attributes that act as goal states for the entire ensemble. The low-level performer agents autonomously adjust their internal states to achieve these goals. The only low-level tasks for the composer are to specify the agents' initial states, specify the mapping of the internal state to the musical output and to specify the mechanism used to traverse the state space.

The role of Strategies/Heuristics

Exactly how a state space search proceeds is determined by the mechanism chosen for changing from one internal state to the next. There are a number of common approaches found in the literature, including: depth first search, breadth first search and iterative deepening search (Russel & Norvig , 2003). The ability to plug in different strategies, changing the behaviour of the agent, is fundamental to the functionality of the multi-agent composition framework. The modular structure of an agent enables processes taken from disparate musical contexts to be combined in one environment. Each agent performer can choose its next state with its own individual strategy, which could be designed to realise some idiomatic compositional process drawn from any

arbitrary musical style. This implements a type of heuristic state space search where a heuristic can be thought of as a rule of thumb. This type of search strategy incorporates domain specific knowledge that either helps it to find its goal state more efficiently, or to exhibit particular behaviour. Heuristic search is a type of informed search. Many of the idiomatic techniques used by composers, arrangers and improvisers to create musical material can be considered to be rules of thumb. State transition mechanisms can be derived to implement established idiomatic (genre specific) practices from mathematical functions, from particular configurations of voltage controlled synthesiser modules, or from other formalisms such as stochastic procedures and chaotic generators. The strategies will determine the specific musical attributes of each voice in the musical texture by determining the agent's path through its state space. The choice of strategy used by each agent in the ensemble is one of the primary compositional decisions that has to be made when composing with this framework, and care must be taken to make sure that the strategies chosen complement each other. Strategies that have proven themselves to work together in other contexts are likely candidates that will be effective in the multi-agent situation. The works in this portfolio implement state transition strategies that I have found to be effective in popular music and in experimental electronic music.

Summary of the Composition Framework

The overarching idea underlying all the works in this composition portfolio is to construct a collection of sound creating agents, and these will collectively construct the materials that will constitute the piece. Each agent

may exist in a number of distinct states, and each agent's musical output will reflect its current state. During a performance, agents explore their state space guided by strategies that are implemented as formal procedures. Most of the agents produce the raw musical material, but a small number may control high-level attributes of the ensemble output. Different types of agents embed different music generation strategies and produce musical output in different ways. The agents, in some cases, may be affected by using humans either as performers interpreting a score or as improvisers improvising within a particular set of constraints, or could be acted out by the composer adopting specific roles. In most of the works, agents are manifest either as software, which implements particular algorithms, or as electronic hardware configured to behave in a particular way. The essential common factors, in all cases, are that the agent performers have a substantial degree of autonomy when deciding the details of each successive note/sound that they produce, and that they faithfully implement the required strategies to generate their contribution.

Elements that differentiate the works in the portfolio include:

- The state transition mechanism.
- The sound production mechanism available to the agents.
- The stylistic idioms that are imposed on each agent.
- The mechanism and degree of agent interactions.
- The method for specifying goal states, if any.
- The mechanism and degree of high-level control over agent behaviour.

A personal perspective of the project

I am interested in music as a process-centric activity. When creating a piece of music my focus is on assembling a collection of procedures that will eventually enable me to communicate some sort of sonic narrative to listeners in some imperfect way. These procedures can range from mathematically precise formalisms implemented in software to informal heuristics implemented as rules-of-thumb that can be realised in real-time by instrumentalists. The three broad autonomous agent-based compositional frameworks in this project try to achieve this end in different ways, and each framework lends itself to a finished work having particular musical attributes. The Painting Approach is suitable for producing fixed works with a static compositional structure. The sonic elements are the basic building blocks, so they need to be created first. In contrast to this, the Goal-Led approach constructs a dynamic environment where a unique instance of the piece will unfold during each performance. This type of piece is initially framed in abstract form by articulating the specific compositional and sound production strategies available to each agent. This type of work instantaneously takes on the form, at the note-level and at the macro-levels during a performance. The Pruning Approach represents a middle ground between these two approaches. Despite these differences, my aim is always to create an acoustic environment that is comprehensible and draws the listener through specific sonic territories.

While I have primarily framed my compositional approach in terms of the way that an agent paradigm allows the integration of a variety of techniques from my various musical activities, there is another aspect of the project that makes this approach very interesting to me. The agent-based state space

traversal model can be likened to a musical application of the exploratory approach that I have adopted in my daily life. In some way, this project aligns my pattern of behaviour with an approach to composing music. I typically operate by consciously establishing the direction I want to go in the various domains I am involved with by formulating long and medium term goals. I am deliberately unspecific when it comes to short-term planning. My short-term plans are usually frameworks that outline intent, a point of departure and an exploratory strategy that is constrained by the specific requirements of the current circumstances. However, this open strategy does need to be managed. I do make extensive use lists and notes to ensure that everything that needs to be done is accomplished. These serve as a repository of ideas and a record of my activities. It is ironic that I find that I need to be extremely organised in order to have the opportunity to be spontaneous. With the exception of regular tasks, where I work out a routine to accomplish them efficiently, I do not usually engage in planning at a fine level of detail because I want to be able to capitalise on whatever unforeseen opportunities may come along. I proceed to explore the possibilities that are presented to me, and consciously look out for what I think are interesting options. Decisions, with potentially far-reaching consequences, frequently need to be made with incomplete information and, because this involves an element of risk, a deliberate conscious effort is required. In order to optimise my chances of success within the particular context, I usually adopt heuristic strategies that I think will enable me to recognise and avoid non-viable paths and that help to manage the complexity of the current problem. These heuristic approaches usually involve applying a simplified problem-solving framework to the problem. This framework involves working out how to

generate a set of possible actions, analysing these and comparing the options against some relevant measurement criteria. This approach to daily life is analogous to the agent-based approach to music composition in this project. In recent years, since reading (Kahneman, 2011) I have been made more conscious of the potential pitfalls of reflex thinking, what Kahneman calls System 1 thinking, and I am careful to avoid oversimplifying problems and becoming a victim of my personal biases. I sometimes explain this overall approach using a surfing analogy: Each wave is an individual and has its own evolution. A surfer has to ride the waves that are there on a particular day, and make the most of their potential, regardless of whether they are perfect or not. I have been actively involved with sailing and windsurfing most of my life, and this attitude of making the best out of the prevailing conditions has become infused into my worldview, and consequently my creative activity. This adaptive and opportunistic approach emphasises spontaneity and minimises constraints, which ties in with the improvisational elements of the agent-based compositional approach. My general approach to operating within in the world has meant that I have adopted a practice of acquiring a relatively high level of technical skills and background knowledge in the various areas I am interested in. This practice acts as a safety net and enables me to adequately execute most technical demands that might be required as they occur without extra preparation. It helps to balance my innate aversion to risk with the pleasure I get when exploring the unknown. So far, this approach has led me through a rich catalogue of experiences in my both my personal and professional life, with new possibilities always opening up before me.

In addition to helping me to operate effectively as an individual, this exploratory worldview has informed my approach to interacting with groups of people, especially in a work environment. Over the years, I have led many creative teams and have always managed to produce viable, and usually successful, products in fields including software development, music, and formal education courses. The lessons I learnt from my experience in professional bands, and from reading about how Miles Davis (Davis, 1990) and Frank Zappa (Zappa & Occhiogrosso, 1989) ran their groups, have had a significant impact on my thoughts on how to effectively lead and manage creative people. It is a cliché to say that the role of a team leader is to provide an environment where creative individuals have the mandate to make the best use their unique talents to make a contribution to the final product, but I think that it takes a conscious effort to give individuals autonomy and not micromanage them. While this can't be considered a ground-breaking new idea, this view of group dynamics is a foundation of my approach to creative work. To create a musical work, I assemble a team, an ensemble of autonomous agents. I carefully consider what each agent brings to the work in terms of specific skills or strategies that they can contribute to making their particular part of the final work. The role of a leader, or composer, in this type of situation, is to manage the interactions between the participants, to provide direction and feedback, and to manage the operating environment. Because this state of affairs involves delegation, the choice of which individual is assigned to each role is crucial. In the real world, a leader has to trust in the skills of the person delegated to fulfil each specific role and has to be prepared emotionally to accept that the delegated work may not be exactly as they would have done it themselves. The leader/composer needs to

turn this potential liability into a strength. The agent-based composition methodologies in this portfolio represent some attempts to apply this concept to creating electronic music. As such, the composition methodology is built around distributed creative decision-making and puts an emphasis on spontaneity, exploration, strategy, and analysis at multiple levels. The composer sets up the environment, designs the participants and the relationships between them, orchestrates their interactions, and maps all of these to create organised sound that is in some way meaningful to some listeners.

One aspect of this compositional approach that I find intriguing is that the final sonic result is an emergent phenomenon. I am often surprised by the way that hierarchical combinations of simple things (autonomous low-level entities) can combine to create high-level elements that exhibit complex properties through their interactions. I see this manifest in the physical and biological world all around me, and in recent years have come to realise that it is a fundamental aspect of living things, and specifically has informed my view of what it is to be a human being. My thinking on this has been influenced by the computational theory of the mind, as explained by (Pinker, 1997). Pinker's book drew my attention to the connection between our physical being and our ability to process information to become what we perceive to be ourselves. From this perspective, we consider beliefs and desires to be information that can be encoded as symbols, which can be represented by the physical states of neurons in our brains. Neurons symbolise particular things because they can be triggered by the presence of those things via our sensory organs. Once they are triggered, neurons that symbolise one belief can trigger other neurons, enabling information to be processed in various ways, giving rise to related beliefs etc.

This physical connection allows “meaning to cause and be caused” (Pinker, 1997 p. 25). What struck me about these ideas was the articulation of a vision of many subsystems of symbolic processors (neurons) working in parallel to somehow (we don’t know how) give rise to the emergent phenomena that is us. People exist physically as a combination of many autonomous sub-systems that process information in the brain at a level below consciousness.

In the outside world, emergent phenomena are manifest everywhere, in the many layers of the ecosystem, in the artificial world of machines and in the information space, where simple programming structures are combined to produce complex software. These ideas have directly influenced my creative approach, and it is my aim to apply this modular construction methodology to the creation of music that will create an engaging experience for the audience.

This project combines my enthusiasm for the idea of emergence with my penchant for building things and for my preoccupation with creating music in a way that is consistent with how I operate in the real world. The notion of emergence unifies some of the musical ideas, mathematical ideas, computational ideas and methods that I have an affinity with to create electronic music compositions.

Chapter 2: Outline of some Theoretical Concepts

Overview of artificial agent concepts from Artificial Intelligence

The application of the agent concept in the Artificial Intelligence community was an inspiration for the compositional approach used for this portfolio. In the 1990s, the agent approach was adopted because it enabled the many different subfields of Artificial Intelligence to be unified. In this approach, Artificial Intelligence was defined as “the study of agents that receive percepts from the environment, and perform actions” (Russel & Norvig , 2003, p. vii). This parallels this project, where the agent approach becomes a framework to facilitate the combination of a number of compositional processes.

Performance Measure

In Artificial Intelligence, the aim is to create a rational agent, which is one that does “the right thing”. To do this, some sort of performance measure is obtained to quantify the degree of success of the agent. Different performance measures are needed for different types of agents, but for software agents, the performance measure needs to be objective and it will directly affect the agent’s behaviour. This means in colloquial terms ‘What you ask for is what you get’. A general rule in the design of performance measures is to define the performance measure in terms of what is actually required, rather than how the agent behaves. This rule is also applicable when using human agents, but in musical situations, there can be more uncertainty of what “the right thing” is, and human agents will have more subjectivity in their decision-making. They are going to

choose what sounds right to them, which will largely be influenced by their background and personal taste. Miles Davis has said that when putting together a band, choosing the appropriate people is extremely important when choosing musicians for recording sessions (Davis, 1990), which corresponds to choosing the right combination of agent designs.

Rationality and Autonomy

The rationality of an agent, which is its ability to do the right thing, is determined not just by the performance measure but also by the prior knowledge of the environment that is embedded in the agent, the actions that the agent has the ability to perform and the percept sequence presented to the agent.

Russell & Norvig define a rational agent as:

For each possible percept sequence, a rational agent should select an action that is expected to maximise its performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has. (Russell & Norvig, 2003, p. 36)

A rational agent optimises the expected performance of the agent and relies on the percept sequence to make its best attempt at maximising the performance measure. In addition to the percept sequence, it also has prior knowledge programmed into the agent by its designers. The degree of reliance on this prior knowledge determines its degree of autonomy. The less an agent relies on pre-programmed knowledge, the more autonomous it is. This implies that an autonomous agent has the ability to learn from its experience, beginning only with the pre-programmed knowledge and building on this with the experience gained from its percept sequence. In this project, the degree of autonomy desired

in the agents varies from piece to piece. In some situations, the degree of autonomy is limited, as the aim is to delegate only some of the compositional choices to the agent. This project aims to produce musical output with the characteristics specified by the composer and these are determined via the computer supplying appropriate pre-programmed knowledge, or by limiting the actions an agent can take.

The Task Environment

The design of an artificial rational agent is directly affected by the Task Environment, which can be described as the 'problems' to which rational agents are the 'solutions' (Russel & Norvig , 2003, p. 38). Task environments can be specified using what is referred to as the PEAS description. It consists of:

- Performance Measure
- Environment
- Actuators
- Sensors

There are a very large variety of task environments possible, but they can be categorised along six dimensions. To a large extent, these dimensions determine the applicability of the different techniques available in implementing the agent design. The Fully Observable vs. Partially Observable dimension is a measure of how complete the agents' knowledge is of the relevant aspects of the task environment required to make good decisions. The Deterministic vs. Stochastic dimension measures how much uncertainty is caused by factors outside the agents' control. The Episodic vs. Sequential dimension is concerned with whether the decisions made by an agent will only have an impact for a finite time

duration, or will have a long-term effect on future behaviour. Agents also may inhabit either a static or a dynamic environment, which determines whether the state of the environment can change while the agent is processing its current percept and deciding upon its action. There are some agents that work in a Discrete vs. Continuous manner in terms of the way time is handled, the representation of the state of the environment and the percepts and actions. The final aspect is whether the system is a Single Agent or Multi-Agent system. Agents that coexist in the same environment may compete with each other or cooperate and work together to achieve a collective goal.

Basic Artificial Agent Design Approaches

The agent program is the part of a software agent that determines the agent's behaviour. There are several approaches commonly used in the design of the agent program, depending on what the agent is designed to achieve. The simplest useful agent design is the reflex agent, which consists of a collection of condition-action rules in the form of: 'if this occurs then do that'. This approach embeds knowledge into the agent and only makes use of the current percept to make a decision as to how it acts. There are some instances of reflex agents in this portfolio, such as those implemented in the automation system of a Digital Audio Workstation. An example of when this type of agent occurs is in *Sky Castles*, where reflex agents whose behaviour is completely determined by the current playback time, articulate the high-level form.

A slightly more complex approach is an agent design that maintains an internal state, which has a role in determining an agent's action. This is a type of Finite State Machine. Historically, Finite State Machines have often been applied

in computer games to create non-player characters (LaMothe, 1999), because it is a computationally efficient way of creating relatively complex behaviour. *Pandan Musings* relies on state-based agents to set high-level goals for other agents.

Another agent design, the goal-based agent, requires a more sophisticated approach. This design needs some way of representing a goal and of periodically evaluating the various actions that it can take to see which action will most likely contribute to achieving that goal. Some compositions in this portfolio, such as *Pandan Musings*, are reliant on goal-based agent designs to implement the strategies required to create the details of the musical texture. The behaviour required of low-level agents is articulated by specifying goals. The low-level agents autonomously modify their internal states to collectively achieve their goal and exhibit the required behaviour. The mechanism to achieve this is a very basic machine learning algorithm implemented inside the agents that enable them to progressively modify their internal state, to bring them closer to the goal state.

More generally, due to the complexity of the agent programs required to implement agent functions, many programmers now approach building agents in a way such that the agents can learn. Some initial knowledge is pre-programmed in, and then the agents are trained to produce the required agent function. Learning agents have four main components:

- Learning element – responsible for making improvements
- Performance element – responsible for selecting actions (this is the ‘agent’)

- Critic – determines how well the agent is doing with respect to an external performance standard.
- Problem Generator- responsible for suggesting actions that will lead to new and informative experiences, so it might discover better actions in the long run.

The details of these components are determined by the representation used. The overall idea is to modify the various components in an agent to bring it in closer agreement to the available feedback information, to achieve an improvement in overall performance.

The artificial agents used in this project are deliberately kept relatively simple. The overarching goal was to produce usable musical results in real-time, on readily available hardware. Reflex agents, state-based agents and goal based agents form the basis of the most of the works, but in some of these systems, a simple learning algorithm is implemented within a goal-based agent.

Interactive Performance Systems

Several compositions in this portfolio build on the interactive composing approach developed by Joel Chadabe. (Chadabe, 1984) Exposure to this work as an undergraduate sparked my interest in working with algorithmic composition systems, particularly the aspect of balancing an improvisatory real-time element within a prebuilt environment.

Chadabe defines Interactive Composing as a two-stage process consisting of:

- Creating an interactive composing system
- Simultaneously composing and performing by interacting with that system, as it functions.

An interactive composition system consists of a computer program that links a synthesizer to some sort of performance input device. The computer program contains a collection of real-time algorithms that interpret the human performer's actions as partial controls of the music. It generates controls for those aspects of the music that are not controlled by the human performer and then directs the synthesizer to generate the sounds.

An interactive composing system has the characteristic that it responds to the performer in a "complex, not entirely predictable way." (Chadabe, 1984) The system and the performer share control of the music produced and the performer needs to react to the system's unpredictable elements in his/her performance. "The computer responds to the performer, and the performer reacts to the computer, and the music takes its form through that mutually influential, interactive relationship" (Chadabe, 1984). Chadabe illustrated the structure of the system using the diagram shown in figure 1.




Figure 1 Block diagram of a system for Interactive Composing (adapted from Chadabe)

This diagram breaks the process into distinct blocks. The composition algorithm receives two sets of inputs. One set comes from the human performer, via the performance interpreter algorithm, which transforms the input from whatever controllers are used into a form that can be used by the composition algorithm. The other set comes from the intelligence algorithm, which defines the aspects of the music that are not directly controlled by the human. These are the most important elements affecting the behaviour of the system. The precise allocation of which parameters are controlled by the human or the computer is determined by the particular composition being implemented. The output from the composition is then translated into the appropriate synthesizer control codes by the sound algorithm.

Some of the works in this portfolio apply a multi-agent paradigm to implement Interactive Composing systems. In these cases, the human performer guides and reacts to the musical output of an entire ensemble of agents. This has the potential to generate a complex musical output and there is a high chance of agents getting in each other's way. The approach of applying complimentary music generation strategies in the different agents was adopted to alleviate this risk.

Strategies/heuristics from popular music, with particular reference to the production process in Redgum

One element of this project is the integration of some work processes common in popular music into an electro-acoustic musical context. Popular music arrangements are created in many ways, but several strategies for creating textural layers tend to recur. My methodology is derived from the work

in the commercial music industry in the 1980s, most notably with the successful Australian folk-rock band Redgum. A number of simple strategies to quickly and effectively develop musical material were used, which enabled us to collectively create intricate arrangements of quite simple songs, particularly on the albums *Frontline* (Redgum, 1984) and *Midnight Sun* (Redgum, 1986). Usually, an individual band member would present a new song to the band to develop. Initially, particularly in this band, there was well developed lyrical content, but only a skeleton of the musical content. The musical details evolved once the band started to work on it. The methods applied by the band members to develop their parts had not been formally taught to them but were informally learnt through interacting with or imitating other musicians. These techniques are applicable to a large variety of popular music styles ranging from folk music to rock and jazz. (This approach to developing material is now described in books and websites, often in the context of helping novice jazz players work in a collective improvisation setting. A good example of this is the online jazz primer by Sabatella (Sabatella, 2000)) In Redgum, each player initially created their own part, making sure to not lose the essence of the song. The process of working out arrangements started in rehearsal and was further refined in live performance or in a recording studio. Generally, individual performers varied their strategies through the course of the song to shape the overall dynamics. Rarely would two players use the same strategy at the same time. Strategies for creating instrumental parts when developing arrangements in Redgum included:

- Pads
 - The role of a pad is to provide harmonic “glue” that helps to outline the harmony. Pads are usually mixed at a fairly low volume. Often

the average listener may not be consciously aware of it, but would notice that it was missing if it were to be turned off. Pads usually consist of between one to three layers of voices that play long (mainly) consonant notes, with minimal movement, often in the middle register. Often each voice in the pad has the same timbre and contains mainly harmonic partials. Pads on the two albums mentioned above were created with a variety of instruments, such as analog synthesizers, samples of recorders, samples of voices and FM synthesizers. Sometimes pads would make use of slow timbral evolution to provide a subtle sense of movement. In more recent times, it has become common to use a low-frequency oscillator synchronised to the song tempo, to modulate some aspect of the sound to add a subtle rhythmic component. A common pad technique, often not thought of as such, involved the use of sustained distorted guitar power chords. The concept of the pad can be generalised to be any unobtrusive sustained background part which has a small amount of movement. In electro-acoustic music, there may be no harmony, but pads can still act as a kind glue.

- Interacting rhythm layers
 - One technique that was well suited for creating arrangements in a multi-track recording/sequencing environment was the creation of several layers of different instruments/timbres that work together to create one complex part. These layers tend to be characterised by short, repeating patterns. There may be many

layers contributing to this part, but usually, there is only one or two notes sounding simultaneously. Each layer tends to be made up of short notes, distributed on different rhythmic subdivisions, with different note rates. This type of part can occupy a wide range of pitches and may include harmonic and inharmonic timbres.

- Arpeggio based parts
 - Arpeggio parts are usually repeated patterns of notes that articulate the notes of the current chord. They may contain the occasional passing note or neighbour tone, but they are characterised by disjunct motion. The pitch contour typically does not simply ascend and descend but tends to exhibit a distinctive shape. Arpeggio parts tend to have a moderate range centred in the middle register. In Redgum, they were often performed on guitar, as they are an idiomatic element of various folk guitar styles, but they can be played on keyboards. Usually, only one arpeggio part will occur at a time, but they may be used in parallel as components of interacting rhythm layers.
- Bass
 - Bass parts consist of low notes often played on bass guitar or synthesizer that provide a harmonic and rhythmic foundation for the rest of the musical elements to build upon. A prominent feature is the use of simple repeated rhythm patterns, largely played on the tonic note of the current chord.

- Comping (short for accompanying)
 - An instrument plays chords with a simple, cyclic rhythmic pattern. Occasionally arpeggio elements are added. This strategy provides a firm harmonic and rhythmic foundation. Common examples found in the Redgum music feature cyclic rhythm patterns articulated by strumming a variety of stringed instruments or playing block chords on acoustic, electric and electronic keyboard instruments. Typically, a comping part will consist of at least four voices, possibly more, each having the same timbre. Quite frequently several instruments would play different complimentary comping parts simultaneously.
- Melody
 - Most of the music made by Redgum predominantly exhibits a homophonic texture so there is a focus on the melodic line. Redgum's vocal and instrumental melodies tend to have quite different characteristics.
 - The vocal melodies tend to have simple rhythms, a very narrow range, with mainly conjunct motion and many recurring pitches. They tend to be very repetitive, easy to sing, and are designed primarily to deliver the lyrical content.
 - Redgum's instrumental melodies have greater variation; they have a greater range and are more complex rhythmically.
 - There is a distinct influence of Irish and English folk music perceivable in the instrumental lines, as well as some blues elements.

- Counter melody
 - Counter melodies occur quite often in Redgum's music. They are melodic lines, often played on tin whistle, violin or guitar, that are designed to complement the main melody. They often reduce their activity, or rest, when the main melody phrase occurs, then become more active in the gaps between phrases.
- Parallel harmony
 - A common approach, especially when arranging vocals, is to have a few lines that parallel another part, usually the lead vocal. These parts would typically have the same rhythm as the other part, but are played with a relatively constant pitch offset.
- Drone
 - A drone consists of (usually) one long note, held constant as the other elements such as melody and harmony progress. In Redgum, I played a synthesizer or a length of PVC pipe played as a didgeridoo to perform this role. Usually, drones are low pitched, but are occasionally in middle or upper registers. A drone, by definition, has no pitch movement, but it can have complex timbral evolution. This opens up possibilities for creating cyclic patterns, such as those typical of the didgeridoo. Synthesizers are particularly well suited to this approach due to the complex modulation effects that can be created with low-frequency oscillators, envelope generators or step-time sequencers.

- Textural Strategies
 - The number of layers in the musical texture of a typical Redgum song changes throughout the piece. A song may start with few layers and gradually build up to many layers, adding one layer at a time. To a large degree, the emotional dynamic of the song was articulated through the textural layering: adding more layers increased intensity, abruptly dropping many layers created an element of surprise etc.

It is important to note that although the individual players developed their parts using the above approaches and when recording, there was a review process that involved the band members, the producer and the recording engineer. The recorded musical part was agreed upon through a process of negotiation. There was no arranger in the traditional sense. The adoption of this strategy-based process facilitated the distribution of notes so the different layers did not get in each other's way. It helped create a sense of space and transparency in the final arrangement so that frequencies from one instrument didn't mask those of another. The emergent complexity that arises out of the interactions of a number of simple musical elements has had a lasting impact on my approach to creating musical material and is a feature the composition framework adopted for the works in this portfolio.

Application of ideas from Spectromorphology

Denis Smalley has extended the ideas of Pierre Schaeffer regarding *Musique Concrète* (Schaeffer, 2012) to develop a descriptive tool, based on aural perception, referred to as Spectromorphology (Smalley, 1997). This has had a

significant impact on my approach to electro-acoustic music composition, as it provides a concrete, easily understandable conceptual framework in which to work. Some of the ideas within Spectromorphology parallel the popular music textural strategies discussed outlined above. Although Spectromorphology was conceived as an approach to help a listener grasp electro-acoustic music, it articulates ways of shaping a sound that can also be helpful when composing. Blackburn (Blackburn, 2011) outlined how these concepts could be applied to composing fixed form electro-acoustic music. Spectromorphology contains a catalogue of strategies that are expressible as formalisms and these can be embedded inside an agent. A summary of some of these is shown in Table 1 below. The strategies for creating material for popular music, outlined above, can be considered a subset of these approaches and are illustrated in Table 2.

temporal phases	Motion Processes	growth processes	Characteristic motions	Textural motion qualifiers	Behaviour	Spectra	Spectral space Frame	Spectral space qualifiers	Spectral density
Attack alone	Unidirectional • Ascent • Plane • Descent	Bi/ Multi-directional	Fly	Streaming	Motion coordination	Note • Harmonicity • Inharmonicity	Canopy	Emptiness—plenitude • Covered • Large gaps	Distant — close continuum
Attack-decay	Reciprocal • Parabola • Oscillation • Undulation	Agglomeration Dissipation	Float	Flocking	Motion passage	Noise • Granular • Saturate	Center	Diffuseness — concentration • Spread • Concentrated	Filled
Graduated continuant	Cyclic/centric • Rotation • Spiral • Spin • Vortex • Pericentric • Centrifugal motion	Dilation Contraction	Drift	Convolution			Root	Streams — interstices • Narrow streams • Broad streams	Packed / compressed
		Divergence Convergence	Throw/ fling	Turbulence				Overlap — crossover	Opaque
		Exogeny Endogeny	Rise	Continuous Discontinuous • Sustained/ • iterative • Granular					Translucent
			Flow	Periodic					Transparent
			Push/drag	Aperiodic					Empty

Table 1 Useful concepts drawn from Spectromorphology.

	temporal phases	Motion Processes	characteristic motions	Textural motion qualifiers	Spectra	Spectral space Frame	Spectral space qualifiers	Spectral density
Pad	Graduated continuant	Plane	Flow	Continuous	Harmonicity	Center	Covered Narrow Streams	Distant Opaque
Interacting rhythm layers	Attack-decay	Cyclic	Flow	Iterative cyclic	Harmonicity	Center	Gaps Narrow Streams	Mid-distance
Arpeggio based parts	Attack-decay	Cyclic	Flow	Iterative cyclic	Harmonicity	Center	Covered Narrow Streams	Mid-distance
Bass	Graduated continuant	Cyclic	Flow	Cyclic	Harmonicity	Root	Covered Narrow Streams	Mid-distance
Comping	Attack-decay	Cyclic	Flow	Cyclic	Harmonicity		Gaps Broad Streams	Mid-distance
Melody	Graduated continuant	Undulation	Flow	Continuous	Harmonicity	Canopy	Covered Narrow Streams	Close
Counter melody	Graduated continuant	Undulation	Flow	Continuous	Harmonicity	Canopy	Gaps Narrow Streams	Close
Parallel harmony	Graduated continuant	Undulation	Flow	Continuous	Harmonicity	Canopy	Covered Narrow Streams	Close
Drone	Graduated continuant	Plane	Flow	Continuous	Harmonicity	Root Canopy	Covered Narrow Streams	Mid-distance

Table 2 Spectromorphology Strategy mapped to pop music functions

The portfolio works in the relation to other agent-based approaches to composition.

The application of software agents to solve complex problems is now quite common in many areas. Whalley, in his overview of the application of software agents to music and sound art (Whalley, 2009), notes that the application of software agent technology to create music and sound art is part of a wider interest in the development of interactive music systems, along with some other interactive music approaches, such as evolutionary systems and A-Life. Whalley categorises agent based interactive composition systems into four basic types. Two of these types focus on structure, while the other two focus on process.

Structural Focus Approaches:

One category of multi-agent systems, which focuses on structure, is Simulation systems. These can naturally be applied to the creation of linear tonal music, or of set forms realised by the manipulation and transformation of motives. The agent approach works well in real-time data and multi-causal situations where the various elements constantly need to react to each other as they construct a structural whole.

Whalley's other structural category is the Performance/Reactive Systems approach, which extends the Simulation approach by creating open systems that can accommodate real-time human input, enabling increased control over aspects of the musical output. Whalley calls this a type of performance Command and Control. In this type of system, agents tend to have low levels of autonomy and learning ability and are generally based on traditional notation.

Process Focus Approaches:

The first process based approach to multi-agent systems is Generative Systems (Dorin, 2001) and (Miranda, 2001). Generative Systems can be applied to both note-based and sound-based composition, and agents create content and structure based on the dynamic interplay of parts. Instead of creating predetermined forms, the composer specifies a range of agent behaviours which are the result of particular logical processes. A variety of particular instantiations of the piece can be produced, rather than creating a fixed form.

Whalley states that the most extensive deployment of agent technology in a musical context is the Generative/Improvisational System, which allows real-time improvised human input and human and machine adaptation (autonomy and learning). Little of the final output is prescribed but is the result of a dialogue between human and machine.

With respect to Whalley's four categories, most of the real-time works in this portfolio would be classified as either the Performance/Reactive approach or the Generative/Interactive approach. The basic premise of these works is that generative procedures can be a mechanism to navigate a particular state space and that these can be implemented so as to provide real-time handles that can control their trajectory, providing a means of interaction. Because these real-time works are inherently generative, they are conceived from a process perspective. In fact, the fixed form pieces could be considered to be instances of the output of various generative processes that have been pre-chosen by the composer to be presented to the audience. In this view, the primary compositional action for all of the works in the portfolio is the choice of the processes that will be the basic building blocks for the generative subsystems.

There is an aspect of simulation that underpins how this generative methodology is applied in these works. Simulating the distinct types of complimentary musical textural approaches found in popular music and jazz provides a departure point for exploring ways that the generated musical layers may be successfully combined without masking each other.

The fact that many of the process embedded in the generative systems have some simple intuitive mechanism, usually a collection of variables, for interacting with their environment makes it easy for a human performer, or the action of other agents, to influence an agent's output in some semi-predictable way. The interactivity that this enables was one of the motivations for adopting an agent-based approach for this project.

One community of composers and engineers that are interested in the application of agents for the creation of music are those interested in Musical Metacreation (Eigenfeldt, Bown, Pasquier, & Martin, 2013). Musical Metacreation is concerned with endowing machines with creative behaviour. There is a focus on the degree of autonomy and agency in the operation of composition and performance systems, and the computer is viewed as an active participant in the creative process.

Musical Metacreation can be divided into a number of distinct problems:

- Composition
 - The creation of a series of performance instructions for performers
- Interpretation
 - The performance of a musical composition, producing an audio rendering

- Improvisation
 - The combination composition and interpretation in real-time
- Accompaniment
 - The situation where a live performer is accompanied, possibly performing pre-composed music
- Continuation
 - The situation where a system continues a given musical input in the same style.

The Musical Metacreation community has developed a taxonomy, which is meant as a classification system for metacreative systems. Systems are categorised based on the amount the human designer/composer controls the final musical output, which is an indicator of the level of musical autonomy given to the system. The different degrees of control that a system has over the final musical artefact can be reduced to:

- How much creative decision-making was left to the system.
- How much influence is required from a human in order for a metacreative system to perform musically.

The Musical Metacreation Taxonomy consists of seven levels, ordered from the least autonomous to the most. The levels are:

- Independence:
 - The application of some sort of process upon a musical work that is beyond the control of the composer.
 - Delegates some creative responsibility to the system.
- Compositionality:

- The process to determine the relationships between predefined (fixed) musical gestures.
- Could include initiating the playback of pre-generated or pre-recorded material, or complex signal processing that is heard as a separate textural layer.
- Generativity:
 - The process of generating new musical gestures.
 - It includes:
 - Reactive systems that use live input to generate material.
 - Feedback based systems.
- Proactivity:
 - The ability to initiate musical gestures without the presence of an external trigger.
 - This is the first level of the taxonomy that exhibits autonomous behaviour.
- Adaptability:
 - Agents can change their behaviour over time due to their own internal evolution.
 - Agents can interact and influence each other.
- Versatility:
 - Agents determine their own content without predefined stylistic limits.
 - Agents generate new compositions each time rather than creating variations of the same composition.
- Volition:

- Agents can decide when, what, and how to compose/perform.

This taxonomy is concerned with indicating the level of autonomy of a system rather than the level of complexity of a system, so it is possible for a complex system to have less autonomy than a simple one.

The Musical Metacreation group contends that improvisation provides insights to musical metacreation when considering the varying relationship between musical intent (composition) and musical practice (performance). In the case of improvisation, they are interested in the relationship between the underlying composition and its realisation, in particular, the influence that the performer/improviser has on different performances. They see a parallel in Musical Metacreation, where they consider the relationship between the composer/system designer's direct influence on the system and the differences between performances of the same work, which indicates the system's autonomy.

This taxonomy is meant to classify metacreative systems' musical autonomy, in terms of their relationship to the human designer/composer's control over the final musical result. Viewing the real-time works in this portfolio from this perspective, they would be placed into the Adaptive category or the Generative category, and one is a Proactive system. The higher-level categories, Versatility and Volition, do not align well with the aims of this project, as they place more creative responsibility on the computer system than I am interested in. For me, the aim of the agents is to realise the musical ideas within the constraints I set them. They exhibit agency in the sense that an agent is somebody or something that acts on behalf of the composer. They are, in some

sense, musical helpers. The musical output, in some way, reflects the will of the composer, not a software entity. In light of this, the Adaptive agent system is the most appropriate for this project, as it optimises the utility of the agents whilst realising the composer's musical concepts.

Chapter 3: The Portfolio Works

Aesthetic placement and stylistic predecessors

All of the compositions in this portfolio encapsulate a variety of musical approaches inside an agent-based paradigm. They range from the composition approaches of some twentieth-century art composers to the production processes of various genres of popular music, especially in a computer-centric production environment. The agent abstraction enables these approaches to be integrated into any arbitrary combination within the constraints of the agent implementation medium. Musical ideas that had an influence on the making of these works are drawn from a number of historical practices which are outlined below.

Tape Music

Along the Corridor, *A Painting in Sound* and *In Transit* are consciously descended from the fixed form tape pieces produced with electronic sound sources from the 1950s onwards, such as Stockhausen's *Studie 1* and *Studie 2*. The tape editing/assembly techniques (Dwyer, 1971) of the tape music pioneers have been replaced by digital editing techniques. An important consequence of this, in addition to the extreme reduction of labour involved, is the amount of extra signal processing that it is practical to apply. Because signal processors now exist as software, it is possible to run many reverb units, delays, equalisers, modulators and distortion devices than is typical when everything is discrete hardware. Additionally, working in the digital domain enables various

transformation techniques that are not possible with tape. For example, in *Along the Corridor*, some sounds have their own internal rhythm. Digital time stretching techniques enabled this rhythm to be synchronised with the tempo of the piece without any change in pitch.

Live Signal Processing

Many of the works, such as *Sky Castles* and *Pandan Musings/Meditations*, are examples of the common practice of applying interactive real-time signal processing to a live instrumental performance. This has been practiced in a variety of forms for over half a century, in many types of musical context. The level of sophistication varies widely, ranging from simply manipulating the controls on a collection of simple guitar stomp boxes, to complex computer-based systems controlled with sophisticated user interfaces, hardware controllers or artificial intelligence systems. Often, when working with systems of this type, a large amount of attention is focused on operating the signal processing, and one or more performers may focus on this exclusively. In these works, the signal processing control task has been delegated to autonomous agents.

Sky Castles is a signal-processing piece that is descended from the many tape-based works that construct interactive performance systems around tape delay and feedback. Warren Burt's performances in 1978 at La Trobe University, which featured very long tape delays combined with feedback, introduced me to this approach. The idea of creating echoes to produce repetitive structures goes back to the early days of *Musique Concrète* and *Elektronische Musik* (see chapters 2 and 3 of (Manning, 2004)). Rock musicians, such as the Beatles, Brian

Eno, Robert Fripp and Pink Floyd, have adopted these ideas and taken them into popular culture. Looping pedals, which include features such as tap tempo delay time and overdubbing, are now standard equipment for many people. The configuration of the signal processes in *Sky Castles*, which transform the flute sound into a slowly arpeggiated chord, is more complex than the common looping systems. It includes a number of feedback delay systems, each with different delay times, which may have other signal processing devices inserted into the feedback loops. In this way, other textural layers such as drones and harmony lines can be created, in addition to simple repeating. Robert Fripp's Soundscape system is a similar extension of the tape loop idea, which is also realised with digital signal processors, but *Sky Castles* includes the additional element of agent control.

Fixed Form Computer Music

The sound of *A Painting in Sound* and *In Transit* was influenced by the many fixed form electroacoustic concerts I experienced at the International Computer Music Conferences that I attended in 2003 (Singapore), 2004 (Miami), 2005 (Barcelona), 2006 (New Orleans), 2008 (Belfast) and 2014 (Athens). In addition to hearing the work of the many practicing electroacoustic composers whom I met at these conferences, I have also been exposed to the works of computer music pioneers such as John Chowning, Jean-Claude Risset, Max Mathews, Trevor Wishart, Curtis Rhodes and Iannis Xenakis in an authentic concert context. This music is characterised by drones, granular clouds of sound, a general lack of rhythm and melody and complex gestures with evolving timbres. This is the type of music that Smalley was attempting to help listeners to

understand when he developed the descriptive tool he calls Spectromorphology (Smalley, 1997). One important element of this style of music that was deliberately not adopted for the works in this portfolio was what Smalley calls Spatiomorphology, which involves the placement of sounds in space. All the works in this folio are in stereo, instead of any of multi-channel formats common to this genre, in order to limit the compositional focus.

Sky Castles prominently exhibits the technique of granular synthesis, which is used in many modern computer based electroacoustic works. (Roads, 2004). In this technique, a composite sound is constructed by combining many short sounds. It is a common technique in this genre, and it produces the distinct, cloud-like, sounds that have become one of the common features found in electroacoustic music.

Synthesizer Technique

Analog synthesizers, or their predecessors, discrete signal generators and processors, have featured in the work of avant garde/experimental composers prominently since the late 1940s. A number of their techniques have been applied to construct the works in this portfolio. I have come to realise that over the years, I have been more influenced by the ideas concerning shaping musical sounds, picked up through reading books and magazines, than by listening to the music that I was reading about. In fact, I have often been disappointed when I heard the sonic result I had read about. Especially influential were the interviews and 'How to' columns in *Keyboard* magazine, and its predecessor *Contemporary Keyboard*, and books such as (Strange, 1983) and (Wells, 1981). I was also influenced by personal contact with other people I interacted with in Melbourne

who worked with synthesizers in the late 1970s. Probably the most important factor that has shaped my thinking about composing with electronics was having free access to large modular synthesizer systems, such as the large Serge system that was at La Trobe University. This gave me a platform to experiment with the ideas I had been reading about, and to develop an intuitive understanding of how disparate simple processes can interact to create a controllable complexity, which eventually led to the agent-based approach. In addition to making it very natural to adopt a processed based workflow, working with the early complex analog modular systems cemented ideas about how to exploit various approaches to sequential control, which could later be applied to MIDI sequencing in the popular music context that I was working in a few years later.

Along the Corridor exhibits some specific examples drawn from these experiences, which are reflected in the processes that realise this piece. These include:

- Rhythmic interplay between textural layers.
 - This was a technique that arose as a natural consequence of the design of the sequencers and multi-track tape recorders available at the time. Data input approaches such as step-time sequencing, and cycle (loop) recording made it very natural to produce several layers in the texture that consisted of simple patterns that interacted rhythmically, producing a distinct sonic entity. This approach is exhibited in English electro-pop in the late 1970's and early 1980s, such as in the work of Tears for Fears, Heaven 17 and Peter Gabriel.

- The music produced by the BBC Radiophonic Workshop.
 - BBC music from the 1960s and 1970s, such as Doctor Who and The Hitch Hikers Guide to the Galaxy, provided exposure to, and an appreciation of, some of the modulation based inharmonic electronic timbres.
- The music produced by local electronic composers.
 - Composers such as Warren Burt, David Chesworth and Graeme Gerrard were important in that they provided a network of other people using the same type of equipment and were open to sharing their ideas and techniques.
- 1970s popular music synthesists such as Larry Fast, Roger Powell, Jan Hammer, Wendy Carlos and Rick Wakeman.
 - The work of these people, and the interviews where they discuss their work, is directly reflected in the details of the audio signal chain aspect of synthesizer programming. Techniques such as audio feedback, modulating oscillators and filters with red noise, treating the Minimoog as an exponential FM synthesizer, and paying very close attention to the fine detail of the sounds as they evolve can be attributed to exposure to these people's work.

A Painting in Sound draws on this same electronic music background, but its implementation focuses on another aspect of synthesizer programming which is very different from that of *Along the Corridor*. In *Along the Corridor*, the agent function was realised in a constrained improvisatory manner: a human performer agent manipulates a specific configuration of the Minimoog Model D synthesizer. In contrast, the creation of the agent function for each agent in *A*

Painting in Sound was closer to building a pre-programmed analog computer. The individual modules that constitute the agent interact in real-time, collectively creating the agent's contribution to the musical texture. In some ways, this is like the synthesizer approach people have historically associated with Buchla and Serge synthesizers, sometimes referred to as West Coast Synthesis. Good examples of this approach include Morton Subotnic's *Silver Apples of the Moon*, created with an early Buchla system.

Stochastic Processes

Many of the works in the portfolio, such as *A Painting in Sound*, *In Transit* and *Pandan Musings/Meditations* draw on the stochastic compositional procedures that have been prominent in the work of Cage (Cage, 2011), Xenakis (Xenakis, 1992) and others since the 1950s. The nature of computer and analog synthesiser environments means that they each lend themselves to different flavours of stochastic procedures. It is easy and efficient to implement algorithms to produce an output that exhibits specific probability density functions on a computer, as described in (Lorrain, 1980) and (Dodge & Jerse, 1997). These are based on simple pseudo-random number generator algorithms. In the analog domain, streams of random values are often created with noise generators, which can be easily constructed with transistors, and the resulting output is a truly random signal. Various types of random generators can be derived from these noise generators that exhibit different spectral densities. A very common procedure is to patch an analog noise generator into a sample and hold module. Every time a new random value is required, a trigger signal is sent to the sample and hold, and it measures the instantaneous noise signal at that time and outputs

that value. The shape of the resultant probability density function can be intuitively manipulated by altering the colour of the noise with some sort of filter. (Strange, 1983). Specific variations of this design have been packaged into modules since the 1960's, perhaps the most famous one being the Buchla 266 Source of Uncertainty module, released around 1970. My analog works rely on this approach for the creation of stochastic generators.

Several of the portfolio works create stochastic generators in a software environment with a technique I developed in 2001 (Spicer, Tan, & Tan, 2003) (Spicer, 2003). In order to create stochastic procedures that enable a controlled exploration of the state space, I adapted the analog sample and hold approach, discussed above, to sample a digital wavetable which had been filled with a random number sequence generated by the desired stochastic algorithm. This resulted in sub-sampling the random number sequence at easily controlled intervals, enabling it to be easily modified so its statistical attributes, such as its average value, could be adjusted to achieve particular goal states. This is an example of a digital process that was influenced by the experience of working with analog systems.

Interactive Composing Systems

The interactive composing systems that form the core of *Pandan Musings/Meditations* and *Momentary Diversions* are all directly influenced by Joel Chadabe's ideas about interactive composing developed in the late 1970s (Chadabe, 1984). A human performer controls high-level aspects of the piece, while the compositional algorithms make the low-level, note-to-note, decisions which are realised in real time with a synthesizer.

Momentary Diversions is an attempt to create an analog agent based interactive composing system. It could be considered to be an agent based take on Salvador Martirino's *Sal-Mar Construction* (Illinois Distributed Museum, n.d.), combined with Joel Chadabe's interactive composing idea, which realised works such as *Play* (Chadabe & Meyers, 1978). One consequence of the real-time requirement is that the complexity of the synthesizer voice patches that system is capable of realising is limited by hardware constraints. In a multi-track studio situation, the same module can be reused in as many textural layers as is required, but in a real-time context, it can only perform one role at a time. This means that there is a restriction on the complexity of the sounds the analog system can produce, without becoming prohibitively expensive. *Momentary Diversions* relies on a limited collection of synthesizer voice patches, each of which requires just a few modules. This limited sound palette can be alleviated to a certain extent by the application of external digital signal processing, but it does create a characteristic sound quality. The synthesizer voices are mainly constructed by applying various amplitude envelope shapes to triangle, pulse and sawtooth waveforms. While simple, each synth voice does have a certain distinctive presence, enabling it to stand out in the mix. Several people have said that they associate these timbres with techno music. One of the distinctive, but slightly unusual approaches to creating a synthesizer voice is one that exploits the electrical characteristics of a circuit component called a Vactrol, a type of light controlled resistor. This technique, associated with Buchla systems and the work of Subtonic, is extremely efficient for producing percussive sounds that exhibit a very natural decay. It alleviates the need for a separate envelope

generator, amplifier and, usually, filter in the audio signal chain controlled by an agent, so it is very economical in terms of hardware requirements.

One aspect of *Momentary Diversions* that separates it from many other performance systems built with modular analog synthesizers is that there is no step-time sequencing. A typical way to create cyclic generative patches is to store the basic voltage patterns in one or more step sequencers and to periodically step through the successive values. This technique has been ubiquitous since the 1960s and is the basic compositional process of several popular styles of music. Some obvious examples would be the work of Tangerine Dream, Kraftwerk and Giorgio Moroder. In the portfolio works, a decision was made to base the analog generative patches on either low-frequency oscillator or sample and hold based techniques instead of relying on sequencers. This gives an agent more autonomy to create or modify its material, because the low-frequency oscillator and sample and hold techniques more easily lend themselves to voltage control, enabling the agents to easily modify their output in response to modifications of their internal state or to the activities of other agents.

An attractive feature of the interactive composing approach is the immediacy of the systems response to a performer's gestures. The *Pandan Musings* and *Pandan Meditations* signal processing systems are attempts to create responsive systems that create the material for various musical layers with algorithmic composition techniques and realises them by processing the input signal in real-time. A feature of these works is that the sound of the instrument performance acts as both the high-level performance controller and the origin of all of the sound heard in the piece. Analysis of the human performed musical material controls the compositional algorithms, creating the high-level

musical features of the agent ensemble output, while subtle instrumental performance nuances can still have a significant impact on the sound of the ensemble output.

Motivic Transformation

Both of the works in the portfolio that feature conventional notation, the *Flute Trio in the Diminished Scale* and *Pandan Musings*, reflect the influence of Motivic Transformation on my approach to creating melodic material. I was exposed to this idea in series of university lectures, derived from Schoenberg's *Fundamentals of Musical Composition* (Schoenberg, 1970). The lectures were largely concerned with how to develop a melodic phrase with a good balance of coherence and variety. Two distinct phrase structures were explained: the period construction and the sentence construction. Motivic Transformation has been a way for me to create coherent melodic material on demand, especially in time-critical situations, such as improvising or under pressure to create alternative melodic ideas in the recording studio. The same techniques are the basis of some of the strategies that guide agents through their space of possibilities in these two works.

The flute trio is a conscious attempt to integrate the traditional techniques outlined by Schoenberg into the agent search compositional framework in a fully notated work. Motivic transformation techniques also provided the strategies to create the live instrument parts in *Pandan Musings*. Flute works written by composers such as Debussy, Fauré, Poulenc and Hindemith consciously influence these parts. The intent was to provide a lyrical

character and exhibit clear sectional fixed forms that anchor each movement, while the agent ensemble autonomously generates an accompaniment.

Popular Music Structure

One of the premises of creating this portfolio is the conscious attempt to integrate some of the processes and organisational elements that are prevalent in popular music with some techniques I found interesting from art music. The one aspect of composing and arranging that making commercial recordings really brought home to me, was an appreciation of how the combination of a number of simple musical fragments can create an emergent character of which there is no hint of when the patterns are heard in isolation. My interests in art music mainly lie in some of the generative processes that have been developed by composers, and I feel this is a fertile territory for creating the musical ideas that can contribute to the emergent musical phenomena.

The *Flute Trio in the Diminished Scale* adapts some popular music practices into a non-tonal chamber ensemble context. The popular music elements are manifest principally in the functional roles that agents take on throughout a piece. These roles correspond to search heuristics. Each agent (manifested as a flute performance) takes on different roles as the piece progresses. An agent may take on a bass role, then provide some stability by playing an ostinato, then take on a lead role or echo the part played by another agent. This is the sort of behaviour that I have often adopted, either consciously or unconsciously when I am part of a group that is collectively creating a spontaneous arrangement. Adopting this approach is one way, when combined with idiomatic knowledge, to increase the probability that the various

improvised parts will complement each other, creating a transparent arrangement.

The overall shape of *Flute Trio in the Diminished Scale* is consciously adapted from an archetypical shape of long, progressive rock pieces, particularly *Close To The Edge* by Yes and *Echos* by Pink Floyd. This is characterised by a fast, active beginning, then a slower section in the middle that eventually loses the sense of pulse, and then coalesces and builds up to a climax at the end. This type of overall organisation (fast slow fast) is also very common in the baroque and classical repertoire, especially symphonies and sonatas. The first movement of the flute trio is constructed along these lines, and the overall form of the three movements combined also exhibits that shape.

Summary of influences

A common thread that runs through this section is my interest in assembling various discrete processes into systems, which become vehicles for the creation of music. The discrete processes in themselves are relatively simple, but their interactions produce an emergent complexity that might not be quite what is expected. This reflects my prevailing interest in treating music composition as an exploratory activity. Each composition is a system that possesses properties that can be in a number of different states, and the shape of the musical work is dictated by a journey through the space of possibilities that each system can be in. The agent concept provides a degree of abstraction that is encompassing enough to integrate processes drawn from electronic music, classical music, and popular music and provides a practical context where these processes serve as navigational tools for musical exploration.

Compositions that use the Painting Approach

Four of the works in this portfolio employ the Painting Approach, described above. The role of each agent is to independently create a collection of sound objects, and these are assembled to form the finished piece. This is analogous the way an arbitrary collection of shapes could be assembled to form a picture. While it is possible to create improvisation-based pieces using this approach with, for example, a sample triggering system, all of the applications of this approach in this portfolio are in a fixed form. Each work illustrates how the agent idea can be applied to create works that fit into an established style. *Along the Corridor*, *A Painting in Sound* and *In Transit*, are all examples of fixed form electroacoustic music, often anachronistically referred to as tape music. They are typical of works historically associated with Computer Music. They are mixed in stereo, and are experienced by the audience as a fixed recording, but it is possible to add a live improvisatory element by diffusing them during playback in the performance space via a multi-channel playback system. The *Flute Trio in the Diminished Scale* is a piece of chamber music, and exists in a two forms: one for three concert flutes, and one for contra bass flute, bass flute and alto flute.

Description

Along The Corridor is a timbre-centric piece for a multi-tracked Minimoog Model D analog synthesizer. The initial brief for the piece was “Have the Minimoog explain itself”. The piece consists of a number of improvisations that exploit idiosyncratic aspects the Minimoog’s design. The improvisations are recorded, combined, edited and/or transformed in a Digital Audio Workstation.

A Painting in Sound is also realised with analog synthesizers, but takes a very different approach to implement an agent based composition model. By chance, I was reading a collection of essays by (Feldman, 2004) around the same period I read an article on Spectromorphology. (Smalley, 1997) This prompted me to think about the various ways artists apply paint, and creating their analogs in sound became the approach to developing the agent function designs in this piece. *A Painting in Sound* was created with a combination of modern and vintage analog modular synthesizers, along with some digital signal processing. Each layer is created by an agent that contained an autonomous system set up to create a distinct type of musical gesture, possessing an identifiable shape and timbre. These were designed to correspond to different types of brush strokes that a painter might employ. The strategies implemented in each agent were adapted from the sound shapes classified by Smalley (Smalley, 1997). The approach literally parallels painting, in that just as different types of brush strokes can be combined to form an abstract painting, a varied collection of distinct sound shapes can be combined to form a piece of music. It begins with the sonic equivalent of a collection of dots, then moves to various types of lines and textures, and finishes by converging on a slightly fuzzy line.

In Transit was created in a similar way, but was realised with a very different technology to form the agents and create the low-level musical material. The agents were software objects running in a music programming language / computing environment developed by Wang (Wang & Cook , 2003) called ChuckK. The agent's sounds were generated with the STK physical modeling synthesizer package (Cook, 2002), included in the ChuckK distribution. These sounds were subsequently significantly transformed by a variety of signal

processing techniques. The piece is predominantly organised around timbre and texture and does not exhibit a clearly articulated rhythmic pulse. A characteristic feature of the piece is the diffuse clouds of notes that fade in and out, which was the only element of the piece that was predetermined when the assembly of the piece was begun. The overall form grew out of the composition process, but effort was made to make sure the piece exhibited a sense of progression, as if the listener was on a leisurely journey, hence the name *In Transit*.

In contrast to the other works that instantiate the Painting Approach model, the *Flute trio in the Diminished Scale* is a fully notated, three-movement piece for three flutes. Each flute part has been composed as a discrete entity, with the composer consciously acting out the role of its composition agent. A variety of different strategies are used to generate the parts, which results in a succession of different musical textures as the piece progresses. This purely linear compositional approach results in music whose harmony emerges as a result of the coincidence of the notes being played simultaneously. There is no conscious harmonic organisation (tonal chord progressions, etc.), although, since the piece uses diminished (octophonic) scales exclusively, its symmetrical structure does impart a particular character.

Compositional approach

Every work formed with the Painting Approach involved assembling a collection of agents, each of which systematically explored its state space and sonically rendered this exploration. These renderings were then assembled to create a fixed musical structure. The differences between the works arise from

the choice and realisation of the state space exploration strategies, and the specifics of the sound production mechanism.

Hardware constrained improvisation strategies form the basis of the agent functions in *Along the Corridor*. During the pre-composition phase, certain categories of timber and synthesizer performance techniques were identified, and these dictated the strategies for creating the individual lines that make up the layers in the final musical texture. These categories also provided a handle for sculpting the overall form of the final piece.

In the production phase, a series of solo improvisations were recorded, whilst listening only to the timing reference. At least one improvisation was made for all of the identified strategies. A common element of all improvisations was that the performance only involved manipulating the synthesizer's front panel knobs to control the interactions between the electronic signal generators and processors. There was no conventional keyboard performance.

Improvisations were classified into categories that shared some common characteristic. The categories used were:

- Abstract decaying inharmonic/harmonic sounds.
- Imitative percussive sounds, such as gongs, bells and drips.
- Sounds created using overdriven feedback. These often have rhythmic and melodic patterns associated with them that are a result of the distortion and the feedback. Changes in amplitude can result in changes of pitch, as the higher level sounds force the synthesizer to produce sub harmonics.
- Vocal-like sounds that feature a highly resonant filter tuned to a fixed frequency.

The next phase of the production involved extracting segments of each improvisation to create distinct sound objects that were assembled to form the different sections of the piece. The length of these extracts varied considerably, ranging from single events to phrases lasting tens of seconds. Each section has an emphasis on different sound objects drawn from one of the above categories, and the transition from one set of timbres to the next is meant to create a sense of progression as the piece evolves. The final form can be discerned in the screenshot of the Cubase project shown in Figure 2.




Figure 2.

Screen capture of the Cubase Project showing an overview of the structure of *Along the Corridor*.

Along the Corridor employs agents at three hierarchical levels. At the lowest level, constrained improvisation strategies produce recorded performances. At the middle level, some of these improvised fragments are

selected and combined to create distinct musical objects. At the highest level, these objects are combined into distinct episodes or sections. In each case, the composer consciously enacts the agent functions of all of the agents involved. Even though the piece consists of an assembly of real-time improvisations, these are recorded one at a time, and the assembly process is out of real-time. The composer is free to take on the different agent personas as required, depending on the current production phase.

One of the low-level idiomatic improvisation strategies involved is configuring the Minimoog as a feedback system. One of its outputs is fed back into the external audio input so that the output is fed back into the filter/amplifier audio chain. The high levels of feedback force the pitch of the output signal to step through the sub-harmonic series of the oscillator's current frequency, producing arpeggio like effects. Different improvisation strategies were developed around this approach, which involved different ways of modulating the feedback gain. This included obvious approaches, such as directly adjusting the external input control and directly adjusting the amplitude envelope attributes, as well as less obvious techniques, such as creating cyclic amplitude variations by detuning two oscillators, thus producing amplitude modulation due to first order beat phenomena, where the exact output is determined by the frequency difference between the two oscillators, their relative amplitudes, and their waveforms. This technique is capable of producing a variety of phrases/rhythm patterns, where very slight changes in the front panel controls will produce quite significant changes in the audio output. The result is an engaging interactive performance system, exhibiting a good balance of predictability and unpredictability. The timing of the resulting cyclic patterns

was a bit unpredictable, so I frequently needed to apply a time stretching algorithm to synchronise the internal rhythm of the patch to the tempo of the piece.

Another low-level improvisation strategy was to create a distinct formant, setting the filter to be highly resonant at a fixed frequency, and then performing long continuous phrases of gliding tones. Many of the improvisation strategies involved developing patches that enabled the change in a single knob to result in a large variety of timbres. The obvious, overused example of this method is to manually change the cutoff frequency of the resonating low pass filter, but there are other parameters that can have a large timbral impact, such as adjusting the amount of modulation in a frequency or filter modulation patch, or adjusting the modulation VCO/noise generator mix knob. Often, easily perceivable timbral aspects of the output are dependent on phase and frequency relationships of the oscillators, over which the performer has limited control. This contributes another level of unpredictability.

At the mid-level, one strategy was to combine isolated events extracted from some of the improvisations, to produce distinct, interlocking rhythmic patterns. This parallels the idea of a drummer creating a pattern from the different elements of a drum kit. Also, at the mid-level, some disparate medium short improvised fragments are juxtaposed to form distinctly recognisable gestures, such as the collection of improvisation fragments heard at the beginning and the end of the piece. The middle level is also where signal processing that contributes a structural role, such as tempo synced echo, time stretching and transposition take place.

At the highest level, the overall form is actualized. The mid-level episodes are placed in order, applying appropriate transitions and controlling the overall mix and dynamics. This is where the piece gets a sense of progression and a sense of travelling along a corridor.

The composition approach to *In Transit* was very similar to *Along the Corridor* and is also constructed in three distinct phases.

- Creation of the basic material using a real-time improvisation system.
- Application of signal processing and editing the basic material
- Assembly of the basic material into the larger form

The impetus for creating *In Transit* came from a class demonstration I did of Probability Jammer, a probability based algorithmic performance system I had developed using the ChuckK programming language (Wang & Cook, 2003). Each student ran Probability Jammer on his or her laptop through his or her laptop speakers. The class was divided up into groups and was directed to adjust the parameters in particular ways. This created an interesting cloud of notes in the room that changed as the students performed. Creating a composition featuring this cloud effect was the aim of creating *In Transit*.

Phase 1 – Real Time Improvisations

Like *Along the Corridor*, the low-level material was created through constrained improvisation strategies. There are recordings of fourteen short solo (unaccompanied) improvisations performed on Probability Jammer. Probability Jammer contains two probability distribution generators which control the pitch and duration of notes. There are a variety of preset sounds available, produced with the STK physical modelling synthesizers (Cook, 2002). Probability Jammer has a very simple user interface, shown in figure 3. In performance, each

onscreen fader is controlled via a MIDI fader box, which effectively allows all parameters to be altered simultaneously.




Figure 3.

Probability Jammer user interface.

Before a performance of Probability Jammer, the performer needs to specify the type of probability distributions that will determine the pitch and durations of the notes produced and select the sound. Various probability distributions, taken from the standard literature (Lorrain, 1980) (Jones, 1981) (Winsor & DeLisa, 1991) (Dodge & Jerse, 1997), such as the Gaussian, Linear, Exponential and Triangle probability distributions can be applied to determine the pitch and duration of the generated stream of notes. In addition to the Probability Jammer improvisations, one other improvisation was recorded with a variation of this instrument that is altered to produce multilayered, drone-like, textures. In a few improvisations, once the instrument settings were producing a distinctive output, the audio was simply recorded for about thirty seconds, with no alterations. Most of the improvisations, such as the solo melodic line in the middle of the piece, were created by exploring different control settings until I felt it was generating something I considered to be a good starting point and

then recording the improvisation by adjusting the probability distribution parameters in real-time. My performance manipulations were guided by the attributes of the resulting melodic line, all the while aiming to create a distinctive pitch contour, with some rhythmic variation.

Phase 2 – Processing and Editing

The second phase of the composition process for *In Transit* involved enhancing the timbre of the raw improvisations (which sound very raw indeed) to create the basic sonic palate with which to assemble the piece. The recordings of the ChuckK improvisations were imported into the DAW and signal processing was applied to each recording with the aim of accentuating what I considered to be its character. This resulted in a collection of distinct musical gestures that would be recognisable to a listener, even after they had undergone further transformations in the assembly phase, because I feel that the recognition of recurring elements is important to provide a sense of coherence.

Phase 3 – Assembly and Mix

At the highest level, the work was assembled. This was an exploratory trial and error process. I felt that it was important that the final form of the piece exhibited a certain clarity/transparency. To achieve this, a number of experiments were undertaken, creating different textural layers, with various degrees of timbral contrast. Eventually, nine distinct textural layers and four distinct sections emerged:

- High and low organ like flourishes over the multilayer drone.
- Swirling “cloud-like” textures.
- A melodic passage played on the bowed STK instrument.
- A noise/percussion section.

In this phase of the construction of the piece, the texture and timbre of large groups of notes were the main concern, rather than the details of the individual notes. I thought of pitch only in terms of the predominant register (high, medium, low).

The primary concern was to achieve a sense of cohesion and progression through the piece, exhibiting a convincing sense of continuity. In the finished form, the piece starts out quite lively and progresses towards an (anti) climax, where it virtually stops, and then returns to some state similar to the beginning (a variation of a ternary form). The final form is clearly discerned in the screenshot of the Logic Pro, project shown in figure 4.




Figure 4. Logic Arrange page, indicating the overall structure of *In Transit*.

A Painting in Sound largely parallels the compositional approach outlined in the discussion of the two works above. Each low-level layer was separately recorded, and the recordings were edited and assembled in a DAW to create the

final form. The big difference is in the low-level implementation strategies. Agent performers were constructed with a modular analog synthesizer. In contrast to the agent functions in the previous pieces, where a human improvises within a set of constraints, the agent functions in *A Painting in Sound* are generative synthesiser patches. The generative patches are designed to create musical gestures with distinct shapes, which correspond to a painter's brush strokes and are self-running. They do not need to be initiated using a controller, such as a keyboard, they run continuously and need no human intervention. While there are many approaches to creating generative patches (Strange, 1983), in this piece, most of them include a clock generator to control the overall timing, and the shape of their pitch contour, dynamic contour, timbral evolution, as well as event timing, are the result of the interaction of control voltages. The generative patches in this work fall into two categories: those that employ random generators in some way, and those that make exclusive use of cyclic generators. When viewed from the agent perspective, the generative patches implement the agent function, the agents' effectors are any of the signal outputs (the most obvious being an audio output, but also include control signal outputs), and the agents' percepts are the control voltage inputs of the modules involved.

In *A Painting in Sound*, the generative processors are tuned to create distinctive musical gestures that are analogous to different types of brush strokes that could be employed by a painter. These were consciously influenced by the various Motion Processes articulated by Smalley (Smalley, 1997), in his discussion on Spectromorphology. The various control voltages were shaped to create unidirectional, reciprocal or cyclic motion as appropriate. In this piece, unidirectional and reciprocal shapes usually involved creating generative

patches based on random voltage sources, derived by some configuration of analog noise generators and sample and hold circuits. These were augmented using voltage processors such as inverters, attenuators, slew limiters and balanced modulators. Cyclic shapes were created by constructing the generative patch with combinations of low-frequency oscillators (or cycling envelope generators) that are tuned to different frequencies and output different waveforms. In every case, the resultant control voltages were applied to oscillators, filters, low pass gates, amplifiers and envelope generators to create the pitch, timbre and dynamics contours of the discrete musical gestures that make up each layer. This approach enabled a very intuitive approach to guiding the exploration of the state space of the agent. It made the composition task of setting up an agent, which could explore distinct paths through its state space while exhibiting specific attributes, to be an intuitive, interactive process. In addition to all of this structural analog control voltage processing, the resulting audio generated by each generative patch was subject to a variety of digital signal processing techniques to further shape the low-level musical output produced by each agent.

An important consideration that helped to guide the final assembly of the piece was Smalley's concept of Spectral Space, which is also drawn from Spectromorphology (Smalley, 1997). Some agents were specifically designed to provide a root or a canopy, creating a frame for various other layers to occupy. When assembling the final mix, considerations were made along the continuums articulated by Smalley, such as of emptiness-plenitude, diffuseness-concentration, streams-interstices and overlap-crossover. Spectral density was also considered in the decision-making process, and it affected decisions made

concerning the apparent distance the musical layers were from the listener, and how the sonic space was filled and layered.

The *Flute Trio in the Diminished Scale* was created in a way similar to *Along the Corridor*, *In Transit* and *A Painting in Sound*, but with acoustic instruments. The initial idea was for the agent compositional framework to encapsulate well-known motivic transformation techniques, to generate low-level musical materials that would be assembled to create the piece. In contrast to the other three painting approach works, the *Flute Trio in the Diminished Scale* involved the human composer taking on the role of four composition agents, one performer agent for each of the flute parts and one supervisor agent, who determines the overall flow of the piece. Unlike the electronic works, there were no external systems to provide built-in hardware or software constraints to restrict the state space search of compositional possibilities. The constraints to guide an agent through the possibility space in *Flute Trio in the Diminished Scale* were purely mental constructs. Because the composer is working out of real-time, it was possible to consciously apply constrained search frameworks arbitrarily. This is quite different to the other three works, where low-level agents search their state space in real-time. The real-time performance constraint imposed practical compromises, resulting in a reduction of the number of choices possible at every decision point. In *Flute Trio in the Diminished Scale*, the adoption of relatively strict motivic transformational processes as state space navigation mechanisms created a similar reduction in complexity.

The procedure to create *Flute Trio in the Diminished Scale* parallels that of the electronic pieces. Low-level activity generates a pool of musical material that

is further refined, processed or combined to become middle-level structural building blocks. The high-level supervising agent assembles these into the final fixed form.

As the human composer is sequentially adopting the role of all the agents, he is aware of what all the agents are doing. In order to decide what material to play, low-level and mid-level agents might assume various functional roles, such as creating a melody line, a bass line, contributing to the creation of a pad in different sections of the piece, as directed by the supervising composer agent. These functional roles act as heuristic state space search strategies and simplify the decision process by eliminating obviously unsuitable possibilities (pruning the state space). An example of this would be the construction of a melody. A mid-level agent will construct an eight bar melody by taking two bars of basic melodic material and transform them into a sentence or period construction. Other examples include creating a canon by making agents imitate what another agent was playing one or two bars earlier, or simply repeating a melodic figure to create an ostinato accompaniment.

A significant contribution to the overall character of the *Flute Trio in the Diminished Scale* is the sound of the half step/whole step diminished scale, as it is known in Jazz (Coker, 2010) - a member of the octatonic family of scales. This scale has several well-documented properties, such as having only two modal forms, but I chose it for this piece because its harmonic ambiguity as a result of its symmetrical intervallic structure, which creates a consistent flavour when arbitrary agent generated monophonic lines are combined.

During the initial low-level phase, a collection of two bar motifs was developed to provide the basic motivic material that would be further

transformed. Initially, the intention was to create basic motivic material with an adapted version of Probability Jammer, which would output MIDI events that could be recorded into a MIDI sequencer. Variations of this material could then be easily produced with common motivic transformation techniques. Several performances of the Probability Jammer were recorded, and the resulting recordings were reviewed and small some small fragments were identified to be further refined. This material was further processed, but the resulting material bore only a slight resemblance to the Probability Jammer output. The opening motif for the first movement is an example of the output of this process. The original part was transformed twice to create three layers and was then modified to reinforce a sense of the three lines converging on one final pitch. Some of the notes were removed so that the part was more playable.

Another approach to creating low-level material for *Flute Trio in the Diminished Scale* was a procedure that applies a pitch contour to a given rhythm pattern. It consists of developing a simple but distinctive rhythm pattern, playing only one note, and then modifying the pitches within the constraints of the diminished scale, to create a balanced contour possessing distinct high and low points. A number of transformations of the resulting basic motifs were made, which added to the palette of possible variations.

The process of assembling the large collection of small melodic fragments into a larger structure had three distinct elements. The first involved expanding some of the motifs into eight bar melodic phrases consisting of either a sentence construction or a period construction. The second involved adapting some of the patterns to create accompaniments for the melodies, and the third involved ordering these phrases to create a larger sectional form.

A non-musical element guided the fashioning of the overall shape of the piece. This shape mirrors the process that creators sometimes go through as they struggle to complete a work. The piece starts with an initial flurry of activity, which reflects the enthusiasm that accompanies the initial idea of the work. Then the piece proceeds systematically, reflecting the first stages of the construction of the project, until slowly running out of energy, getting lost, questioning the value of the original idea until, finally, after resolving all of the doubts, the work gains momentum and the project is eventually completed. The first movement, *Impetus*, exhibits of this entire structure, the second movement, *Stasis* reflects the questioning phase, while the third, *Resolution* represents the resurgence of confidence and completes the work.

Personal reflections

The Painting Approach is my default development approach for fixed form electronic music and is similar to the workflow of many other electroacoustic composers, as described in (Roads, 2015). The key element of my realisation of the painting approach is that the works are assembled out of constrained, usually monophonic, improvisations or are algorithmically generated. As mentioned earlier, improvisation is an integral part of my daily modus operandi and formed a large part of my commercial music practice. The improvisation approaches adopted to create material for *Along the Corridor*, *In Transit* and *A Painting in Sound* were largely dictated by my choice of the implementation technology underlying each performance system. Such systems usually have two distinct components: a sound generator and a controller/interface, and often there is a trade-off between a player's preference

for particular sound qualities and for controllability. My priority is usually biased towards the ability of the performance interface to facilitate expressive tactile control, which reflects my desire to create conditions ripe for spontaneous performance possibilities. All of the fixed form electronic works in the folio required me to invest a significant period of time in practice and experimentation before I could achieve a level of performance quality that I considered acceptable.

Two of the fixed form electronic works were realised with analog synthesizers, but they were of very different designs, so each demanded a very different improvisation strategy. The common element of both strategies was the emphasis on shaping the performance via continuous controllers such as potentiometers, faders and pressure sensors. The other fixed form electronic work features algorithmically controlled software physical modelling synthesizers. Control signals from a MIDI fader controller were mapped to the appropriate compositional and synthesis algorithm parameters through an iterative process of trial and error until I achieved what I felt was the right balance between computer and human control.

Sonically, all of the fixed works reflect my preference for music that has a distinctive electronic sound. By this, I am referring to several sonic characteristics that are typical of music in which electronics are an inherent part of the production process. These sonic characteristics can be broken down into two categories; those that are derived from electronic tone generation techniques, and those that emerge from processing during the assembly process. In his discussion of Spectromorphology (Smalley, 1997), Smalley talks of source bonding and gestural surrogacy in relation to tone generation. Although

Spectromorphology mainly describes intrinsic sound events and relationships within electroacoustic music, Smalley argues that an extrinsic foundation in culture is necessary for the intrinsic sounds to have meaning. Source bonding represents this extrinsic-intrinsic link. It refers to a listener's ability to perceive a real or imagined extrinsic connection between a sound and its cause. Smalley states that the ultimate goal for a listener is to be able to interpret these links in terms of meaning, expression and psychological significance. Personally, I am interested in sounds that are not easily identified as having natural sources. I do not usually want to create obvious sonic mappings of electronic sound to real-world causes. Smalley refers to the differing degrees of mapping extrinsic cause to the musical sound as gestural surrogacy, which indicates how remote a sound is from our perception of its physical cause. He classifies our perception into five classes, ranging from proprioceptive perception (the primal gesture) of the actions that will eventually cause a sound through different orders of surrogacy, to remote surrogacy, where the source and cause of a sound are unknowable. In Smalley's classification scheme, I tend to favour sounds that are either remote or third order gestural surrogates, where a causal gesture is inferred or imagined. A consequence of this is that I tend to favour generating sounds with either electronic oscillator circuits or by transforming real-world sounds into something that is not recognisably related to their source.

At the sonic assembly stage, I have developed a preference for particular signal processing techniques that enable me to tame the wilder amplitude and spectral components of the various sonic elements and help me to achieve transparency in the resulting mix. I always try to make sure that all of the sounds exist in their own sonic space and do not accidentally interfere with each other.

This preference has meant that I have developed a repertoire of habits for shaping sounds with compressors, parametric equalisers, modulated delays and artificial reverberation units that implant a particular sonic signature on my work. Frequently occurring signal processing practices include:

- Compressing the dynamic range of the signal with a compressor design that also adds some colouration to the sound, usually making the sound seem fuller.
- Applying a convolution reverb algorithm to the audio signal, usually emulating a hall or a spring reverb. The reverb processor is installed either on an auxiliary bus, if it is being used to place several sound sources in the same space, or is inserted into the mixer channel strip if its role is to isolate or defocus a sound.
- Equalising a sound to reduce muddiness and avoid sounds masking each other. I almost never boost frequencies; I mainly attenuate specific frequency bands.
- Apply short delays modulated (Chorus, Flanging, Leslie Cabinet simulation) to create a sense of animation to a sound.
- Apply echo and feedback to either provide a sense of rhythmic pulse or a sense of blurring the sound.

Together, these processes enable sounds to be combined to form a cohesive and transparent mix. The natural dynamics, frequency balance and sonic blending that would occur if several sources were to sound in a shared acoustic space are usually deliberately distorted.

My preference for the sonic characteristics of the electronic sound assembly process is also present in my acoustic instrumental work.

Electronically amplified sound has become so natural for me that it took me years to notice this bias. I always perform in an amplified context, so consequently that is the way I perform *Flute Trio in The Diminished Scale*. (I decided not to specify that the piece must be performed amplified because I think that amplification could be a barrier for other performers). This sonic preference was unconsciously reflected in the way I chose to record the piece, which captures this amplified flavour. The recording does not sound quite like a recording in a natural acoustic space because it was recorded the same way that I amplify the flute in concert situations; a dynamic microphone is placed very close to the embouchure hole of the flute, capturing the direct flute signal and minimising the sound of reflections from the room. Artificial reverberation, equalisation and compression were added during the mixdown, typical of the way it would be in an amplified concert setting. I could take this electronic production idea further in future performances, and add a fourth performer whose role would be to distribute the flute sounds via a multichannel speaker array. This could make the emphasis on the counterpoint between the three agents clearer, and might decrease the expectation of some sort of harmonic progression that some listeners seem to expect.

One inherent aspect of the workflow of the Painting Approach is the forced separation between controlling an ongoing real-time process, when creating the low-level surface material, and the assembly of this material to create the higher-level structure. I find that this separation between the real-time and non-real-time processes is very comfortable, and is inherently much less risky than the Goal-Led and Pruning Approaches. The Painting Approach is a reframing of the traditional tape music assembly process, with the additional

element of clear consideration of how complementary generation strategies can provide an extra level of security. Having a higher probability of a successful completion of the work has obvious advantages in terms productivity, but, because the result is a fixed form, I feel a bit too comfortable when the works are presented in a concert situation. While I very much enjoy developing the low-level components, and the assembly process, publicly performing these pieces does not give me the same sense of anticipation that I get when performing with the other two frameworks. I think this also affects the expectations of the audience. An audience that is aware that the music being presented to them has been carefully predetermined and refined will have a different set of expectations to an audience that knows that it is witnessing the spontaneous creation of the music in real-time. The other approaches have a greater element of risk because there is more uncertainty as to exactly what will happen. The fact that most of the decisions have already been made generally gives each performance of the Painting Approach pieces less sense of occasion in my mind, and I think in the mind of the audience because the only uncertain elements are at the detail level. Currently, only basic elements, such as manipulating the volume or equalisation in real-time to accentuate the dynamics to suit the performance space, are possible when the works are pre-mixed to stereo. In the future, like many electroacoustic composers, I plan to insert a more of active performance element into my fixed form pieces. Instead of producing a stereo mix, I will premix the piece into several intermediate tracks, a process usually referred to as creating stems, and will then manually mix these during the performance. This methodology will allow more spontaneity, will better enable

me to take advantage of any multichannel speaker options available and will facilitate adapting the mix to the acoustics of the space.

While assembling these works, a conscious effort was put into creating a sense of sonic narrative. The non-real-time aspect of the assembly process enables each work to be carefully constructed with consideration to the way sonic events connect to each other and cumulatively form phrases and sections. I am acutely aware of the role that memory plays in the listener's ability to make sense of real-time phenomena like music, so I focus on the construction of perceivable patterns of sounds that connect to each other or group together, in some way. Smalley (Smalley, 1997) and (Roads, 2015) discuss how many acousmatic composers employ recognizable sounds to create some sort of concrete narrative thread, but in my work, I prefer not to use any sounds that are obviously referential to sounds in the real world. (Inevitably, some of the electronic sounds evoke natural sounds.). *Along the Corridor* does contain some sounds that are reminiscent of water dripping, and a number of sounds have a metallic character, but these have no extrinsic meaning in the sonic narrative that I am conscious of. Their distinctive timbres merely help the listener to mark out the territory that they are currently traversing in the imaginary sonic topography.

Compositions that control the ensemble by pruning

Description

Four compositions in this portfolio are examples of the Pruning model of the multi-agent framework. Performer agents independently explore their state space, whilst producing musical output that reflects their trajectory. A high-level

agent acts as a gatekeeper and determines which agents the audience will hear. *Sky Castles* has a fixed form, while *Sonic Escapade* and *Momentary Diversions* are improvisation environments. All are intended for live performance.

Fixed form pieces

Sky Castles features a series of conventionally notated long notes performed on the flute, which are transformed by a system of delays and pitch shifters producing a dense series of evolving chords. As the piece unfolds, a number of synthesised sounds are progressively added in a predetermined way.

Improvisation Environments

Momentary Diversions is a real-time performance system for modular analog synthesiser, augmented by digital signal processing. It is characterised by the approach it takes to the way the modules are connected, rather than the specifics of the particular instance of the piece. The available hardware is distributed to instantiate at least six compositional agents. The musical contribution of each agent exhibits a particular characteristic, which reflects the strategy that is implemented in the agent function. *Momentary Diversions* shares many surface characteristics of *A Painting in Sound* because it is realised with the same hardware, but *Momentary Diversions* is conceived as a specific type of interactive composing environment, rather than as a fixed form piece. This implies that there may be many realisations of the piece, which could be created with a variety of synthesiser hardware. The example included in the portfolio, which was recorded live in the studio, has an additional textural layer from an electronically processed flute part, typical of my public performances.

The non-electronic example of the pruning approach is *Sonic Escapade*. Originally, *Sonic Escapade* was a conceptual framework for creating improvised ensemble performances. It was based on the premise that players act as an autonomous agent, possessing a clearly articulated collection of compositional strategies with which to spontaneously create their individual parts. It has since turned into an ensemble with a repertoire of one work, and the ensemble performs regularly in Singapore in performance situations ranging from small private workshops to large public events. The improvisation framework is designed to accommodate the skills and stylistic preferences of young musicians drawn from the different musical communities that exist in Singapore.

Compositional Approach

Each composition discussed in this section implements the Pruning model in a distinctly different way. The composition model is implemented in *Sky Castles* with three distinct agent types that control different aspects of the piece. Some agents exist as ChuckK software objects, some are manifest in the DAW's automation system, and one is the human flute player.

The score creation agent to generate the pitch series for the flute was implemented as a ChuckK program similar to Probability Jammer. The flute player needs to instantiate two agents simultaneously. One agent adopts the simple strategy to perform the sequence of notes presented to him, in written order, for the duration of a full breath, with short pauses in between. The pauses in between notes are quite significant, as they provide 'windows' in the texture to feature the details of the signal processing, as well as leaving space for the layers

of synthesisers. Whilst playing these long flute notes, the performer must instantiate the other agent. The agent strategy for this agent is to hum a note. Combining the two agents essentially deploys a common timbral technique used by rock and jazz flute players, most commonly associated with the band Jethro Tull. The pitch contour of the humming part is improvised, and may vary through of the duration of the note. The pitch may glide towards the flute pitch, may hover around it, or may glide around, depending on the player's preference.

The automation system of the DAW implements several simple reflex agents whose only percept is the current playback time. These agents shape the resulting output by determining what is audible. In *Sky Castles*, the automation system manipulates a relatively complex system of delays and pitch shifters, by interpolating between successions of goal settings that are specified to occur at particular times. This is illustrated in figure 5. The real-time flute audio signal is sent to nine different signal processing sub-systems, each consisting of some combination of plug-ins to create delay, pitch shift, modulation (such as chorus/flanging) and distortion. The automation system dynamically controls the output level of the signal processing, the synthesizer channels, the amount flute signal sent to the signal processing channels, and the amount of feedback in some sub-systems.




Figure 5

The screen capture of the automation tracks showing the trajectories of audio levels of each track in *Sky Castles*.

A number of synthesiser parts are gradually added to the texture as the piece progresses. These are recordings of an ensemble of software agents that are implemented as ChuckK objects. The pitch and duration values are determined by a Gaussian probability strategy, the sound is produced with the STK physical modelling synthesisers combined with a granular processor, which produces a cloud-like effect.

Improvisation Environments

Momentary Diversions is another improvisation environment, this time for a modular analog synthesiser. It adopts an approach to constructing synthesiser patches that instantiate a collection of autonomous musical performer agents. Each agent contributes a sonic layer to the musical texture. The final music output heard by the audience is determined in real-time by controlling the output level of each agent with a mixer. This means it is a composition where the ensemble is largely controlled by pruning.

Based on their compositional strategies, these agent designs can be broadly classified into two categories: stand-alone generative patches and interactive patches. These, in turn, can be further subdivided. Generative patches can be grouped into those that involve random processes and those that are cyclic. Interactive patches can be similarly divided into those that feature random a random element and those that are largely deterministic. This combination of strategies allows each agent to exhibit distinct and complimentary musical characteristics, enabling each of them to carve out their own space in the sonic spectrum. Some examples of typical strategies include producing low drones exhibiting complex movement through alterations of their harmonic content, creating numerous short sounds that collectively exhibit a sense of unity, creating layers that exhibit cyclic properties that manifest at two or more hierarchal levels, creating interlocking rhythm layers and creating layers of sustained gliding tones.

A foreground layer is implemented as an interactive patch that includes a significant random element with a continuous control mechanism. This type of patch ensures a sense of two-way interaction between the agent and the human

performer because the human performer will not know precisely how the agent will respond before the gesture takes place. To achieve this, I use a pressure-sensitive touch sensor system, similar to those developed by Serge and Buchla in the 1970's. This produces a gate signal every time it is touched, plus a control voltage that is proportional to the surface area of the finger in contact with the surface. The gate signal triggers the generation of two new random control voltages that typically control the frequencies of two oscillators configured to modulate each other in some way, producing a different complex timbre for each gesture. The pressure sensor typically controls the volume of this signal. The result is an expressive system whose output is a collaboration between the agent and the human performer. Additional melodic lines can be added in whatever way is convenient. This could easily be a conventional instrument played by another musician, enabling an ensemble improvisation setting.

An important aspect of *Momentary Diversions* is that because everything is generated in real time with modular equipment, it is easy to set up agents which can have sensors that gather input from other agents. Various types of interactions between voices are possible via control voltage inputs of the modules concerned. This means *Momentary Diversions* can be a true multi-agent system, demonstrating autonomy that is achieved by the complex multi-agent interactions that can take place.

Like *Momentary Diversions*, *Sonic Escapade* relies on the choice of complimentary musical material creation strategies to enable the combination of independently produced material in a coherent way. The significant difference is that the agents in *Sonic Escapade* are people who are drawing on their musical backgrounds. In order to accommodate the expected variety of idiomatic music

practices that different musicians would bring to the ensemble, a set of guidelines, referred to as the Rules of Engagement, was developed. These guidelines set out the parameters of the piece and are what separate *Sonic Escapade* from just being an eclectic jam session, in the same way Zorn's rules constrain the performers in his game pieces, such as *Cobra*. The Rules of Engagement articulate how musicians should interact, outline the type of strategies they could adopt to create their part, specify a minimum set of instrumental resources that have to be available and outline the approach to creating the overall form. At the macro level, three inter-agent interaction models are articulated which apply when a player is either leading the ensemble or is being lead. Agents will usually cooperate, but they may be deliberately disruptive, especially when they feel it is time to move onto another musical episode, or agents may occasionally ignore what the others are doing for a short while. Rehearsals of this piece largely consist of learning how to abide by these rules, and a large emphasis is placed on refining the performers' listening, communication and musical interaction.

Personal reflections

The Pruning Approach lies somewhere between the static fixed form oriented Painting Approach, and the dynamic system that characterises the Goal-Led approach. A very basic process that can be considered a version of pruning is frequently applied to create contemporary commercial dance music. Many tracks, for example, *Lucky* by Daft Punk, can be thought of as being constructed by the superposition of many layers of cyclic patterns of various lengths. What the audience hears is determined by turning tracks on and off with a mixer at

various times as the song proceeds. Loop-based software, such as Ableton Live and Garage Band, are optimised to make the creation of this type of music very easy. This methodology, which is characterised by the selection, juxtaposition and superposition of cyclic musical material, can operate at different hierarchical levels.

The difference between this common commercial approach and the Pruning Approach works in the folio boils down to the dynamic nature of the material that makes up each textural layer. Instead of simply repeating cyclic patterns, the low-level agents produce the musical material in their textural layer instantaneously as they explore their state space via their clearly articulated strategy. This means there is a decrease of the amount of predictability in the music, from the audience's perspective. As a result, more effort from the listener could be needed to process the piece because of the reduction in repetitive material. In addition to the rudimentary pruning mechanism of turning tracks on or off, potential exploratory paths in the state space search are being pruned because the musical material for each layer is generated in real-time and each step along the path precludes all of the options along the paths that were not taken. The real-time nature of the decision-making process also makes it is possible for a human or other agents to interact with the generation mechanisms, thus influencing their path through the state space. This gives a composer more options to guide the evolution of the musical output, potentially resulting in opportunities for greater spontaneous expression. *Momentary Diversions* has obvious high-level control mechanisms which can spare the performer the responsibility of low-level, moment-to-moment note-level decisions. These give the performer the option to exclusively concentrate

on shaping higher-level structures if desired while keeping mechanisms accessible that allow the performer to alter the details of the note generation strategies whenever they see fit. The multi-level control mechanism is designed to help the performer make the high-level trajectory of the piece more apparent to the audience while maintaining the option of low-level interaction. I feel that this is an important element of the agent approach because it lets the performer pursue unexpected options that may emerge from the ensemble, creating unanticipated avenues for expression and opportunities for communication with the listeners. The mindset of the performer is in an exploratory mode where there is limited prior low-level knowledge of the musical terrain that will be traversed, and in each performance, new possibilities are presented. As a consequence, the function of rehearsing the piece is to help the performer effectively interact with the control mechanism and to be tuned to the flavour of the various agents' output, which should help the performer better exploit any opportunities that are presented and take the listeners on a unique journey. Because a performer hears each instance of the work only during a performance, I feel that the inherent uncertainty helps to make a generative agent performance more intriguing for the audience.

Sky Castles takes a different approach to implementing the pruning concept. The performer contributes the raw materials for the foreground layer, which contain a constrained improvisatory element. The other layers are either deterministic transformations of this constrained improvisation or are simple recordings of generative systems that are pre-programmed to become audible at specific times. The pre-programmed mix frames of the macrostructure. This is a top-down compositional approach where the form is specified at the macro level

before the piece begins, and the lower level details are filled in during a performance. Because the structure is so clearly articulated, all of the performances have a large degree of resemblance, and this predictability means that the environment is relatively risk-free for the performer. It is very likely that the audience will experience the intended musical outcome. Performing this work is similar to performing well-rehearsed music with a fixed score. I think this process is analogous to traversing a familiar path: you proceed forward along the route, with only a few variable options. Because the way is preordained, you are free to concentrate on the low-level variables instead of overall the direction, and your pre-knowledge of the future can inform your instantaneous decision-making.

Momentary Diversions shares characteristics with the work of many of the people currently participating in the current resurgence of interest in modular analog synthesizers. Some well-known composers that work with similar equipment and adopt similar approaches include Todd Barton (Barton, 2016), Keith Fullerton Whitman (Whitman, 2016) and Richard Devine (Devine, 2016). A significant factor in the musical result produced by working this way lies in the choices made by the composer when assembling of the synthesizers themselves. It is the equivalent of deciding the instrumentation of an acoustic ensemble work, and I think is something that is overlooked by most listeners because I do not think they are not aware of the huge variety of possibilities that are now available. When I first began working with modular synthesizers in the late 1970s, a composer was usually restricted to whatever was available in a given studio, but now, it is not uncommon for the composer to own an idiosyncratic collection of modules drawn from the hundreds that are available, and to

assemble a synthesizer specifically for a particular work. Unlike many composers, I choose not to include any sample playback modules in my systems. This has a significant consequence on the music I produce because it precludes the category of sounds that are created via voltage controlled digital transformations of natural sounds. While I think this is a good approach conceptually, I find the sonic results less appealing, possibly because these sounds are currently so ubiquitous. Consequently, I focus my synthesizer work more on analog processes implemented via analog processors. Digital signal processing is usually applied at the mixing stage via plugins or in signal processing software I have written myself. Like many synthesizer-based composers, such as those mentioned above, generative patches are the foundation of a lot of my synthesizer work and this influences my choice of modules. When I assemble the synthesizer, I tend to favour fairly dense modules and configure them as relatively small, comprehensible systems. Whilst developing the *Momentary Diversions* approach, I noticed that the agent paradigm produced a significant change in my workflow. Usually, composers that perform live with modular synthesizers maximise the capabilities of their hardware by extensively changing the patching during a performance, and they trigger many events manually. This approach is common in the work of the composers mentioned above and has always been central to my approach to working with modular synthesizers. When I adopted the agent centric approach for this project, I have felt much less need to do that. Conceiving synthesizer patches as agents that implement specific improvisational strategies has had several benefits for me as a composer/performer:

- The available hardware capabilities can be maximised by providing clear demarcations between the components of separate subsystems and the ways they can interact.
- The agent approach enables a more deliberate performance practice by placing clear boundaries on the improvisational possibilities.
- The deployment of agents helps to make the results more coherent by making it easier for the various textural layers to complement each other.

The consequence of this is that the daunting act of modular synthesizer performance is much less stressful, allowing me to concentrate on shaping the musical structure at multiple levels and making it easier to create a sense continuity and flow. Because of this, I now find that the live performance works created via the Pruning Approach can now easily have several textural layers present simultaneously without them masking each other, and the works often exhibit more gradual textural changes. This is a contrast to the popular approach of building narrative structures by juxtaposing different short gestures produced by different patches. I now feel that there is a danger with the agent approach that the textural changes become too gradual and predictable, thus lacking an element of surprise and contrast, so I consciously guard against this.

Compositions that lead the ensemble by setting goals

Description

The software agent design deployed in *Pandan Musings* and *Pandan Meditations* has the ability to adjust itself to achieve certain target average pitch and duration values. These compositions are based on multi-agent systems that

include a mixing/pruning capability as one of their functions, in addition to the ability to autonomously adjust their musical output to exhibit specified goal attributes. The choice of a system's current goal state is determined by analysis of a live audio input stream. This gives a composer an element of predictability when designing this type of system, to some extent, because the system will exhibit specific behaviour when presented with specific audio material. A distinctive aspect of these systems is that all of the sounds heard by the audience are altered versions of the audio input stream, so the sound presented to the system is both a control source and the raw material for all the audio produced. This provides a tight coupling between the input and output of the system.

Pandan Musings is a series of fixed form pieces for Flute (or another woodwind) and Computer, and is notated using conventional notation. I have written five pieces with this approach, three of which are included in the portfolio. In each piece, the computer runs a specific Pd patch, which implements a particular multi-agent interactive composing system. This system transforms the sound of the instrument performance in a variety of ways, specific to each work, creating a multi-layered musical texture. Each layer contributes its own acoustic signature, which compliments the other layers present. An important aspect of the piece for me is the expressive relationship between the human performer and the response created by the multi-agent system. Each Pd patch has its own distinct flavour that is the result of the combination of strategies embedded in the agents and the type of audio transformations that are applied to the live audio signal.

The Pd-based multi-agent systems developed for *Pandan Musings* are also free improvisation environments in their own right, and I refer to improvised

performances with these systems as *Pandan Meditations*. Because all of the systems employ substantially different algorithms, each one provides a distinctly different interactive experience for the performer. The systems are easily adapted to work with any type of sound source and have been played with a range of instruments, from traditional acoustic instruments such as flute, saxophone and bassoon, to sound sculptures constructed out of found objects with piezoelectric pickups. An example of one instance of Pandan Meditations, recorded with a saxophone, is included in the portfolio. For this recording, I chose a very experienced improviser, Dr Timothy O'Dwyer, to work with the most complex interactive composing system. I am very familiar with Dr O'Dwyer's work, having witnessed several performances in a variety of contexts. I knew he has a very highly refined ability to make creative decisions in real-time, has an acute awareness of what was happening in the surrounding musical environment, and has a wide repertoire of performance techniques at his disposal. These attributes are just what is needed to work with this type of interactive composing system. Just before the recording, I prepared him by explaining how the multi-agent system worked and gave a brief overview of the strategies embedded in the agents and their sound production mechanisms. I asked him to respond to the agents output as if it were an ensemble and told him to "Play your natural game" the way I had witnessed him do in the past. After the recording, he told me that he enjoyed the experience. He could sense the system responding to his actions, and he, in turn, felt the need to respond to the agents contributions.

Compositional approach

The multi-agent systems deployed in *Pandan Musings* and *Pandan Meditations* evolved out of many years of improvising with multi-agent interactive composing systems. In an improvisation context, my objective for the multi-agent system development was always to achieve what I considered to be an appropriate balance between the predictability and the amount of surprise in the system's response. This balance is important when trying to create a sense of collaboration with an artificial agent ensemble. A distinguishing factor of *Pandan Musings* that differentiates it from my previous work (Spicer, 2005), is that the element of uncertainty has a different role to that when working in a free improvisation context. The basic premise of *Pandan Musings* is that the live performer will be affected by the musical contribution made by the multi-agent system, whilst interpreting the pre-composed score. The multi-agent system is configured and fine-tuned, during the composition phase to strike a balance between consistency and variety from performance to performance. This variation in the musical accompaniment encourages a fresh interpretation of the fixed score every time it is played.

To create an intimate, expressive feedback loop that engenders a feeling of influence by the human performer on the multi-agent systems output, a collection of signal processing systems for transforming the live input signal was embedded into the performer agents. This creates a situation where the multi-agent system controls the pitch, duration and timbre of the sonic events that make up the textural layers that it contributes, but enables expressive control which is directly linked to the dynamics, articulation and phrasing of the human performer.

There is a non-musical aspect to *Pandan Musings*. The title of the composition alludes to the Pandan Reservoir, an artificial reservoir created by flooding a coastal mangrove swamp in the western part of Singapore in the 1970s. It is at an intersection between several different ecosystems. On the east side, there is still an area of mangrove and a short estuary leading to the sea. The north side has a high-rise housing estate, while to the south and west there is a dense industrial estate containing a variety of businesses including logistics hubs, small-scale manufacturing and some oil and chemical processing. This combination results in an unusual intermingling of natural and man-made acoustic phenomena. *Pandan Musings* attempts to reflect this interaction between the sounds of nature and the sounds of human industry to some degree, in a non-imitative way. The sound of the woodwind instrument specifically played without employing any extended performance techniques represents the natural world, and the multi-agent system contributions represent the man-made sounds.

The procedure adopted to create each movement consisted of creating an initial draft instrument part, and the assembly of a number of pre-made agents to establish a preliminary multi-agent system. This initial combination is refined through an iterative process of testing and adjusting the system response. During this phase, the instrumental part, and the agents, may be substantially modified. The instrumental melody lines were deliberately fashioned by applying common motivic manipulation techniques, in a similar way to some of the melodic material in the *Flute Trio in the Diminished Scale*. The melodies of each movement are differentiated by a distinctive rhythmic profile, and each is built

on a particular scale. The resultant melodic material consciously adopts some idiomatic characteristics common to early twentieth-century flute repertoire.

Agent Design.

Pandan Musings and *Pandan Meditations* have two broad categories of agent in the multi-agent interactive composing systems: performer agents and controller agents. Performer agents transform the input sound in some way and create an independent layer of the musical texture. Controller agents either determine which agents the audience hears, or specify goal states to drive the behaviour of the performer agents.

Performer Agents

Performer agents continually monitor the live audio signal. The instantaneous pitch and amplitude data, as well as cumulative statistical data, are percepts for the agent and determine an agent's behaviour as it transforms the live input signal. Performer agents fall into two broad categories.

The first class of performer agents simply apply common signal processing effects such as echo, chorus, flanging, spectral delay and ring modulation to the live signal, and their primary role is to shape the overall ambience of the ensemble output. Most of these agents produce timbres that are recognisably transformations of the live input signal.

The second class of performer agents plays a more structural role. They transform the live audio signal into new, distinct musical lines and are optimised to perform particular musical functions. They may significantly alter the pitch and timbre of the input waveform, so much so that they don't seem to be recognisably derived from the input signal at all. Signal processing techniques

embedded into this type of performer agent include pitch shifting (implemented by a granular synthesis algorithm), modulation synthesis and the Karplus-Strong plucked string algorithm. These processes are capable of transforming the live signal into drones, parallel harmonies, arbitrary melodic lines, ambience effects and clouds of sound. Because the resulting audio output is still reliant on the live input signal, the human performer maintains an element of control over the entire ensemble sound through choices made with respect to dynamics and articulation.

The basic operating principle of this type of performer agent is:

- Determine the current pitch of the live signal.
- Determine the next pitch the agent will play.
- Transform the live signal so that the output has the required pitch.
- Determine the length of the new note.
- Play the note for the required duration.

The structure of this type of agent is shown in figure 6.




Figure 6. Structure of a performer agent.

Performer agents implement a variety of algorithms to determine the pitch and duration of their output. These include stochastic algorithms, such as random number generators exhibiting Gaussian and Exponential probability distributions (Dodge & Jerse, 1997), and tracking and evasion algorithms (LaMothe, 1999). The algorithms provide the means to realise strategies that

implement a variety of musical functions, such as pads, drones, clouds, interlocking layers etc.

An important aspect of the design of the performer agents in these works is their ability to adjust themselves to exhibit specific musical behaviour. This allows agents to vary their functional roles as the piece progresses. Because they are implemented as finite state machines, this ability to change their note generation strategy can be achieved by agents autonomously modifying their internal states in order to meet externally specified targets. In this design, the specification of two targets, the average pitch, and the average duration leads the musical output of agents. The mechanism to achieve this is a simple form of gradient descent learning, where each agent periodically calculates the current state of the system and compares this with the target state. This produces an error measurement that is used to alter the internal state slightly to reduce this error. By repeating this process several times a second, agents soon converge on the target behaviour.

Controller Agents

The multi-agent systems in these works contain two high-level agents that affect the overall musical output. One of these controller agents specifies the current pitch and target parameters of the various performer agents. The other agent determines which performer agents are heard by the audience, essentially carrying out a pruning process. Both of the agent's agent functions are implemented as finite state machines, which can be in one of eight states. The current target state for the system of agents is determined periodically by an analysis of the live input signal and is the mechanism that enables the live audio

to shape the musical output at the macro level. Each of the eight states in the goal setting agent is configured in the composition process to implement different note creation strategies, providing a different path through the agent's state space, thus producing a different musical effect. Similarly, each of the eight states in the mixing agent specifies a different blend of performer agents heard by the audience. When creating a new composition with this type of system, a composer needs to specify the details of the eight different target states for every performer agent and for the controller agent that controls the mix. In operation, these agents periodically analyse the audio input stream and use the results of this analysis to set the agent to its new state. The fact that this is a deterministic process is what endows this type of system an element of predictability. Similar musical input gestures tend to produce similar musical ensemble output. This enables the composer and the human performer to learn how to shape the response of the agent ensemble.

Restrictions and possibilities

The interactive composing systems in *Pandan Musings/Pandan Meditations* have been tuned to achieve my compositional aims of this piece. This involved imposing many constraints in order to create a particular musical character in system output. Another composer could take the same system and produce something quite different. For example, any type of input waveform can drive the system. The multi-agent system can respond to any type of input and its response can be easily tuned to exhibit particular behaviour. *Pandan Musings* was tuned to get a particular sound when driven by woodwind instruments, particularly the standard Boehm flute. This meant that the analysis algorithms

can be optimised to work effectively with this type of signal. I have used these systems with many types of sound sources, including voices, saxophone, bassoon, analog synthesizer and sound sculptures made of collections found objects. It is also possible to adapt this type of system to work with multiple audio inputs, where one audio signal provides the raw material for the audio output; another signal drives the analysis to determine pitch and duration targets, while another audio signal can determine the blend of performer agents that are heard by the audience.

Another effective application of this type of multi-agent systems is as an interactive feedback instrument. The system output becomes the live signal input via a unidirectional microphone. This creates a feedback loop which can be controlled by adjusting the position and orientation of the microphone in relation to the speaker system.

In the *Pandan Musings* recordings included in the portfolio, the pitch of the system output is constrained to a particular musical scale. This restriction was a deliberate compositional choice that required additional programming to achieve. Typically, when analog synthesizers and found objects generate the live signal, the scale mapping system is turned off, allowing any arbitrary frequency to be generated.

There is a restriction that is inherent to the system design, which is that the timing of the output of the multi-agent system is quantised to a time grid. In some movements of *Pandan Musings* this was adjusted so the system provides a clear pulse for the performer to play with (or against). Again, the timing resolution is a compositional choice. It is easy to set the temporal resolution fine enough so that the time quantisation is not perceivable.

Personal reflections

Of the three frameworks I have adopted in this project, the Goal-Led approach is the least conventional and is the one I most prefer in a live performance context. My approach to creating goal-driven systems is inherently real-time and creates immediate opportunities for spontaneous expression by the performer at the note level. Producing music in real-time has an element of uncertainty that I find helps me, as a performer, to cultivate the very focused mental state. I interpret this to be the flow mental state described as “the state in which people are so involved in an activity that nothing else seems to matter” (Csikszentmihalyi, 1990, p. 4). While I am performing with goal-driven interactive composing systems, I find the sense of cause and effect to be very engaging and I am totally focused on the present moment. Because it is so natural to focus at the note level when interacting with these systems, I do find that it can be a challenge to simultaneously create a higher-level structure and to engage a listener at the macro level, despite the presence of the high-level goal setting mechanism.

One of the original conceptual aims for writing *Pandan Musings* was to overcome this issue by articulating the macro structure in a conventionally notated instrumental part. In addition to outlining an overall form, the notated part creates a constrained circumstance where the performer can concentrate on capitalising on opportunities for spontaneous expression without having to decide what notes to play. The performer's role is to interpret the written score in real-time within the context of the musical output of the interactive performance system. Because the system responds to the performer's nuances, a feedback loop is created, giving the performer some degree of direct control. The

combination of stochastic generative processes and intuitive performer interactive control means that the musical result is different each time the piece is performed. The audible result should clearly be a collaboration between the human and software participants. Each movement of *Pandan Musings* has an instrumental part that is similar in style, but the interactive composition systems vary considerably in their scope and complexity. I have always considered the main compositional activity in this piece to be the development of the software environment. The software environment determines the nature of the feedback system and hence determines the expressive possibilities available to the player. Reviewing the finished work, I think that this methodology seems to work best in the first two pieces. In both of these, the interactive composing system realises conventional signal processing transformations on the live instrumental signal, and the interplay between the live performer and the agent system is clear. This makes it easy for the player to interact with the feedback network in an expressive way. In the later movements, the more complex software environments apply signal-processing techniques that create audio output that is much less obviously related to the input. They go a long way beyond simple audio transformations and are really synthesis techniques that derive only a small element, typically an excitation signal, from the live performance audio signal. This results in musical output that has far less correlation with the input signal. It is much more difficult for the performer and the audience to perceive the connection between the performance gestures and the output of the agent ensemble, especially in cases where an agent can continue to produce sound for several seconds after the input signal has decayed to silence. Examples of this behaviour in the later pieces include the various agents that drive plucked string

physical models or FM synthesizers, which may exhibit very long decay times. While these agents are completely dependent on the input material to determine their pitch and rhythmic attributes, not to mention the energy required to produce their sound, the connection is not necessarily easily perceivable to either the player or the audience. (It is interesting to note that I have had several conversations after performances with listeners that have had trouble believing that all of the sounds in the work originate from a flute.)

The artificial agent performers that have been equipped with synthesizers have been delegated the responsibility to create additional musical lines. The melodic shape and rhythm of these lines are determined by a combination of their internal algorithms and what they have been exposed to in their recent past. Technically, this system works, but the more I perform with it, the more I think that there are several problems, particularly in the context of working with a fixed score. Perhaps the biggest issue with *Pandan Musings*, in my opinion, lies in the original idea of trying to drive the interactive composing system via a conventionally notated score. I now feel that this situation can be too constraining. It restricts the control possibilities and severely reduces the performer's ability to explore options that might emerge through agent interactions. In some of the pieces, I realise that the interactive composing system can easily degenerate into a sub-optimal auto-accompaniment generator. In addition, the written notation constrains the form of the piece irrespective of the music output of the agent ensemble. For example, some of the systems tend to produce a regular grid-like beat, without much textural variation. In the free improvisation context, this is not so much of a problem because the performer has the opportunity to react to this situation and can change the instrument part,

which will force the system to react and to change its behaviour. In this context, the limitations of the system can be viewed as a positive feature, and the musical output of the agent system can be harnessed by the performer to evolve an appropriate form.

It is clear to me that the attempt to try to combine an interactive composing system with a fixed score, and hence, a fixed macro-level form, was not as effective as I thought it would be. I have always been aware that this approach contradicts the spirit of interactive composing, but I thought that it would produce something interesting. Initially, the notated part and the agent system for each movement were developed together, but now I consider the interactive composing systems to be the legacy of this compositional work. The collection of multi-agent systems has become a separate entity of their own, which I now refer to as *Pandan Meditations*. Over the last few years, these systems have proved themselves in a variety of performance situations, ranging from solo improvised concert performances, installations and as part of a live accompaniment for dance. Customised versions have been integrated into performances of *Sonic Escapade* and *Momentary Diversions*.

Another original conceptual element of *Pandan Musings* has been abandoned in recent performances with the interactive composing systems. Initially, the multi-agent systems were optimised for the flute to be played with a conventional flute timbre. Extended performance techniques were avoided in an attempt to provide a clear sonic contrast to the timbre of the electronic parts. I think this approach was effective in the two sound-centric pieces contained in *Pandan Musings*, but in the context of the other interactive composing systems that are more note-centric, I now feel that restricting the timbral range has

become a drawback. In these three systems (the *D Lydian*, the *E Neapolitan Major* and the *E Flat Lydian Minor* systems), the agent ensemble creates sounds that are reminiscent of plucked strings and mallet instruments, and the overall effect can easily become an imitation of a conventional ensemble. Because the additional textural layers that are created are influenced by the instrumental part, many interesting opportunities that might emerge from the agent ensemble can be hard to be framed as foreground material when a conventional flute timbre is maintained. Extended techniques, such as whistle tones, breath noise, multiphonics and harmonics enable many more opportunities for textural interplay between the instrument and the electronics. The woodwind sound can more effectively become part of the background, allowing the possibility of the agent generated parts to become the focus of the listeners' attention.

Superficially, these goal-driven works resemble the many compositions created for a live performer and electronics in recent years, but few pieces share the same intent. *Pandan Musings/Meditations* are inherently interactive, so are in a different category from all works that feature a fixed electronic part. And, because they are built around the interactions of only one human performer, they are in a different category from the many of the interactive works that involve real-time manipulation of an acoustic instrument during a performance by someone else. In that type of work, the composer usually performs any sonic transformations manually. This common approach shifts the burden of decision making from a collection of software agents interacting spontaneously with an arbitrary live signal, which is the case with *Pandan Musings/Meditations*, to a human agent. A human can see, can plan with the instrumentalist, can hear the

acoustic and electronic sound in the room, and take advantage of his or her sophisticated human brain to enhance the performance.

This leaves *Pandan Musings/Meditations* in the category of autonomous interactive systems. Many of these types of performance system create a sense of cohesion by constructing the computer part by extracting elements directly from a live instrumental/vocal performance and incorporating these into the sonic output. Systems deployed in many pieces, such as *The Air Inside Our Heads* by Charles Nichols, sample, and often loop, short fragments of the live performance enhanced by some sort of signal processing. Another common approach is to extract the pitch/amplitude parameters of a live signal and to transform this data to make musical material via a compositional algorithm. Prominent examples of this approach from the literature include the various implementations of the *Voyager* system developed, over many years by George Lewis, and Robert Rowe's *Cypher* system. The *Pandan Musings/Meditations* systems fall into this general category, but they are far more autonomous in terms of how they produce their material. Data derived from analysis of the live input is not directly transformed to generate musical material; instead, it determines the current overall goal state of the system. The individual software agents autonomously apply their embedded strategies to produce sonic material that satisfies the requirements specified by the current goal state. This is a much less restrictive approach, but the inherent cohesion obtained by producing material that is directly derived from the input signal is lost.

Chapter 4: Concluding Summary and Future Work

The compositions in this portfolio are instances of a compositional framework where formal generative processes are embedded inside agents that exhibit a degree of autonomy. The generative processes are mechanisms that enable each agent to navigate through its space of compositional possibilities. The compositions inhabit a number of different musical contexts or environments and demonstrate the viability of this approach to composition. Although this portfolio is primarily concerned with producing electro-acoustic music and there is an emphasis on agents that are designed for interactive real-time operation, the composition framework was also applied to create chamber music. In this portfolio, there was no attempt to create completely autonomous systems, rather, a collection of agents work in tandem with a human being, who provides guidance and/or decision making in some way. Taking the human composer out of the system, while an interesting idea, would have shifted the focus of this project more towards artificial intelligence issues instead of the utility of multi-agent systems as practical compositional tools. In some sense, the compositions in this portfolio try to carve a middle path between the idea of compositions as structures built out of sonic materials, as espoused by Varèse (Varèse, 1966) and others, and the idea by Cage, of compositions being containers that frame something that is already there (Cage, 2011).

By encapsulating various formal processes inside an agent's decision-making mechanism, where they serve as navigation aids that guide individual

agents as they traverse their state space, it is easy to build multi-layered musical textures with arbitrary combinations of generative processes. Consequently, the conceptual approach of building a musical texture out of layers, each designed to fulfil a complimentary musical function, can integrate disparate formal generative techniques such as those that were used by various twentieth-century composers.

Time, Form and Improvisation

One aspect of an agent-based approach that became very clear to me during this project, and is a common element found in other people's musical work related to agents, is how well agents are suited to working in real-time improvisational contexts. By giving a degree of autonomy to an agent, allowing it create its own contribution to the musical output, the composer gives up some control, which results in a greater opportunity for a human performer to respond to the new possibilities created by the agents.

Several pieces of the portfolio, such as *Along the Corridor* and *In Transit*, instantiate a model derived from the multi-agent compositional framework, which I call the Painting Approach, which includes a partial real-time component. The basic material was created in real-time via improvisation, and these improvisations were later assembled, out of real-time, to create these works. The low-level agent creating the musical material was realised by a human improviser, who guided a software, or hardware, performance system that implemented the formal generative processes. The resulting audio recordings were then edited and assembled to create the fixed form of the piece. This workflow is analogous to that of *Musique Concrète* and its descendants. Applying

formalisms that can consciously affect specific musical functions, or note choice strategies that have proven successful elsewhere, increases the likelihood of creating a complimentary collection of musical material that can be arbitrarily combined. This also reduces the complexity of the decision-making process during the assembly phase, because selections will be made between recordings with clearly demarcated attributes. In *Along the Corridor* and *In Transit*, the assembly process was very improvisational and intuitive, drawing on a wide palate of distinctive recordings. The form of these works grew out of the attributes of the sounds themselves, and the compositional decisions were made by a process of mixing and matching different parts, making choices that produced what I considered to be a convincing sense of continuity, and engendered a sense of motion through each work.

Instead of restricting the real-time component to one phase of the compositional process, other works are completely focused on real-time performance. These works are created by constructing a community of agents, specifying their internal structures and relationships, and then creating an instance of the work completely in real-time, potentially in a concert situation. This type of work may have a preconceived structure which needs to be realised by guiding the agents in some way, or may have an open form that is completely improvised during a performance.

In this portfolio, preconceived structures are articulated with some form of notation. Traditional notation poses some challenges in a multi-agent situation because it is difficult to anticipate the contributions of the agents on the final sound that the audience will hear, without unduly compromising the agents' autonomy. Creating fixed forms was approached in several ways in this project.

One approach, applied in the *Flute Trio in the Diminished Scale*, simply eliminated uncertainty altogether. Because all of the parts are strictly notated, I manually implemented the various formal processes out of real-time, executing the Painting Approach to create a traditional score. The formal processes that were chosen to generate material for this piece are drawn from the conventional catalogue of motivic transformation processes, and result in a traditional, note-centric (as opposed to a sound-centric) acoustic work.

The Pruning model, where a group of agents simultaneously create their potential contribution to the musical texture while a high-level agent determines which layers are heard by the listener, proved to be very effective for creating fixed forms in real-time. The high-level structure of *Sky Castles* was articulated with this approach, and, in several of the other works, the pruning model was combined with other approaches. *Sky Castles* articulates its fixed structure through a combination of conventional notation and a fixed schedule specifying which agents are heard as the piece progresses. A small element of constrained improvisation is present, which adds a significant element to the overall sound of the piece because its presence fundamentally alters the timbre of the piece. Although the performer is presented with a conventionally notated fixed series of pitches, he or she is required to improvise its timing (based on the length of a breath) and, most significantly, to improvise a simultaneous sequence of pitches that the player must hum whilst playing the notated pitches. The combination of humming and the natural flute sound results in a buzzy timbre, which establishes the characteristic timbre of the work. Concurrently, a collection of agents transforms the flute performance into a mass of sound, as more agents are progressively added to the mix. This combination of factors produces an

easily identified result that is predictable within set boundaries, while still providing considerable scope for variation. The performer is encouraged to interact with the sound produced by the agents via the improvised humming part.

A variation of the Pruning approach is found in *Sonic Escapade*, where human performer agents themselves determine if the audience will hear them. They determine when and what they will play and decide their own dynamic level. A certain amount of pruning occurs before the performance when the instrumentation is decided. This occurs because, while each performer plays a variety of instruments, there are many potential instrumental configurations of the ensemble instrumentation that can never occur due to constraints placed upon them by the composer.

Pandan Musings creates a fixed form for each for each piece which is articulated in a fixed score and is performed in a specific multi-agent environment. Some of these environments consist largely of agent controlled effects processors whose role it is to cast the instrumental performance in a distinctive sound world. More complex environments feature agents that can transform the performed instrumental part into completely autonomous textural layers which perform different musical functions. In all cases, agents decide on the details of each individual layer in real-time and determine which textural layers will be heard at any given time, basing those decisions on the output of other agents that provide an analysis of the live performance. The performer needs to spontaneously react and shape the ensemble output by shaping the live instrumental signal.

The remaining works in the portfolio have a more open structure and are either constrained, collective improvisation environments, or may even be instantiated as installations, allowing members of the general public to interact with artificial agent performers. Each work gains its distinctive identity from the choice of formalisms embedded inside the agents that create the musical material, the means of interaction between agents and human performers, and the instrumental resources employed.

Conclusion

The contents of this portfolio demonstrate that three basic composition models that were derived from the multi-agent framework can be applied in a variety of different musical contexts. The Painting, Pruning and Goal-Led approaches all rely on independent collections of formal processes to generate the various layers of the musical texture that the work. It is clear that the choice of the combination of formal procedures that are embedded in the various agents is crucial. The formal processes chosen have a large influence on the balance of predictability and uncertainty, as well as the ability to achieve a sense of continuity and transparency. When creating these works, I found that it was important that the musical output derived from the formalism was distinctive and could be controlled sufficiently to fulfil whatever musical function I felt was required. After experimenting with many different algorithms, of varying sophistication, the most frequently implemented algorithms were simple Gaussian distribution generators. One reason for this was that its two controls, the mean and the standard deviation, facilitate an intuitive way to effectively achieve a good predictability/uncertainty balance, which yielded the creation of

textural layers that could fulfil specific musical functions. The other reason for this choice was that the mean and the standard deviation can be quickly calculated, which makes it easy to implement simple error correction algorithms which enable simple machine learning to take place. This means that agents containing Gaussian agent functions can easily adjust themselves to meet externally specified goals. The other common algorithm deployed in the portfolio was piecewise linear functions, which are made up of lines connecting points, due to their simplicity of implementation and their intuitive nature.

It was also very clear that the goal of making high-level compositional decisions by software agents is very difficult to do well, and is an area that needs much more research and computing power. This is an area I will be working on in the future. The creation of sophisticated high-level agents is beyond the scope of this project, so all high-level decisions are left to either a human, or a finite state machine. In most of the compositions in this portfolio, the focus was on the construction of suitable the performer agents, to meet the needs of the piece. These agents really are agents in the traditional sense, that of “A person that acts on behalf of another” according to the online Oxford dictionary (Oxford Dictionary, n.d.), because, although they are delegated to perform particular tasks, in the end, it is the composer, and any human performer present, that shapes the experience for the audience. The compositions that make the most use of artificial agents to determine high-level musical goals are *Pandan Musings* and *Pandan Meditations*. These are real-time pieces for a solo instrumentalist, and the high-level musical tasks are delegated to agents so that the player can focus on listening and interacting with the ensemble output.

Involving agents in composing music inherently means that there is an act of delegation of some aspects of the compositional process to someone, or something, else. The advantages and disadvantages of using an agent based compositional approach both stem from this aspect of delegation. The prime advantage I noticed when creating the works in the portfolio, was that multi-agent designs imposed a modular approach, which allowed me to focus my attention at different levels of detail, at different stages of the development of each piece, whilst allowing me to go and easily revise anything along the way. The multi-agent framework provided the conceptual scaffolding that allowed me to arbitrarily work on any aspect of the piece to whatever degree of refinement that I thought was appropriate at that time. I perceive that this modularity helped me to create an overall sense of narrative, flow, journey or high-level progress as a work unfolds, which I feel is important in these works. The modularity provided a level of abstraction, turning the results of the agent's actions into structures that exist at a higher level than individual notes or sounds. The tactic that the agents were producing musical material that fulfilled distinct musical functions made it easier to get a clear idea of how the piece could be shaped.

The principal disadvantage of the approach was also the result of delegation. Delegating low-level surface details of the music to agent performers meant that I had to accept whatever output they delivered when they constructed the surface details of the piece. While this is not necessarily a bad thing, I considered it important to maximise the probability of achieving something I felt was suitable. Moderating this risk was one of the rationales for selecting formal processes capable of reproducing functional behaviour, such as

creating pads, drones etc, because these have proven to be effective tools for constructing complementary musical textural layers in other contexts.

To conclude, the works in this portfolio verify the validity of the three assumptions stated in the research question and demonstrate some of the ways that an agent based compositional framework can facilitate the creation of musical works that exhibit coherence and a distinctive identity. All of the works were created with an approach that segregated high-level structural concerns from low-level surface details. The works demonstrate a variety of approaches of imbedding formal processes inside low-level agents in order to generate the elements that will make up the musical surface. A number of formal processes were applied which contain some mechanism to adjust their output, thus altering the musical material produced by the agent. In some agent designs, this mechanism was directly accessible to higher-level agents, human or otherwise. In other designs, this mechanism specified particular goal states that indicate the specific musical behaviour that the high-level agent requests, and the low-level agent autonomously adjusts its internal state so as to align its musical output with this behaviour.

The agent approach has demonstrated its suitability for real-time applications and shares many characteristics of improvisation. Most of the compositions in this portfolio feature at least some real-time improvisatory elements. In each case, the question regarding how to balance the improvisatory element of surprise, while still producing a distinctly identifiable work, had to be resolved and was addressed in different ways. The works produced with the Painting approach simply restricted the improvisatory elements to an initial exploratory phase, where all the low-level material was created. The one

exception to this was a deliberate attempt to transfer the agent approach into a conventionally notated, fixed form via a process of the composer deliberately acting out the roles of an imaginary ensemble of agents. The fixed form works created with the other two models are realised by treating the multi-agent systems as instruments and articulating some performance instructions in some sort of notation, to constrain the performances. I think that each work in this portfolio succeeds in exhibiting a distinctive quality. At the musical surface level, what differentiated each work were the choices of formalisms and timbres, as well as inter-agent interaction. At the higher level, the modes of interaction that are available to mould the large-scale structure play a large role in the works' distinctiveness. Generally, the identity of a work emerges from the interactions that take place within the human and artificial agents as they explore their space of possibilities, and this a distinctive benefit of composing with a bottom up, multi-agent based compositional framework.

Section 5: References

Bibliography

- Adorno, T. W. (2007). *Philosophy of Modern Music*. (A. G. Mitchell, & W. V. Blomster, Trans.) London: Continuum.
- Auner, J. (2013). *Music in the Twentieth and Twenty-first Centuries*. New York, NY: W.W.Norton & Company.
- Bailey, D. (1993). *Improvisation: its nature and practice in music*. London: British Library National Sound Archive.
- Ball, P. (2010). *The music instinct: how music works and why we can't do without it*. London: Vintage Books.
- Barton, T. (2016). <http://toddbarton.com/>.
- Blackburn, M. (2011). The Visual Sound-Shapes of Spectromorphology: an illustrative guide to composition. *Organised Sound*, 16(01), 5-13 .
- Cage, J. (2011). *Silence: lectures and writings*. Middletown, CT: Wesleyan University Press.
- Chadabe, J. (1984). Interactive Composing: An Overview. *Computer Music Journal*, 8(1), 22-27.
- Chadabe, J. (1997). *Electric Sound: The Past and Promise of Electronic Music*. Upper Saddle River, NJ: Prentice Hall.
- Chadabe, J., & Meyers, R. (1978). An introduction to the PLAY program. *Computer Music Journal*, 2(1), 12-18.
- Coker, J. (2010). *Improvising jazz*. New York, NY: Simon and Schuster.
- Collins, N., & d'Escriván, J. (2007). *The Cambridge Companion to Electronic Music*. Cambridge: Cambridge University Press.
- Cook, P. R. (2002). *Real sound synthesis for interactive applications*. Wellesley, MA: A K Peters.
- Csikszentmihalyi, M. (1990). *Flow : The psychology of optimal experience*. New York: HarperCollins Publishers.
- Davis, M. (1990). *Miles*. London: Picador.
- Devine, R. (2016). <https://soundcloud.com/richarddevine>.
- Dick, R. (1975). *The other flute: a performance manual of contemporary techniques*. St. Louis, MO: Multiple Breath Music Company.
- Dodge, C., & Jerse, T. A. (1997). *Computer Music Synthesis, Composition, and Performance* (2 ed.). New York, NY: Schirmer Books.
- Dorin, A. (2001). Generative Processes and the Electronic Arts. *Organised Sound*, 6(1), 47-53.
- Dwyer, T. (1971). *Composing With Tape Recorders: musique concrete for beginners*. London: Oxford University Press.

- Eigenfeldt, A., Bown, O., Pasquier, P., & Martin, A. (2013). Towards a Taxonomy of Musical Metacreation: Reflections on the First Musical Metacreation Weekend. *Musical Metacreation: Papers from the 2013 AIIDE Workshop* (pp. 40-47). Palo Alto: Association for the Advancement of Artificial Intelligence.
- Emmerson, S. (1986). *The language of electroacoustic music*. Houndmills, Hampshire: Macmillan Press.
- Emmerson, S. (2013). *Living electronic music*. Farnham: Ashgate Publishing Limited.
- Eno, B. (2004). Generating and organizing variety in the Arts. In C. Cox, & D. Warner, *Audio culture: Readings in modern music* (pp. 226-234).
- Farnell, A. (2010). *Designing sound*. Cambridge, MA: The MIT Press.
- Feldman, M. (2004). *Give My Regards To Eighth Stree*. (B. H. Friedman, Ed.) Exact Change.
- Griffiths, P. (2011). *Modern music and after* (3rd ed.). New York, NY: Oxford University Press.
- Grout, D., & Palisca, C.V. . (1988). *A history of Western music*. (4 ed.). New York, NY: W.W.Norton & Company.
- Hiller, L., & Isaacson, L.M. . (1979). *Experimental Music; Composition with an electronic computer*. New York, NY: Mc Graw-Hill.
- Illinois Distributed Museum. (n.d.). *Sal-Mar Construction*. Retrieved June 2015, from http://distributedmuseum.blogspot.sg/p/sal-mar-construction_1.html
- Jones, K. (1981). Compositional applications of stochastic processes. *Computer Music Journal*, 45-61.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York, NY: Farrar, Straus and Giroux.
- Kostka, S. (2011). *Materials and techniques of post-tonal music* (4 ed.). Boston, MA: Pearson.
- Kreidler, J. (2013). *Loadbang: Programming Electronic Music in Pd*. Hofheim: Wolke Verlag.
- LaMothe, A. (1999). *Tricks of the Windows game programming gurus*. Indianapolis, IN: SAMS.
- Landy, L. (2007). *Understanding the art of sound organization*. Cambridge, MA: The MIT Press.
- Levitin, D. J. (2011). *This is your brain on music: Understanding a human obsession*. New York, NY: Dutton Group.
- Lochhead, J. a. (2013). *Postmodern music/postmodern thought*. New York, NY: Routledge.
- Lorrain, D. (1980). A panoply of stochastic 'cannons'. *Computer Music Journal*, 53-81.
- Loy, G. (1989). Composing with computers: A survey of some compositional formalisms and music programming languages. In M. V. Mathews, & J. R. Pierce, *Current Directions in Computer Music Research* (pp. 291-396). Cambridge, MA: The MIT Press.
- Loy, G. (2006). *Musimathics: The mathematical foundations of music* (Vol. 1). Cambridge, MA: The MIT Press.
- Loy, G. (2007). *Musimathics: The Mathematical Foundations of Music* (Vol. 2). Cambridge, MA: The MIT Press.

- Lucier, A. (2012). *Music 109: notes on experimental music*. Middletown , CN: Wesleyan University Press.
- Manning, P. (2004). *Electronic and computer music*. New York, NY: Oxford University Press.
- Miranda, E. (2001). *Composing Music with Computers*. Oxford: Focal Press.
- Moore, F. (1990). *Elements of computer music*. Englewood Cliffs, NJ: Prentice-Hall.
- Nierhaus, G. (2009). *Algorithmic composition: paradigms of automated music generation*. New York, NY: Springer Wein.
- Oxford Dictionary. (n.d.). *Oxford Dictionary*. Retrieved June 2015, from Oxford Dictionary: <http://www.oxforddictionaries.com/definition/english/agent>
- Pinker, S. (1997). *How the Mind Works*. New York, NY: W. W. Norton & Company.
- Redgum (Composer). (1984). *Frontline*. [Redgum, Performer] Sydney, NSW, Australia: Epic.
- Redgum (Composer). (1986). *Midnight Sun*. [Redgum, Performer] Sydney, NSW, Australia: Epic.
- Roads, C. (1996). *The Computer Music Tutorial*. Cambridge, MA: The MIT Press.
- Roads, C. (2004). *Microsound*. Cambridge, MA: The MIT Press.
- Roads, C. (2015). *Composing Electronic Music*. Oxford, UK: Oxford University Press.
- Ross, A. (2007). *The rest is noise: Listening to the twentieth century*. New York, NY: Picador.
- Rowe, R. (1992). *Interactive music systems: machine listening and composing*. Cambridge, MA: The MIT Press.
- Rowe, R. (2001). *Machine Musicianship*. Cambridge, MA: The MIT Press.
- Russel, S. J., & Norvig , P. (2003). *Artificial Intelligence A modern approach*. Englewood Cliffs, NJ: Prentice Hall.
- Sabatella, M. (2000). *A jazz improvisation primer*. Retrieved June 2015, from The Outside Shore: <http://www.outsideshore.com/music/a-jazz-improvisation-primer/>
- Schaeffer, P. (2012). *In search of a concrete music*. (C. North, & J. Dack, Trans.) Berkeley, CA: University of California Press.
- Schafer, R. M. (1993). *The soundscape: Our sonic environment and the tuning of the world*. Rochester, VT: Destiny Books.
- Schoenberg, A. (1970). *Fundamentals of musical composition*. (G. Strang, & Stein, L., Eds.) London: Faber and Faber.
- Scruton, R. (2009). *Understanding Music*. London: Continuum.
- Smalley, D. (1997). Spectromorphology: explaining sound-shapes. *Organised sound*, 2(2), 107-126.
- Spicer, M. (2005). SPAA: An agent based interactive composition. *Proceedings of the International Computer Music Conference*. Barcelona: ICMA.
- Spicer, M. (2006). Agents in ChuckK A Timely Programming Experience. *Proceedings of the International Computer Music Conference* . Miami: ICMA.
- Spicer, M. J. (2003). A Real-Time Agent Based Interactive Music Performance System. Singapore: National University Of Singapore.
- Spicer, M., Tan, B., & Tan, C. L. (2003). A Learning Agent Based Interactive Performance System. *Proceedings of the International Computer Music Conference*.
- Stockhausen, K. (1989). *Stockhausen on music*. (R. Maconie, Ed.) London: Marion Boyars.

- Strange, A. (1983). *Electronic music: systems, techniques, and controls* (2 ed.). Dubuque, IA: Wm. C. Brown Company Publishers.
- Varèse, E. (1966). The liberation of sound. (C. Wen-Chung, Ed.) *Perspectives of new music*, 11-19.
- Wang, G., & Cook, P.R. . (2003). ChuckK: A concurrent, on-the-fly audio programming language. *Proceedings of International Computer Music Conference* (pp. 219-226). ICMA.
- Wells, T. H. (1981). *The technique of electronic music*. New York, NY: Schirmer Books.
- Whalley, I. (2009). Software Agents in Music and Sound Art Research/Creative Work : current state and a possible direction. *Organised Sound*, 14, 156-167.
- Whitman, K. F. (2016). <http://www.keithfullertonwhitman.com/>.
- Winkler, T. (2001). *Composing interactive music: techniques and ideas using Max*. Cambridge, MA: The MIT Press.
- Winsor, P., & DeLisa, G. (1991). *Computer music in C*. Blue Ridge Summit, PA: Windcrest.
- Xenakis, I. (1992). *Formalized music: thought and mathematics in composition* (Revised ed.). Hillsdale, NY: Pendragon Press.
- Zappa, F., & Occhiogrosso, P. (1989). *The Real Frank Zappa Book*. New York: Poseidon Press.
- Zorn, J. (2000). *Arcana Musicians on Music*. (J. Zorn, Ed.) New York, NY: Granary Books/Hips Road.

Appendix

Appendix A

Portfolio Contents

Track	Title	Dur	Performer
1	<i>Along The Corridor</i>	5:01	Michael Spicer
2	<i>InTransit</i>	4:37	Michael Spicer
3	<i>A Painting In Sound</i>	7:18	Michael Spicer
4	<i>Flute Trio In The Diminished Scale Mov 1</i>	4:21	Michael Spicer
5	<i>Flute Trio In The Diminished Scale Mov 2</i>	1:00	Michael Spicer
6	<i>Flute Trio In The Diminished Scale Mov 3</i>	3:31	Michael Spicer
7	<i>Sky Castles</i>	2:21	Michael Spicer
8	<i>Momentary Diversions</i>	12:10	Michael Spicer
9	<i>Sonic Escapade</i>	6:27	Michael Spicer, Deborah Tan, Christopher Clark, Irfan Rais, Ng Zheng Jie, Xaviour Lee
10	<i>Pandan Musing I</i>	4:18	Michael Spicer
11	<i>Pandan Musings II</i>	3:59	Michael Spicer
12	<i>Pandan Musings III</i>	3:56	Michael Spicer
15	<i>Pandan Meditations E Neo Maj</i>	5:52	Dr. Timothy O'Dwyer

Appendix B

Data Archive Directory Structure

Goal-Led Approach	
Pandan Musings PD Projects	
	Pandan Musings 1 pd C Lydian Dominant (Pd project)
	Pandan Musings 2 pd D Oriental (Pd project)
	Pandan Musings 3 pd D Phrygian (Pd project)
	Pandan Musings 4 pd E Neo Maj (Pd project)
	Pandan Musings 5 pd E flat Lydian Min (Pd project)
Pandan Musings scores	
	Pandan Musings and Meditations Score.pdf
Painting Approach	
A Painting in Sound	
	A Painting in Sound Program Note
Along the Corridor	
	Along The Corridor Description
Flute Trio In The Diminished Scale	
	Flute Trio In The Diminished Scale Full Score.pdf
In Transit	
	In Transit Program Notes
Pruning Approach	
Momentary Diversions	
	Momentary Diversions Rules
Sky Castles	
	Sky Castles (Logic project)
	Sky Castles - Full Score.pdf
	Sky Castles 60x60 Program Notes.docx
Sonic Escapade	
	Sonic Escapade - Rules of Engagement

Appendix C

Australasian Computer Music Conference Paper

Spicer, M. (2010). Composing “In Transit”. Proceedings of the Australasian Computer Music Conference 2010. Canberra: ACMA.

COMPOSING “IN TRANSIT”

Michael Spicer

Singapore Polytechnic/
Monash University Conservatorium

ABSTRACT

“In Transit” is a fixed form electro acoustic piece. It combines of a number of short improvisations performed on an instrument created by the composer using the ChuckK programming environment. It enables the performer to create sequences of notes that display the characteristics of various probability distributions. The performer has macro level control of the output by adjusting parameters that control the algorithms, in real time. The improvisations are imported into a Digital Audio Workstation where they are edited and processed with various signal processors before being assembled to create the finished piece. The assembly makes use of a “bottom up” approach. There is no clear idea as to how the piece will be, the final form emerges from the way the composer combines the improvisations. The composer works largely intuitively, making use of high-level concepts such as progression, continuity and contrast as a framework to help make compositional decisions.

1. INTRODUCTION

“In Transit” is a fixed form stereo playback electro acoustic piece. It was created using a combination of ChuckK, and Logic Pro. All sounds are produced with physical modeling synthesizers, but the timbres have been significantly transformed by a variety of signal processing techniques. The piece is predominantly organized around timbre and texture, and (mostly) does not exhibit a clearly articulated rhythmic pulse. A characteristic feature of the piece is the use of diffuse clouds of notes that fade in and out, which was the only element of the piece that was decided upon when the composition of the piece was begun. The overall form grew out of the composition process, but effort was made to make sure the piece had a sense of progression, as if the listener was on a leisurely journey, hence the name “In Transit”.

The impetus for creating this piece came from a demonstration of the “Probability Jammer” in a music class. Each student was running “Probability Jammer” on his or her laptop and they all played with it through the laptop speakers. The class was divided up into groups, and students were directed to adjust the parameters in particular ways, “conducted” by the composer. This

created an interesting cloud of notes in the room that changed as the students performed. This effect was the only preconceived characteristic that the composer set out to achieve in the final form of this piece.

2. COMPOSITIONAL APPROACH

“In Transit” is part of a collection of pieces that are created using a “bottom up” compositional methodology that attempts to apply some of the approaches the composer has successfully used in popular music to an electro acoustic context. In particular, there is the intention to include an element of improvised performance as the way of creating musical material, which is then further refined, out of real time. The piece is constructed in three distinct phases.

1. Create Real Time Improvisations
2. Apply Processing and Editing
3. Assembly/Mix

To use an analogy of a child building a plastic model kit (such as a plane/car/ship), the first phase involves creating the basic components (done by the model manufacturer), the second phase would be painting the parts, and the third phase, assembly. The big difference between the model building process and the compositional process is that the final form of that model is known from the beginning, while in this composition, the final form emerges in the assembly process, and is not apparent until the end. In that respect, it is like taking a collection of Lego bricks, and, without any preconceptions about what to make, putting pieces together randomly for a while, then assessing the situation to see if the assembly suggests any particular thing, then modifying this to refine it.

2.1. Phase 1 – Real Time Improvisations

In this piece, the basic building blocks were fourteen short solo (unaccompanied) improvisations performed on an instrument that the composer created with the ChuckK programming language, called “Probability Jammer”. This instrument makes use of two probability distribution generators to control the pitch and duration of notes that are produced using one of the STK physical modelling

synthesizers [1]. All of the STK instruments supplied with ChuckK are available for use with the instrument. “Probability Jammer” has a very simple user interface, shown in figure 1, implemented using the MAUI widgets, part of the MiniAudicle ChuckK development environment. Each onscreen fader can also be controlled via a MIDI fader box, which effectively allows all parameters to be altered simultaneously.




Figure 1. “Probability Jammer” user interface.

An example of the output of one of the probability generator functions (the Gaussian) with different settings is shown in figures 2, 3 and 4. The left hand window is a plot of the output produced when the generator is run 100,000 times without changing the parameters. The right hand window displays the corresponding (normalized) probability distribution function, and the relationship between changes in the parameter settings and the output produced is quite apparent.




Figure 2. Output of the Gaussian generator, with its probability distribution for 100,000 samples. The mean is 0.5 and deviation is 0.2.




Figure 3. Output of the Gaussian generator, with its probability distribution for 100,000 samples. The mean is still 0.5, but the deviation is reduced, producing a series of outputs that are closer to the mean.




Figure 4. Output of the Gaussian generator, with its probability distribution for 100,000 samples. The mean has been raised, while the deviation has been kept the same as figure 3. Now the outputs are clustered around the new (higher) mean.

In order to perform with “Probability Jammer”, the composer needs to specify the type of probability distributions to be applied to pitch and duration, select the synthesizer (all specified using the ChuckK language), and initialise their controlling parameters. Various probability distributions, taken from the standard literature [2,3,5], can be applied to the pitch and duration, such as Gaussian, Linear, Exponential and Triangle probability distributions (Settings for using Gaussian are shown in figure 1). One other improvisation was made with another ChuckK instrument that is a variation of “Probability Jammer”. This instrument is optimized for making multilayered

drone like textures. It makes use of a Gaussian probability distribution to choose pitches that are at harmonic frequencies of a note specified by the user, via a keyboard, using a number of STK synthesizers.

Different strategies were employed by the composer to create the different improvisations that form the raw material for this piece. In a few improvisations, there were no changes in the controls during the recording phase. Once the instrument settings were producing a distinctive output that the composer found interesting, the audio was recorded for around thirty seconds. One example of this approach is the improvisation used at the very beginning of the piece. It makes use of a Gaussian distribution to control the pitch, with the mean pitch set quite low and a medium deviation around this pitch. The duration is set so as to produce a constant, fast, note rate. As these parameters do not change, all of the notes generated exhibit these characteristics. Most of the improvisations, such as the solo melodic line in the middle of the piece, were created by initializing the system to a configuration that the composer felt was a good starting point, and then performing the improvisation by adjusting the probability distribution parameters while the part was being recorded. In this situation, the composer was usually thinking in terms of melodic line, aiming at creating a recording with a degree of variation in the note rate and the pitch. There was an attempt to create a distinct pitch contour, with some rhythmic variation.

2.2. Phase 2 – Processing and Editing

The second phase of the composition process involves “dressing up” the timbre of the raw improvisations (which sound very raw indeed) to create a basic palate with which to assemble the piece. The recordings of the ChuckK improvisations were imported into Logic Pro8 and signal processing was applied to each part so as to accentuate an aspect of its character that appealed to the composer. The aim was to create a collection of distinct musical “gestures” that would be recognizable, even after undergoing various transformations, as these will be the elements that provide coherence and can help a listener make sense of the piece.

All of the standard modern signals processing options, such as equalisation, reverb, compression etc, were utilized as seemed appropriate, during this phase of production. It is worth noting that a large part of the characteristic surface “sound” of the piece came from extensive use of the “EnVerb” plug in (a combination of a reverb and an envelope shaper), the Rotary Speaker simulator, as well as various types of delay and distortion. Offline Pitch Shifting was applied to some of the audio files, to create some of the low pitched parts. Additionally, in order to enhance the “sound cloud” effect that was apparent in

some of the improvisations, those recordings were duplicated, cut into sections and reordered, and layered with some timing offsets. This resulted in a more diffuse cloud of notes with the same overall pitch characteristics.

2.3. Phase 3 – Assembly and Mix

Once the palate of parts was in place, the piece was assembled. This took place over a period of one month. It was done using a trial and error process involving experimenting with different combinations of parts so as to create a distinct character that appealed to the composer. It was important to the composer that the final form of the piece exhibited a certain clarity/transparency in its overall sound. To help achieve this, experiments creating different textural layers, with various degrees of timbral contrast, were undertaken. After some time, nine distinct textural layers and four distinct sections emerged:

1. High and low organ like flourishes over the multilayer drone.
2. Swirling “cloud like” textures.
3. A melodic section making use of the “bowed” STK instrument.
4. A noise/percussion section.

In this phase of the construction, the texture and timbre of large groups of notes were the main concern, rather than the details of the individual notes. Pitch is only organized in terms of predominant register (High, Medium, Low).

The primary concern was to achieve a sense of cohesion and progression through the duration of the piece, with a convincing musical continuity (with appropriate balance between “flow” and “break”). In its final form, the piece starts quite lively and progress towards an (anti) climax, where it virtually stops, and then returns to some state similar to the beginning (a variation of a traditional ternary form). The final form is clearly seen in the screen shot of the arrange page of Logic Pro, shown in figure 5.




Figure 5. Logic Arrange page, indicating the overall structure of the piece.

3. HISTORICAL CONTEXT

“In Transit” is part of a project that attempts to apply, in an electro acoustic context, some of the processes that the composer utilized working with rock/folk/jazz groups (and recording sessions for TV and film) in the 1980s. Central to this approach was that players made use of distinct strategies to come up with their own part, rather than playing a set “arrangement”. The arrangement emerged spontaneously, in real time. In that situation, the strategies were often “rules of thumb” for creating common textural elements. An example would be the various different approaches used to create the idiomatic rock/jazz/ballade/country bass lines. Approaches for these could be (amongst many others):

- rock - play the root of the chord in quavers.
- jazz - play a “walking bass” making sure to clearly emphasise the chord tones.
- ballade - play the root of the chord dotted crotchet-quaver pattern.
- country – play the root and fifth of the chord in minimums.

Similar idiomatic approaches were applied to the creation of “comping” patterns, pads, arpeggio parts and counter melodies etc. Often, this approach is moderated (often by a producer) either in rehearsal, while the parts are being recorded (especially in a multitrack recording situation) or after the performance was recorded. Mixing, signal processing and editing techniques, such as copying/cutting and pasting/splicing, have been utilized extensively to modify recorded improvised performances, since the mid 1960s. A well-known example is “Bitches Brew” by Miles Davis, where the raw recordings of the musicians were heavily edited to create the final record. (Ironically, for this project, “Bitches Brew” was influenced by the work of the Musique Concrete composers.)

The process used to create “In Transit” is an obvious descendent of the above process. The “rules of thumb” are replaced by the distinct improvisation strategies used to create performances with the “Probability Jammer” (directly inspired by Xenakis), and the finished composition emerges through choices made by the composer in the assembly process, making use of extensive studio production techniques.

4. CONCLUSION

In this piece, the composer has made use of a compositional process that attempts to create an electro acoustic piece that combines the elements the spontaneity of an improvised performance with the careful consideration of a pre-composed piece. The process provides a lot of opportunities to create interesting musical results through the serendipitous juxtaposition and

superposition of material. The approach also enables the possibility of musical decision making skills developed in other musical genres to be applied in an electro acoustic context, which may be considered, by some people, as a desirable attribute.

5. REFERENCES

- [1] Cook, P (2002), *Real Sound Synthesis for Interactive Applications*, A K Peters Ltd, Massachusetts.
- [2] Dodge, C. and Jerse, T.A. (1997). *Computer Music Synthesis, Composition, and Performance*. Schirmer Books, New York NY, U.S.A.
- [3] Lorrain, D. (1989). “A Panoply of Stochastic ‘Cannons’”. In Rhodes, C., editor *The Music Machine*, pages 351-379. MIT Press, Cambridge MA, U.S.A.
- [4] Wang, G and Cook, P.R. “ChucK: A Concurrent On-The-Fly Audio Programming Language” In *Proceedings of the International Computer Music Conference*. Singapore, 2003
- [5] Xenakis, I. (1971). *Formalized music; thought and mathematics in composition*. Indiana University Press, Bloomington, IA, U.S.A.

Appendix D

International Computer Music Conference Paper

Spicer, M. (2014). A multi-agent interactive composing system for creating expressive accompaniment. Proceedings of the International Computer Music Conference . Athens: ICMA.

A Multi-agent Interactive Composing System For Creating Expressive Accompaniment.

Michael Spicer

Singapore Polytechnic, Monash University

ABSTRACT

This paper describes the approach and an application that the author has adopted for creating real time performance systems whose musical output is created by the interactions of a human performer and a multi-agent system that acts as an ensemble of software "performers". The music produced typically consists of several distinct textural layers, where all the sounds produced are transformations of the sound made by the human performer. This type of system can be thought of as an "extended" instrument, where the performer effectively "plays" the ensemble. This approach has been used with notated compositions, improvisation performances and for creating installations. This paper focuses on a composition that utilises a notated score, and is concerned with how the score is interpreted in the context of the musical output of the agent ensemble. This system makes use of two broad categories of agent: performers and controllers. Performer agents transform the live sound in various ways, while controller agents work at a higher structural level. They specify goal states and determine which agents are currently heard. Each performer agent has a way of transforming the audio input, and has its own internal strategies for determining what it does. The complexity of the performer agents note choice strategies ranges from simple harmony generators, to algorithmic composition systems.

1. INTRODUCTION

"Pandan Musings" is a notated, five-movement piece for Flute and Computer. The computer is running a Pd [1] patch that implements a multi-agent based interactive composing system. See Whalley [2] for an overview of musical applications of multi-agent systems for music and Chadabe[3] for information about interactive composing. The Pd patch transforms the sound of the flute performance in a variety of ways, creating a multilayered musical texture. Each layer adds its own acoustic signature which compliments the other layers present. An important aspect of the piece is the expressive relationship between the human performer and the response created by the multi-agent system. Each movement has its own distinct flavour, which

is the result of the particular strategies implemented in the agents, the type of timbre the flute sound is transformed into, and the use of different musical scales.

2. COMPOSITIONAL APPROACH

2.1 Overview

This piece grew out of many years of improvising with multi-agent interactive composing systems. [4] [5] [6] When working in an improvisation context, the focus of the multi-agent system development was always on achieving a balance between the amount predictability and the amount surprise in the way the system responded. Achieving the right amount of unexpected behaviour in the multi-agent systems response was important in creating a sense of collaboration with the ensemble of software agents. "Pandan Musings" frames this element of uncertainty in a slightly different way from the improvisation context. A primary element of this piece is that the live performer is affected by the musical contribution made by the multi-agent system when interpreting the pre-composed score. The multi-agent system is set up during the composition phase so as to strike a good balance between consistency and variety from performance to performance, forcing a fresh interpretation of the score every time it is played.

Early versions of the piece made use of software synthesizers to realise the parts composed by the agent software, but it didn't create the intimate, expressive feedback loop that lent itself to expressive interpretation. It did not create the necessary feeling of control over the multi-agent systems output. To overcome this problem, the original synthesizers were replaced by a collection of signal processing systems that transform the sound of the live flute signal. This allowed the multi-agent system to control the pitch and duration of the sonic events in the textural layers, producing a variety different timbres, while still providing some expressive control linked to the dynamics, articulation and phrasing of the performer.

2.2 Extra-musical Aspects

The piece was inspired by the locale of an artificial reservoir in Singapore, that was created by flooding a coastal mangrove swamp in the 1970s. It forms an intersection between several different ecosystems. On the east side, there is still an area of mangrove, and an short estuary leading to the sea, the north side has high rise housing estates, and to the south and west there is a dense industrial estate contain-

ing a variety of businesses, including logistics hubs, small scale manufacturing and some oil and chemical processing. This combination results in an unusual intermingling of natural and man made acoustic phenomena. "Pandan Musings" is supposed to reflect this interaction between the sounds of nature and man made sounds, to some degree. The sound of the flute is represents the natural world, and the multi-agent system contributions mainly represent the man made sounds. (There are also a few agents whose output are suggestive of natural sounds.)

2.3 Basic Composition Procedure

The basic procedure adopted to create each movement of the piece is an iterative process:

- Create a draft of the solo flute part, based on manipulation of motivic melodic material. Each movement makes use of a distinctive rhythmic profile, and is based on a particular scale. The flute material deliberately adopts some idiomatic characteristics of commonly performed flute repertoire written in the early twentieth century.
- Assemble the multi-agent interactive composing system. The composer has developed a collection of different agents that use different strategies perform different musical roles, implementing a variety of sonic transformation techniques. Appropriate agent designs are chosen to form an initial ensemble of agents for each movement. The choice of agents within the ensemble is a major compositional decision, and is explained further below.
- Use the playback of a recording of the draft flute part to approximately set the various parameters that control the agents behaviour.
- Practice performing the piece with a flute, making adjustments the various agent parameters, aiming for a variety of interesting responses from the system.
- Iterate. Often the agents need to be modified or substituted. Sometimes entire new agent designs were developed. The flute score always needed to be altered in the light of the musical context generated by the multi-agent system.

2.4 Agent Design

There are two broad categories of agent that make up the multi-agent interactive composing systems used in "Pandan Musings": performer agents and controller agents. Performer agents transform the input sound in various ways so as to create a new layer in the musical texture. Controller agents work at a higher musical level, and either determine which agents are heard by the audience, or specify goal states that determine the behaviour of the performer agents.

2.4.1 Performer Agents

Performer agents continually monitor to the live audio signal produced by the human performer. The instantaneous pitch and amplitude data, as well as cumulative statistical data, are used in the agents decision making process to determine how an agent alters the live input signal.

There are two broad categories of performer agents. The simplest agents apply common signal processing effects to

the live signal. These effects include echo, chorus, flanging, spectral delay and ring modulation. Most of these processes produce timbres that are recognisably transformations of the live input signal. Some agents in this category use combinations of DSP processes to produce output that is suggestive of common natural sounds, such as birds, thunder and wind. As each agent makes use of an analysis of the live flute signal to alter the controls of these DSP effects, the musical output of the agent is directly affected by the input signal. Performer agents using this design are frequently used to shape the overall ambience the ensemble output.

The other type of performer agents perform more of a textural/structural role, and have been the focus of much of the system development. These agents take the live flute signal and transform it into new, distinct musical lines. Each agent is optimised to perform particular musical functions, and may drastically alter the flute input waveform creating timbres that are not flute-like at all. The signal processing techniques employed in this type of performer agent range from pitch shifters (implemented using granular synthesis) to wave shapers, various types of filtering and modulation, plus synthesizers that implement the Karplus-Strong plucked string algorithm. Using these techniques, the live signal can be transformed into drones, parallel harmonies, melodic lines, ambience effects, and pointillistic clouds of sound. An important characteristic of these agents is that all of the resulting audio output is a transformation of the live input signal, so the human performer maintains an element of control over the entire ensemble sound.

The basic operating principle of this type of performer agent is:

- Determine the current pitch of the live signal.
- Calculate the next pitch the agent will play, using some strategy specific to that agent.
- Calculate the interval between the input pitch and the calculated pitch.
- Transform the live signal so that it has the required pitch using the signal processing technique built into the agent.
- Calculate the length of the new note.
- Play the note for the required duration.

The structure of this type of agent is shown in figure 1.

There are a variety of algorithms employed in the various performer agents used in "Pandan Musings" that determine exactly how the agent will create the pitch and duration of the notes it plays. The algorithms implemented include stochastic algorithms, such as Gaussian and Exponential random number generators[7], and tracking and evasion algorithms. Many agents are implemented using a design built around two blocks of memory (arrays) which store the outputs of the chosen compositional algorithms. One array contains data representing the note duration, and the other one contains data for the pitch. An algorithm that makes use of the contents of the duration array to derive an index to read data from the pitch array which is then transformed with some simple mapping functions, determine an agents musical output.

An important aspect of the design of the performer agents is their ability to adjust themselves so as to exhibit specific

musical behaviour. They achieve this by modifying their internal states (the contents of the two arrays) so as to meet externally specified targets. The two targets used to control the musical output of agents in "Pandan Musings", are average pitch, and average duration. In order to enable an agent to readjust its internal state so as to meet any particular supplied target, a simple form of gradient descent learning is used. Each agent periodically calculates the current average values of the data stored in the pitch and duration arrays, and compares these to a target average pitch and a target average duration. This produces two error measurements: one for the pitch and one for the duration. The two error measurements can then be used to alter the contents of each array slightly, so as to reduce the error. Periodically repeating this process several times a second, eventually results in the agent to converging on the target behaviour. See [4] for more about this approach.

2.4.2 Controller Agents

The multi-agent systems implemented in "Pandan Musings" also make use of two higher-level agents that affect the overall musical output. One agent supplies the target parameters that are used by the various agent performers to individually adjust their internal state (the pitch and duration targets). Another agent acts as a "mixing engineer", and determines which performer agents are heard by the audience. Both of these agents are implemented as finite state machines that use of an analysis of the live signal to determine their state, and thus shape the musical output. Finite State Machines are very simple to design, and are very efficient and can create the illusion of complexity. They have been used for many years to create simple Artificial Intelligence systems in computer games. The two controller agents in "Pandan Musings" can each be in one of eight different target states, and the parameters of each state are chosen in the composition process so as to create a particular musical effect.

When creating a new composition with this type of system, for each movement, the composer needs to:

- Determine the eight different target states for each performer agent. These are vectors consisting of target average pitches and average durations. Setting these values is an extremely important compositional decision, as they collectively determine the behavioural extremes of the piece.
- Determine the eight different target states for controller agent that controls the mix. This sets the extreme volume levels for each agent and has a major impact on the resulting musical textural possibilities.
- Assign each state to a 3D coordinate. (These correspond to the different corners of a cube and will be used in the decision making process).

In operation, these agents:

- Periodically derive three values from the live signal that will be turned into a 3D coordinate that represents the current attributes of the live signal. This is done using a mapping function, fine tuned for each specific movement, that manipulates the pitch and duration data to produce these values. The exact mapping function is a significant compo-

sitional decision that needs to be made for each piece, and is usually the result of an iterative trial and error process. The basic analysis data, that is fed into the mapping function, is obtained from the input signal is using a fiddle (or sometimes sigmund) object. This data is accumulated for a particular time interval, and then some sort of statistical analysis is performed and then is periodically transformed by the mapping function. For example, the system could be set up so that the mean pitch and melodic range of the flute part, the duration of the last phrase is measured, and are then scaled to become a value in the range -1 to 1, and then combined to produce the 3D coordinate.

- Calculate the Euclidian distance of this 3D coordinate from each of the vertices of the cube.

- Set the system to the state that has the smallest Euclidian distance from the 3D coordinate.

The fact that this is a deterministic processes means that there is an element of predictability in the behaviour of the system. This means that similar musical input gestures will tend to produce similar musical ensemble output, which enables the human performer to learn how to shape the response of the agent ensemble.

2.5 Restrictions imposed on the multi-agent system for "Pandan Musings".

It should be noted that the multi-agent interactive composing systems in "Pandan Musings" have been optimised to achieve the aims of the piece. Many constraints have been imposed to create a particular type of musical result. For example, there is no constraint on the input waveform used to drive the system. "Pandan Musings" was written for flute, but the multi-agent system can be used with any type of input. The author has used these systems in public performance with voice, saxophone, analog synthesizer and collections found objects. It is also possible to adapt this type of system to have multiple inputs. Another effective way to use the multi-agent system is to use its output to be the live signal input via unidirectional microphone, creating a feedback loop which can be controlled by the positioning and orientation of the mic.

"Pandan Musings" also restricts the pitch of the output of the multi-agent system, constraining it to a particular scale. This restriction was a compositional choice, which required additional programming to achieve. Usually, when using this type system with input devices such as analog synthesizers and found objects, the scale mapping system is turned off, producing an output that makes use of the pitch continuum.

The timing of the output of the multi agent system is quantised to a time grid. In some movements of "Pandan Musings" a compositional choice was made to provide a clear pulse for the performer to play with (or against). It is possible to set the temporal resolution fine enough that any timing quantisation on the musical output is not perceivable

3. AESTHETIC PLACEMENT AND STYLISTIC PREDECESSORS

3.1 Motivic Transformation and musical Phrases

Each movement of "Pandan Musings" has a pre-written flute part, whose style is consciously influenced by flute works written by composers such as Debussy, Faure, Poulenc and Hindemith. The aim was to create melodic material with a lyrical character, with clear sectional forms that achieve cohesion by traditional motivic transformation techniques. Each movement gets some its character due to the (almost) exclusive use of less commonly used scales. Rhythmically, the pieces are characterised by frequent use of triplets, quintuplets, sextuplets and septuplets. There is some use of metrical changes, usually to elongate or contract the current melodic phrase.

3.2 Interactive Composing Elements

The interactive composing systems used in "Pandan Musings" are directly influenced by Joel Chadabe's [3] ideas about interactive composing, that he developed in the late 1970's. A human performer controls high level aspects of the piece, while the compositional algorithms make the low level, note to note, decisions which are realised using a synthesis algorithm in real time.

Another very strong influence on the piece is the tradition of applying interactive real-time signal processing to a live instrumental performance. This has been practiced in a variety of forms for over half a century. Techniques range from simply manipulating the controls on a collection of simple guitar stomp boxes, to complex computer based systems using sophisticated user interfaces, hardware controllers or artificial intelligence systems. Usually, a large amount of attention is focused on operating the signal processing, and one or more performers may focus on this exclusively.

An attractive feature of both of the above approaches is the immediacy of the systems response to the performers gestures. The "Pandan Musings" signal processing systems are an attempt to create a responsive system that combines these two approaches. It makes use of algorithmic composition techniques to create the "score" for various musical layers and uses real-time signal processing to realise them. A key feature of the piece is that the sound of the flute performance acts as both the high level performance controller, and is also the origin of all of the sound heard in the piece. Analysis for the musical material performed controls the compositional algorithms, creating the high level musical features of the agent ensemble output, while the subtle flute performance nuances have a significant impact on the sound of the ensemble.

4. CONCLUSIONS

The application of a multi-agent interactive composing system to create an accompaniment for a fixed score composition has been demonstrated to be a viable approach. It enables a useable balance of predictability and surprise,

allowing room in each performance for happy accidents, each performance bringing a fresh interpretation.

5. REFERENCES

- [1] [Online]. Available: <http://puredata.info/>
- [2] I. Whalley, "Software agents in music and sound art research/creative work: current state and a possible direction." *Organised Sound*, vol. 14, pp. 156–167, 2009.
- [3] J. Chadabe. (1984) Interactive composing: An overview. [Online]. Available: <http://www.joelchadabe.com/articles/interactivecomposing.pdf>
- [4] M. Spicer, B. Tan, and C. Tan, "A learning agent based interactive performance system." in *Proceedings of the International Computer Music Conference*. Singapore: International Computer Music Association, 2003.
- [5] M. Spicer, "Spaa: An agent based interactive composition," in *Proceedings of the International Computer Music Conference*. Barcelona, Spain: International Computer Music Association, 2005.
- [6] M. J. Spicer, "Agents in chuck a timely programming experience," in *Proceedings of the International Computer Music Conference*. New Orleans, USA: International Computer Music Association, 2006.
- [7] C. Dodge and T.A. Jerse, *Computer Music Synthesis, Composition, and Performance*. New York NY, U.S.A.: Schirmer Books, 1997.




Figure 1. Structure of a performer agent.

Scores

Michael Spicer

Flute Trio in the Diminished Scale

Flute Trio in the Diminished Scale

1. Impulse

Michael Spicer © 2012

♩ = 180
Energetic

Flute 1

Flute 2

Flute 3

6

11

Fl.

Fl.

Fl.

mf

mf

mf

mp

mf

mf

mp

mf

18

Fl. *mp*

Fl. *mp*

Fl. *mp*

mf

29

Fl.

Fl.

Fl.

mf

37

Fl.

Fl.

Fl.

mf

42

Fl.

mp

f

mp

f

Fl.

p < mf > p simile

Fl.

p < mf > p simile

54

Fl.

mp

f

mp

f

mp

Fl.

mp

Fl.

mp

66

Fl.

f

Fl.

mf

Fl.

mf

74

Fl.

Fl.

Fl.

81

Fl.

Fl.

Fl.

mp

mf

89

Fl.

Fl.

Fl.

f

97

Fl.

Fl.

Fl.

f

6

101

Fl.

Fl.

Fl.

f

The musical score consists of three staves, each labeled 'Fl.' on the left. The first staff (top) contains a melodic line for the first flute, starting with a half note G4, followed by a slur over a quarter note F#4, a quarter note G4, and a quarter note F4. The second staff (middle) contains a rhythmic accompaniment for the second flute, primarily using eighth and sixteenth notes. The third staff (bottom) contains a rhythmic accompaniment for the third flute, also using eighth and sixteenth notes, and begins with a forte (*f*) dynamic marking. The score is divided into four measures by vertical bar lines, ending with a double bar line. The measure number '101' is written above the first staff at the beginning of the first measure.

2. Stasis

$\text{♩} = 80$

Dream like

Flute 1

mp *mf* *mp* *mf* *etc*

Flute 2

mp *mf* *mp* *mf* *etc*

Flute 3

mf *mp* *mf* *etc*

8

Fl.

Fl.

Fl.

rall *rall* *rall*

3. Resolution

8 $\text{♩} = 125$
Resolutely

Flute 1 *mp* *mf*

Flute 2 *mp* *mf non legato*

Flute 3 *mp* *mf*

14

21

28 *non legato* *mf* *f* *mf* *f*

36

Fl.

Fl.

Fl.

mp

mp

mp

47

Fl.

Fl.

Fl.

mf

59

Fl.

Fl.

Fl.

66

Fl.

Fl.

Fl.

mf

non legato

10

75

Fl.

Fl.

Fl.

f

mf

mp

85

Fl.

Fl.

Fl.

mf non legato

mf

mf

94

Fl.

Fl.

Fl.

f

mf

100

Fl.

Fl.

Fl.

Rall

Flute Trio in the Diminished Scale

Flute 1

1. Impulse

Michael Spicer © 2012

♩ = 180 **Energetic**

2

mf

5

8

mp

12

mf *f*

17

mp

5

27

mf

33

40

5

mp *f*

51

mp *f* *mp* *f*

59

mp *f* *mp*

4

68

f

72

76

80

83

5

f

91

96

101

$\text{♩} = 80$
Dream like

2. Stasis

mp *mf* *mp* *mf*

7

etc

11




rall

3. Resolution

$\text{♩} = 125$
Resolutely

Flute 1

3



43 *mp*

50 *mf*

57

64 *mf*

73

78 *f*

84 *mf non legato*

92 *f*

96

100

104 **Rall**

The musical score for Flute 1 consists of ten staves of music. The key signature is one flat (B-flat). The tempo and dynamics are indicated by markings such as *mp*, *mf*, *f*, and *mf non legato*. The score includes various musical notations such as notes, rests, slurs, and dynamic markings. The measures are numbered 43 through 104. The score ends with a double bar line and a repeat sign.

Flute 2

Michael Spicer © 2012

Musical score for "The Rose Tree" in 7/8 time. The score is written on a single staff with a treble clef and a key signature of one flat (B-flat). The tempo is marked "Allegretto". The score includes measures 1 through 38, with dynamic markings such as *mf* and *mp*. The melody is written on a single staff with various note values and rests.

Flute 2

41  *p < mf > p*

45  *simile*

53 

61  *mp*

67  *mf*

72 




77 

82  *mp*

87  *mf*

92 

97  *f*



34 *mf*

39 *mp*

46

53

59

63

67 *non legato*

72

78 *mf*

84 *mf*

91

96



101

Rall

Flute 3


Flute Trio in the Diminished Scale

1. Impulse

Michael Spicer © 2012

Energetic

♩ = 180



23



29



35



39



41



45



53



61



68



74



80



86

mf

92

98

f

102

$\text{♩} = 80$
Dream like

2. Stasis

mf *mp* *mf* etc

7

12



rall

3. Resolution

4

$\text{♩} = 125$
Resolutely


Flute 3



70



76



83




89



94



100

104 **Rall**

Momentary Diversions

Michael Spicer (2014)

Construction Guidelines

Momentary Diversions is an improvisation environment for modular analogue synthesizer. The piece is characterised by the approach it takes to the way the modules are connected, rather than the specifics of the particular instance of the piece. This implies that there may be many realisations of the piece, and the piece can be implemented with a variety of synthesiser hardware.

The piece consists of constructing, and performing with, a multi-agent interactive composing system. The performer builds a collection of autonomous musical performer agents, exhibiting specific characteristics, each manifest as an analogue modular synthesiser patch. Each agent contributes a sonic layer to the musical texture. The final music output is determined in real-time by controlling the output level of each agent via a mixer. The multi-agent system should instantiate at least six agents, and make no use of sequencer modules. Some agents should have sensors that gather input from other agents, allowing various types of interactions between voices, via control voltage inputs of the modules concerned.

Agent designs can be classified into two basic categories, based on their strategies to create musical material:

- **Generative** patches
 - Run autonomously, with no need for human intervention.
- **Interactive** patches.
 - Rely on human interaction to shape some element of the musical output.

Generative patches can be grouped into those that are constructed on

- **Stochastic** processes.
 - These are built around noise generators and/or shift registers.
- **Cyclic** processes.
 - These are built around oscillators, or devices that can be configured to act as oscillators.

Interactive patches can be similarly divided into those that:

- Combine **human control** with some **stochastic** control system.
 - Performance controllers offer predictable control over some aspects of the musical output, while other details of the sound are determined with cyclic or stochastic processes.
- Largely shaped by **human** gestures.
 - Performance controllers provide most of the control of the musical output. Some elements may be predictably controlled with envelope generators and/or low frequency oscillators.

This combination of strategies allows each agent to exhibit distinct and complimentary musical characteristics, enabling each to carve out their own space in the sonic spectrum.

Examples of **Generative Patches** typical of **Cyclic** strategies include:

- The creation of low, or high, drones which exhibit complex movement in the harmonic content, achieved by modulation systems built with one or more low frequency oscillators.
- The creation of layers that exhibit cyclic properties that are manifest at two or more hierarchical levels. An example of such a system is a configuration where a slow low frequency oscillator sweeps the pitch of an oscillator, while another low frequency oscillator modulates the amplitude, with a different waveform at a different frequency.

Examples of **Generative Patches** typical of **Stochastic** strategies include:

- The creation of a number of short sounds that collectively exhibit a sense of unity via some random process patched to control a perceivable element, such as the average pitch of an oscillator.
- The creation of interlocking rhythm layers, realised by triggering the amplitude envelopes of several independent synthesizer voice layers, with a shift register.
- The creation one or more layers of sustained gliding tones, realised by creating oscillator frequency control voltages by slew limiting the voltages generated by periodically sampling coloured noise.

Example of an **Interactive Patch** with **Stochastic** elements include:

- The creation of a melodic, or foreground, element that combines a significant random element with a continuous control mechanism. An example of such system can be constructed with a pressure sensitive touch sensor system that is able to produce a gate signal, plus a control voltage,. The gate signal triggers the generation of a random control voltages, which is applied to control the frequencies of two oscillators that modulate each other in some way, resulting in a different complex timbre for each gesture. The pressure sensor control voltage shapes the volume of this signal. The result is an expressive system whose output is a collaboration between the generative subsystem and the human performer.

Example of an **Interactive Patch**, mainly shaped by **human** gestures.

- Common synthesiser voice patches, consisting of oscillator, filter, amplifier chains, such as those descended from the Minimoog Model D architecture. It could be controlled by any of synthesizer controllers, such as keyboard, wind controller etc.
- Theremin, or other spatial controlled instrument.

Pandan Musings
and
Pandan Meditations

Pandan Musings

Pandan Musings is designed to be performed with the accompanying Pd patches.

Each movement has its own Pd patch.

The computer part should be amplified by two speakers on the stage, either side of the performer.

The instrumental part should be routed to the computer via a dynamic mic and a suitable audio interface. The Pd patches implement a multi-agent based interactive composing system, which transforms the solo instrument into an ensemble.

Each movement is a conversation between the instrumentalist and the agents in the interactive composing system.


It is important to spontaneously make decisions regarding tempo, dynamics and articulation, in response to what the multi-agent system is doing.

Any dynamic or articulation marks indicated on the score are guidelines only.

Extended instrumental techniques should be avoided.

4


47



51




54



57



62



65



Meditatively

 $\text{♩} = 35$

5

9 Theme A

13

17 Trans

21

25 Theme B

29

33 Retransition

37

41

The musical score is written in G major (one sharp) and 6/8 time. It consists of nine staves of music. The tempo is marked as $\text{♩} = 35$. The score includes various rhythmic patterns and melodic themes. The first staff (measures 1-8) features a complex rhythmic pattern with 4 and 5:6 groupings. The second staff (measures 9-16) continues this pattern. The third staff (measures 17-24) introduces a new theme labeled 'Theme A'. The fourth staff (measures 25-32) continues 'Theme A'. The fifth staff (measures 33-40) introduces a new theme labeled 'Theme B'. The sixth staff (measures 41-48) continues 'Theme B'. The seventh staff (measures 49-56) continues 'Theme B'. The eighth staff (measures 57-64) continues 'Theme B'. The ninth staff (measures 65-72) continues 'Theme B'. The score includes various rhythmic patterns and melodic themes. The first staff (measures 1-8) features a complex rhythmic pattern with 4 and 5:6 groupings. The second staff (measures 9-16) continues this pattern. The third staff (measures 17-24) introduces a new theme labeled 'Theme A'. The fourth staff (measures 25-32) continues 'Theme A'. The fifth staff (measures 33-40) introduces a new theme labeled 'Theme B'. The sixth staff (measures 41-48) continues 'Theme B'. The seventh staff (measures 49-56) continues 'Theme B'. The eighth staff (measures 57-64) continues 'Theme B'. The ninth staff (measures 65-72) continues 'Theme B'.

6

44



Pandan Musings

III

Resolutely

To be performed with PandanMusingsDLydian.pdf

7

♩ = 60

5

10

15

19

23

27

31

35

39

43

8

48




52



Pandan Meditations

Pitches contained in PandanMusingsCLydian Dominant.pd

9



3 Pitches contained in PandanMusingsDOriental.pd



5 Pitches contained in PandanMusingsDPhrygian.pd



7 Pitches contained in PandanMusingsENeapolitianMajor.pd



9 Pitches contained in PandanMusingsEFlatLydianMinor.pd



Michael Spicer

Sky Castles

Performance Instructions

Begin playback of Sky Castles Logic Project before commencing to play.

Hold each note for a full breath.

Leave small gaps between each note, to give focus to the signal processing.

Hum whilst playing the notated pitches.

Improvise the pitch contour of the note you hum. It may contain any combination of:

- Hover around the pitch of the notated note, adding small upward and downward inflexions.

- Glide up to the pitch of the notated note.

- Fall away from the pitch of the notated note.

Respond to the signal processing.

Sky Castles

Michael Spicer

Flute

mf

6

Fl.

The musical score for 'Sky Castles' by Michael Spicer consists of two staves. The top staff is labeled 'Flute' and the bottom staff is labeled 'Fl.' (Flute). Both staves are in treble clef. The top staff begins with a mezzo-forte (*mf*) dynamic marking. The bottom staff begins with a measure number '6'. Both staves contain five measures of music, each featuring a half note with a fermata. The notes are: G4 (first measure), E4 (second measure), D4 (third measure), C4 (fourth measure), and B3 (fifth measure). The piece concludes with a double bar line at the end of the fifth measure on the bottom staff.

Sonic Escapade

Michael Spicer (2008)

Rules of engagement.

(Word-based score: set of verbal instruction along the lines John Cage , Alvin Lucier)

General Structure

- This is a piece for 4 -9 players (optimally 5)
- Each player autonomously creates his/her own part.
- The piece consists of a number of distinct episodes.
- Each episode should last less than 4 minutes.
- Each episode should have a distinct timbre.
- Each player can either:
 - Cooperate (the default behaviour)
 - Free spirit
 - Disrupt (especially when used to trigger transitions to new episodes)

Instrumentation

- Each player has a least 1 percussion instrument, 1 chordal instrument, 1 melodic instrument and bass instrument (ideally).
- Players should only play one instrument at a time, but may sing and play.
- There should be only one of any instrument in the ensemble.
- Some players have access to electronically amplified sound sculptures, consisting of a collection of found objects that may be hit, scraped, bowed etc.
- Conventional instruments are drawn from both Western and Asian cultures, and may be electric or acoustic.
- At least one instrument should be hand made (or assembled) by the performers.
- There should be at least one pvc pipe instrument.
- There should be some analog electronic instruments (eg, analogue synthesizers , theremin, home brew circuits etc.) Digital instruments may be substituted, as long as they provide extensive real-time interaction capabilities.

Strategies for choosing musical material.

- For each episode, one player must take the initiative to set the character.
- Each player chooses what they play, according to the chosen strategy, (clarify) but must always be conscious of what the other players are doing.
- Players should be extremely conscious of the resultant ensemble musical texture at all times, and should “rest” (not play) where they feel it is appropriate.
- Some sections may use diatonic scales
- May use cyclic chord progressions, but should be at the complexity that the players can play by ear (dependent on the ensemble)
- May make use of:
 - Bass line idioms
 - Pad
 - Idiomatic “comping” patterns
 - Arpeggios
 - Melodic line
 - Counter melody
 - Parallel harmony
 - Call and response
 - Series of pointillistic events/sound cloud
 - Sound mass
 - Rhythmic layering/Cyclic looping
 - Percussion patterns
 - Drones High/Low
 - Delay/ pitch shift/feedback / signal processing techniques