Chapter 1

Introduction

> "In spite of being scientists, three gentleman consented to an experiment which must have seemed very strange at first sight, namely the marriage between music and the world's most potent machine ... With the help of an electronic brain the composer turns into an astronaut pressing buttons of his musical spaceship to introduce co-ordinates and keep the course of his vessel on its journey through constellations and galaxies of sound, controlling from his easy-chair what the imagination of yesteryear could have envisaged only remotely in its wildest dreams." - Xenakis (1971a) p.124–133

It is an exciting time to be a musician. The potency and portability of Xenakis' electronic brain has increased to the point where we no longer even need to press buttons in the cockpit to keep it on its journey: we can steer it through unimagined realms with remote gestures or sonic provocations; we can programme it to learn from and replicate the works of composers past and present; we can even programme it to to *lead* us around spaces beyond our *own* imaginations. And these new universes are not just of the sonic variety, but inhabited, explored and invented by practitioners from every corner of the arts allowing cross-fertilisation of ideas and techniques and opening possibilities for new practices. What is more, we are not alone in developing navigational strategies: our contemporary scientific colleagues offer a multitude of conceptual and technical know-how to be put to use on our expeditions.

Art has always been driven by an urge to explore, to create, to mimic and to come to terms with the world around us. In this respect the marriage of computers and music is not very strange, but entirely expected. Art, technology and science have always been locked in an intimate coevolution, the products of technological development continually fuel our creative endeavours, that in turn drive new technical innovation. At the same time the processes that support our creative outpourings have become hot topics in scientific research.

Throughout history the development of musical culture has been intimately influence by technical developments with new technologies constantly changing the ways in which music is composed, performed, preserved and distributed. Digital techniques have refined and expedited practice across the arts, simulating existing tools to make them more efficient, more flexible and easier to use: we can record, produce and master an album on a laptop, pitch shifting out some bad intonation; we can clean up photos, removing the blemish on the bride's nose and brightening the sky whilst undoing any errors; we can tween between keyframed poses of an animation, saving hours of time drawing transition frames by hand. But computers are not only good at replicating existing tools: they offer many new ways of working.

The development of creative software is inevitably tied closely to trends and advances in the computer sciences. In the last few decades there has been extensive research into applications of information processing, drawn in particular from Artificial Intelligence (AI) research, to the development of musical composition and interactive performance software. In fact it has been suggested that the development of systems that can create music in established styles is one of the major achievements of AI to date. In this area techniques of mutual interest to musicologists, cognitive scientists and composers have been developed through which computers can not only be imbued with musical know-how in the form of explcit instructions, but can be programmed to derive representations from existing musical works. Learning mechanisms and search strategies from AI have been deployed to develop programmes that can harmonise chorales, improvise bee-bop or extract patterns from existing works and generate convincing pastiches of the supplied material.

Similar techniques have also been developed to assist the contemporary composer's search for new musical idioms: new sounds, new textures, new means of expression. The picture of the composer or artist, set apart from worldly influence, fed by divine inspiration is common but mythical. We might like to uphold the Mozartian image of Composer as Visionary, but in reality most artists are fuelled in part by serendipity, gathering ideas that spring up in their active interaction with the tools of their trade. Composers for instance might sit and improvise open mindedly, saving ideas that arise for future development. Computers can not only help in this process, but offer new ways of systematically exploring new ideas.

Whilst the computer can only do what we tell it to, it can do so very quickly, and the results are not always necessarily something we could have predicted. For composers wishing to systematically expose themselves to novel ideas, or explore whimsical complexities, computers are not only a time saving device, but open up whole new realms of possibility. In wanting to step outside the confines of their own imagination, some composers have employed random processes - more or less sophisticated digital dice rolling techniques - to shake up the material of their practice. Other have implemented probabilistic frameworks, not dissimilar to Xenakis', defining broad fields of possibilities and employing the computer to generate the detail. Approaching composition in this way, the implicit process underlying the act of composition becomes reified in formal language: crudely put, arranging notes becomes designing processes.

The use of processes to induce unknown outcomes is nothing new. Most infamously perhaps advanced by the Experimentalists who dismissed fixed forms and prespecified structural relationships in favour of exploring ways of "outlining situations in which sounds may occur" (Nyman (1999) p.10). John Cage in particular is renowned for his use of various sources of chance to specify processes that bring about "acts the outcomes of which are unknown" (Nyman (1999) p.10). More recently Brian Eno has enthused about his fascination with inventing systems and machines that " ... make music with material and processes I specified but in combinations and interactions I did not ... (Eno (1996)). This interest in relinquishing control to a definite process with an *in*definite outcome characterises much current activity in the broad church of 'generative art'. Whilst the scope and rate of evolution of generative art practice defies any strict definition, Philip Galanter's description of the key elements of the approach is widely accepted:

> "Generative art refers to any art practice where the artist uses a system, such as a set of natural language rules, a computer program, a machine, or other procedural invention, which is set into motion with some degree of autonomy contributing to or resulting in a completed work of art." - Galanter (2003), p.4

On these terms, we can include the use of the computer to facilitate our specific artistic aims, as in Xenakis' use of the electronic brain to expound the sonic virtues of Gaussian Galaxies. We can include the use of the computational procedures to mass produce media, or to demolish the hierarchies of authorship. We could also include the adoption of AI techniques to mimic existing styles mentioned above. But for many generative artists the appeal is not just the ease of production, absence of repetition and never-ending something-else: its the promise of something *more*, something new and surprising.

The promise of excess and emergence of new ideas and new forms is one reason why many generative artists have adopted techniques from Artificial Life (Alife) research. In contrast to traditional AI, which concerns itself with developing representations of a highlevel central processor capable of human-level cognitive tasks, Alife concerns itself with studying the low level interactions of distributed processes from which coherent behaviour emerges as a global product of the system. One of the simplest and most famous examples is the set of cellular automata (CA) rules called the *Game of Life*, devised by John Conway (first described in print in Gardner (1970)).

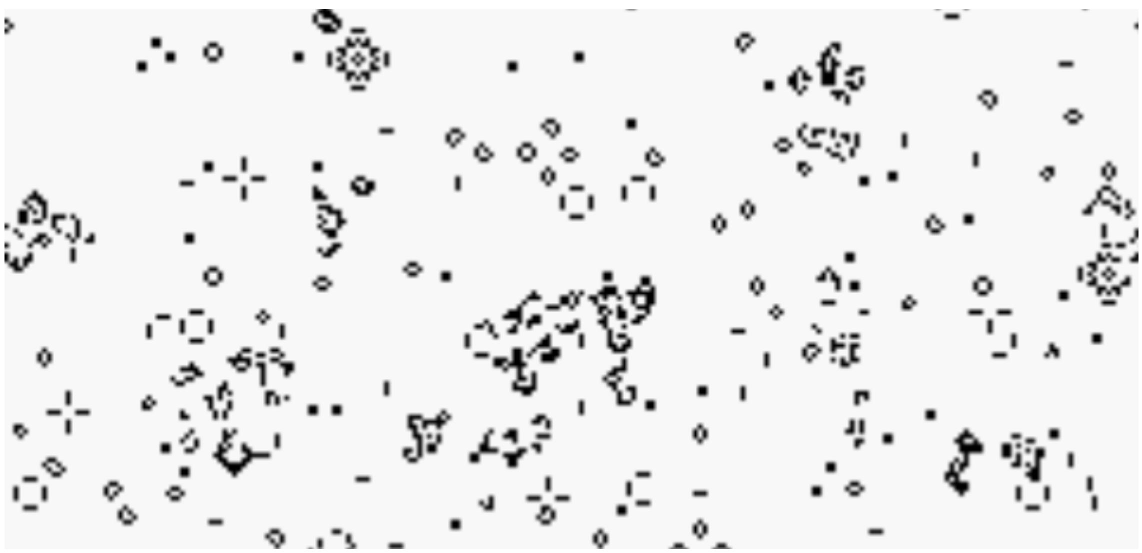

Figure 1.1: Screen shot from the Game of Life, a CA rule set devised by John Conway.

The Game of Life is a set of rules that specify the state of cells on an infinite two-dimensional grid i.e. whether they are 'dead' or 'alive'. In Figure 1.1, black dots are alive, dead cells are white. The state of each cell is determined by its eight immediate neighbours – those horizontally, vertically or diagonally adjacent. At each time step, the following set of rules gets applied to all the cells, setting their state for the next time step:

- Any live cell with fewer than two neighbours dies, as if by loneliness.
- Any live cell with more than three neighbours dies, as if by overcrowding.
- Any live cell with two or three neighbours lives, unchanged, to the next generation.
- Any dead cell with exactly three neighbours comes to life.

From a computer scientific point of view the significance of this rule set was that it had the power of a Universal Turing machine (i.e. anything that can be computed algorithmically can be computed with the Game of Life), but for physicists, biologist, economists, mathematicians, philosophers and Alife researchers, it represents a prime example of emergence and self-organisation.

According to the state of the cells at the start, different configurations and patterns emerge. Imaginatively named after things in the real world ('boats' and 'toads') or worlds of science fiction ('lightweight space ships', 'gliders 'and 'glider guns'), these are not only static forms, but coherent patterns of organisation that move across the grid. The rules describe what happens in each static cell, and yet what we see can only be sensibly described in terms of movement *across* cells. A verbal description of what is going on demands recourse to a formal level of description that is absent in the initial specification.

These systems are not alive in any sense of the word, but emulate operational aspects of biological convention without simply mimicking nature. There is something about their movement – their behaviour – that invites attributions of intentionality. Similar models have been used by visual artists to create forms that similarly display an uncanny agency. We have seen flocks of strange silicon bird-like forms or schools of abstract digital creatures, breeding and chasing one another. Robotic forms that are not explicit simulations, not merely automaton but substantially themselves.

Their sense of presence is due in part to their reactive nature. Many are defined in terms of local rules that are sensitive to their environment, creating digital creatures can be convincingly engineered to recoil at our touch, to follow us through space or hide in dark corners of their silicon worlds when we enter the room. Many visual artists have created virtual critters of varying degrees of abstraction whose movements and behaviours exhibit a degree of agency that belies their formally specified origins. The dynamic and flexible ways in which they respond to us (informed by sensors, video cameras and the like), has led researchers such as Ken Rinaldo (1998) to suggest that Alife techniques present "opportunities for both artists and viewer/participants to develop true relationships with the computer that go beyond the hackneyed replicable paths of "interactivity" which have thus far been presented by the arts community". For Rinaldo the time has come to indulge in a "cybernetic ballet of experience."

An impression of intentionality also arises from their often unpredictable nature. Although every aspect of a system is formally defined, the outcomes can at times be surprising, exceeding the expectation of the programmer-artist. In discussing their work many artists give the impression that they are pursuing a general interest in creating something that goes beyond what they specified: to achieve that "something more" (Whitelaw (2004)). Simon Penny writes, "I am charmed and fascinated by the possibilities of complexity theory and emergent order" (Whitelaw (2004) (p.216)). Robb Lovell expresses his interest in "going after creating something that gives me more than I expected" (ibid).

Similar techniques have been explored to an extent in musical applications. Many researchers have employed biologically inspired search mechanisms such as Genetic Algorithms (GAs), which harness the power of Darwinian evolution to systematically explore a defined, yet vast, space of musical possibilities. Others have used agent-based models of evolving ecosystems to create fluctuating populations of sound. CAs like the Game of Life and models of biological growth such as L-systems have also been extensively explored to generate abstract melodic paths or deployed in sound design tasks. But investigation has been limited to a small handful of models which have almost exclusively been applied in compositional tasks.

Rinaldo used the phrase 'Cybernetic Ballet of Experience' in the context of installation art. The vision behind the current project is to bring this sense of a digital 'other', that has made an appearance in the visual installation arts, into the sonic domain and onto the stage. The techniques that inject life into these visual creatures scurrying about on screens are crying out to be let loose on stage and offer an intriguing alternative to traditional approaches to interactive music. The origins of life and mind are far from solved, but the conceptual and technical tools used to tackle these epistemological issues offer a tool box which inspires a 21st Century self-steering upgrade of Xenakis' space ship. To borrow a turn of phrase from George Lewis, perhaps the time has come when we can programme computers that we not only can play our music on, but which will play their music with us (Lewis (2006)).

1.1 Summary of contributions

This thesis aims to expand upon the dominant design strategies for composing and performing with computers, introducing simple adaptive systems as mechanisms for both generating musical material and affording a novel approach to interaction with digital systems. Whilst many arts practitioners have appropriated Alife models, these invariably focus on agent-based simulations or EC methods, hankering after glamourous 'out of control' properties of emergence and self-organisation. Within music these have tended to remain in the domain of composition rather than performance.

A particular class of formal system is ear-marked as most suitable for this task. These are described as *simple adaptive systems* that are rich in both generative and interactive possibilities. This set of models adds to the compendium of tools available for computer musicians to play with. These tools are of both a practical variety and conceptual utensils which it is hoped will inspire the extension of the current collection.

The thesis is split into two halves. The first half provides a context for the current work and surveys and discusses work in the fields of Interactive Computer Music and Computer Automated Algorithmic Composition. The second half presents a practical exploration of the use of simple adaptive systems in different musical contexts from stand-alone generative music systems, to physical interactive installations, culminating in performance systems for man-machine improvisation. These practical applications represent initial, and minimal implementations of a fresh perspective on digital generative and interactive arts which is developed in the first half.

Chapter 2 examines the notion of interaction and considers how this is affected by the move from acoustic to digital instruments. It is suggested that the *active* nature of the digital medium makes new forms of interaction possible and inspires a model of interaction that is different to that associated with acoustic instruments. A conversational metaphor is introduced highlighting the mutual influence between performer and instrument in contrast to the one-way control we have over standard acoustic instruments.

In order to provide a conceptual framework for understanding these differences, heuristic concepts drawn from behavioural robotics and the philosophy are introduced. Specifically the notion of autonomy is explored and an operational definition raised as a useful conceptual and practical tool for the development of interactive systems. Conceptually this is useful as it provides a perspective for appreciating how a software system can be at once independant from, yet sensitive to its environment. Practically an operational definition outlines the type of system architecture that can realise this, albeit minimally.

With this in mind, Chapter 3 focuses on existing approaches to the generation of musical material and proposes that simple adaptive systems offer attractive features as generative composition tools. Included in this chapter is a gentle reminder to practitioners of algorithmic composition that music is not a natural kind, and the relationship between music as experienced and music theory is not the same as that of theory and phenomenon in the physical world. In this respect a mild warning is raised that music theoretic tenants may not necessarily be a suitable basis for designing mechanisms for generating music.

An experimental offering is made in Chapter 4, with the results of an experimental psychology style study that was designed to investigate whether formal properties of complex systems, which are evident in graphical representations, could be similarly appreciated from an auditory display. The results showed that at least in the current conditions, the highlevel properties of complex systems can be appreciated in audio. This is an assumption which is made by practitioners who employ extra-musical algorithms in music, but has not been experimentally verified in this context in the past.

The practice based investigations of Chapters 5, 6, 7 and 8 contribute finished works to the broad field of digital generative arts in offering a balance between the 'out of control' Alife systems explored most extensively in visual and installation arts, and the use of mathematical models in algorithmic composition. Specific contributions are made by each of the practical investigations to the distinguishable practices of generative art, algorithmic composition, installation arts and interactive computer music performance within which they sit.

The studies presented in Chapter 5 contribute to extra-musical algorithmic composition by bringing a new set of workable tools. Chapter 6 presents two generative music installations which use these components in combinations, and explore the use of multiple mappings from model to sound. Chapter 7 opens these systems up to the real world and investigates both the use of simple adaptive systems in responsive environments and a cybernetic take on providing feedback.

The Self-karaoke Machine presented as both an installation and performance system in Chapter 8 speaks to both the Alife generative art world and interactive computer music players. To the interactive music world, it offers an alternative approach to man-machine collaboration that allows full expression of both generative and traditional forms of improvisation. To the Alife artists preoccupied with emergence and user's creative freedom, it demonstrates the power of bringing the user into the generative loop. In coupling the formal system with the open-ended dynamical system that is the performer's sonic improvisations, possibilities escape the programmer's intentions and open-ended exploration can occur.

The body of work has acted as a spring board for the creation of *behavioural objects*, a project funded by the EPSRC network Live Algorithms for Music, which aims at to develop tools and an understanding of adaptive systems in improvised computer music. This project promises to gather enough momentum to fulfil my main wished-for contribution: that this work may in some way inspire future explorations of the use of adaptive digital processes in conjunction with acoustic instruments in live performance, bringing together the boundless possibilities of computer simulation with the wonderfully productive constraints of performing live music with acoustic instruments.