Adaptive Systems Music:
Algorithmic Process as a Compositional Tool
EASy MSc Thesis 2002

Alice Eldridge
School of Cognitive Studies
University of Sussex
BN12 4PQ
alicee@cogs.susx.ac.uk

September 30, 2002

Contents

List of Figures

Acknowledgments Thanks to:

- Phil Husbands for encouraging me to explore my random ideas and enthusing over their bizarre outcomes.
- Russell Blakeborough for introducing me to the wonderful world of Borland Builder.
- The rest of my band for welcoming the homeostat onto the stage.
- The organisers of Blip for facilitating a public listening.

Abstract

Applying the Artificial Life ethos in the musical domain, the Adaptive Systems Music (AdSyM) project disregards the high level details of musical theory and attempts to generate music as it'could-be' using abstract models of adaptive systems. Three systems are built, the principal component in each being modified versions of the homeostat described by Ashby. The first 2 systems exploit the homeostatic process to generate harmonies, accompanied by a rhythm derived from the states of a one dimensional binary cellular automata. The 3^{rd} system utilises the functional properties of the homeostat to build a prototype 'trainable sequencer module'. This system allows real-time reassignment of samples to particular channels. A self-regulating homeostatic network is developed for use in the AdSyM II which is shown to preserve the essential properties of the Ashbian homeostat. Initial investigation suggests that the system may also be truly auto-regulatory. This warrants further investigation. The music generated received high commendation in public appearances, and listeners in a more rigorous evaluation conclusively supported its musicality. It is suggested that the techniques developed could be usefully applied in more theoretical areas of the field.

Chapter 1

Introduction

> "New means change the method; new methods change the experience, new experiences change man." [Stockhausen, 1989, p.88]

1.1 Overview

The traditional links between music and mathematics mean that the computer has been readily accepted as a compositional tool. As well as serving as a digital sequencer, facilitating the recording and organisation of existing musical material, the computer has enabled the exploration of a previous fascination with stochastic and rule-based processes in the automatic generation of music. There is already a wide range of generative software available. Although some of these systems are very successful, they are almost exclusively designed to mimic existing styles. Those that do produce original material create only small fragments of music and lack larger scale structure.

In the spirit of Artificial Life endeavors, the Adaptive Systems Music project aims to generate music as it 'could-be' (after [Langton, 1992]). Rather than embedding theoretical constraints in stochastic processes, the project aims to capture fundamental properties of music using [adaptive algorithms]. The systems developed are based on abstract models of the processes of homeostasis (the Ashbian homeostat), and pattern propagation (cellular automata).

This section gives a brief history of the exploration of generative processes in composition including an overview of current techniques of algorithmic composition. Examination of the fundamental elements of music in different time-scales follows before consideration of the suitability of CA as the basis for rhythm and a description of homeostasis in both physiological and Cybernetic systems. Chapter 2 describes the development and testing of the constituent algorithms used in the software. Various homeostatic systems are examined, and a self-regulating homeostatic network developed. Chapter 3 gives a conceptual description of the 3 software systems. Implementation details are provided in chapter 4. In Chapter 5, the systems are evaluated, audio examples being provided on the accompanying cd. A review of the public appearances made by the Autonomous AdSyM is given, as well as description and analysis of the controlled listener evaluation carried out. A discussion and conclusion is presented in chapter 5.

1.2 The Origins of Generative Music

Although music is primarily associated with an artistic expression of emotions, the act of composition is an inherently rational process. Composers have long been fascinated by the application of mathematical concepts in music: procedures that entail rules or provisions to govern composition have been used since the Medieval period.

It is no surprise then, that the computer has taken its place alongside more traditional methods of composition. As well as the development of digital sequencers (essentially simulated analogue multi-track recorders), which facilitate the recording and arrangement of *existing* musical ideas and samples, there is an increasing proliferation of algorithmic software designed to *generate* musical fragments. Whilst the technology facilitating these processes is relatively new, the generative processes themselves have been employed by composers since the start of the last century.

In fact rational methods of composition have been employed since the Middle Ages: Guido d'Arezzo, in the 11^{th} century, developed of a look up chart for assigning pitch to the syllables of religious hymns. Much Medieval music was constructed within this essentially parametrical framework: note sequences derived using systems such as d'Arezzo's were combined with rhythms derived separately from a rhythmic pattern. If these sequences were of different lengths, they would be rotated against each other repeatedly until their ends matched. The note sequences were then developed by applying transformations such as inversion and retrograde inversion.

The birth of formal music Although similar formulaic techniques are evident in music throughout the Renaissance and Classical periods, it was not until the $1920's$ that these methods were formalised and extended beyond the pitch domain to the point at which they could be used to automatically generate entire musical works.

The harmonies of Western tonal music are structured hierarchically. The root (bass note) dominates the chord and defines the harmonic progression. In the belief that the interval between all notes should have equal importance, Schoenberg (1873-1951) developed a compositional method called *Serialism* [Schoenberg, 1950]. Rather than being based around notes defined by the current key, a phrase in Serial music is made up of exactly one instance of each of the 12 semi-tones in the equal-tempered scale. This basic phrase is then manipulated using a selection or elaboration of the following formal transformations[1]:

Inversion:	each element e of the basic set is replaced by $12 - e$.
Retrograde:	the ordering of the elements of the entire set is reversed.
Transposition:	an integer t (modulo 12) is added to each element of the set

Schoenberg applied the principles of Serialism only in the pitch domain, other aspects of his work following traditional practices and forms. It was the development of his ideas

---

[1] Note that the first two operations are exactly the same as those used by Medieval composers

by his pupils and subsequent sympathisers of his methods that lead to the first truly generative processes in music.

Composers such as Anton von Webern (1883-1945), Olivier Messiaen (1902-1998) and Pierre Boulez [Boulez, 1963] extended the method to control parameters such as duration, dynamics and attack. This enabled the almost entirely automatic generation of pieces such as Boulez's *Structures* for two pianos, in which the serial processes determine the overall structure by ordering the different series in a combinatorial manner.

Various other generative processes have subsequently been developed. In contrast to the strict control of Serialism, John Cage [Revill, 1992] used random techniques such as coin tossing and dice rolling - as well as consulting the I Ching. Iannis Xenakis [Xenakis, 1971] generated music using statistical techniques such as Markov Chain analysis. These probabilistic and stochastic methods were used in the earliest composition programs and in fact most of the generative software currently available employs some form of probability.

1.3 Current Approaches to Algorithmic Composition

Statistical techniques The use of random processes in composition has long intrigued composers, but the raw outcome of a purely random process rarely produces results of musical worth. Now that algorithmic composition can be carried out on computers however, randomly generated material can be sorted through testing or embedding of musical constraints within the process. Examples of currently available probabilistic software include the Argentinian composer, Luis Maria Rojas' *Texture* and SSEYO's *Koan*.

Texture is designed to compose aleatory music, and essentially generates random MIDI events. The range of the parameters controlling various aspects such as volume, pitch and density of events can be set by the user, as can the probability distribution within this range. This allows the user to shape the random process to personal specification.

Koan Pro is a parametric composition program. Each piece comprises a number of 'voices'. The principle factor determining which notes each voice will play are given by a set of user defined rules. Separate rules exist for scale, rhythm, harmony etc., and determines the probability of certain events occurring. Rules such as 'play in C major' can be defined, and the probabilities for notes outside of this scale set to 0. In this way, the musical outcome can be constrained by music theory and user preference.

Rule-Based Methods Many other automatic composition applications employ rule-based methods to generate entire pieces. In its most elementary form, this process centers around a series of tests or rules through which the program progresses. For example Schottstaedt's automatic composition program [Schottstaedt, 1989] writes music based on rules from Fux's Gradus ad Parnassum [Fux, 1965] [2]. The program is built around almost seventy-five rules, such as "Parallel fifths are not allowed" and "Avoid

---

[2] An in-depth analysis of fugal structure.

tritones near the cadence in lydian mode.” Such systems are generally employed to replicate existing musical styles.

Neural Networks Neural Networks have also been designed that can be trained to generate music in a specific style. A simple example is the network developed by McAlpine (1999) to mimic a single melody line. It consists of one output node, whose outcome is taken to be the next note in the melody. This output is then fed back as 3 separate inputs with consecutive delay values, enabling the network to calculate the current note based on those played in the 3 previous time steps. The network is trained on a set of composed melodies, and produces one in a similar style.

Genetic Algorithm Genetic Algorithms, especially Interactive GAs (where the user acts as fitness function as in Dawkins) have been used extensively for the composition of music mimicking traditional styles. *GenJam*[Biles, 1994] is a typical example. The GA based system evolves Jazz improvisations. Two populations are maintained, one containing chromosomes representing a series of notes or measures, the other represents a higher level phrase, consisting of four pointers to individuals in the measure population. These measures seem to be pre-programmed abstract contours, and are not assigned specific pitch values until they are played. At this point a look up table is consulted so that the note is constrained to the scale currently played by the phrase section. A fitness value (also stored in each chromosome) is assigned by a human mentor. Variation is driven by a standard cross-over operator, but mutation involves operators such as transposition, inversion, retrograde, sorting, smoothing, and sequencing. The level of embedded musical knowledge means that a population suitable for performance is evolved in just 12 generations.

Iterative Processes Many attempts to generate novel small scale structures employ recursive processes. *FractMus*, developed by Spanish pianist and composer Gustavo Diaz-Jerez uses a selection of algorithms such as Morse-Thue logistic map, Henon maps and the Lorenz equation. Music is generated using up to 16 voices, each associated with a selected algorithm. The user specifies a number of parameters such as scale, initial note and rhythmic figure. The algorithms iterate according to the beat of different voices and at each iteration a new note is picked from a respective scale according to the result of the iteration. This system produces some interesting melodic fragments.

Cellular Automata Given the inherent pattern formation properties of cellular automata (CA), they have been surprisingly under used in the generation of music. One of the few examples is *CAMUS* [McAlpine and Hoggar, 1999]. The basic system generates musical fragments by deriving the intervals of a 3 note chord from the coordinates of live cells in the Game of Life [Gardner, 1970]. The duration and placement of the notes are similarly determined from the states of the neighbouring cells. The system produces some interesting original material, but is capable of little more than note sequences.

Some of these approaches have produced some interesting music, and an understanding and acceptance of the use of mathematical formalisms and computer models in music has accumulated. Recent systems such as Cope's *EMI* [Cope, 1991] (based largely on Markov analysis of a bulk of training examples) has produced 'Mozart symphonies' and 'Chopin Nocturnes' that have almost convinced the critics. The success of these algorithmic composition methods demonstrate that computers can compose if they are programmed correctly. But almost all approaches to date rely on hardwiring, training on a bulk of examples or the embedding of stiff constraints based on traditional music theory to reproduce existing musical styles. Those approaches aimed at creating 'new music' - such as the CA based *CAMUS* and the iterative *FractMus*, have been successful in producing interesting novel fragments, but are a long way from producing music with any kind of large scale form.

1.4 Ascertaining the basis for Music as it 'Could-be'

It has been suggested that part of the difficulty of composing 'new music' by computational methods, is that the inevitable lack of cultural reference makes assessment difficult [Miranda, 1995]. Whilst there exist a plethora of theories and methods of analysis of all musical traditions we have not reached a comprehensive understanding of the fundamental universal properties that would enable an abstract formal analysis of 'new music'. Given the inherently subjective nature of music it seems erudite to take an approach similar to that adopted for the testing of the proverbial pudding, and leave the final judgment to the listeners.

However, in attempting to create a system capable of producing music as it 'could be', it is necessary to examine the low-level components and properties in an attempt to identify those abstract elements that make us perceive arrangements of sound in space as music. It is hoped that this level of enquiry will enable the separation of the fundamental from the stylistic.

If there exists an abstract commonality amongst all musics, it is perhaps the presence of patterns. Patterns in music exist in many dimensions and across multiple time frames. In the absence of any formal abstract theory, one means of ascertaining the fundamental elements of music is consideration of the low level emotional or physiological effects they can induce or the cognitive capacities employed in out appreciation of music. [3]. The inclusion of patterns that are quintessential to music, and the exclusion of those that are specific to particular musical styles forms the basis of the current approach.

Understanding of a complex system can often be advanced by examination of its constituent parts and consideration of how these parts interact to form the whole. Given the fundamentally dynamic nature of music, division according to temporal scale offers a means of decomposition. Time-domains in music have be usefully categorised hierarchically according to our perception of time into timbre, frequency, pulse and form.

---

[3] After all, the human appreciation of sound arrangements beyond communications is unique in the animal kingdom, and the faculties with which we currently appreciate music must serve some previous if not concurrent function.

[Orton, 1990], [Stockhausen, 1989, pp.88].[4]

Timbre. Timbre is a complex entity comprising multiple partials (components characterised by relative amplitudes and frequencies). It has been suggested that timbre is the primary aspect of sound to which we react, and is used cognitively for sound recognition and source localisation. Although manipulation of timbre is a very powerful aspect of music, the present concern is with musical structures at higher levels, and leave this matter to the vast literature on sound synthesis.

Frequency. Sound wave oscillations between 16Hz and 18 kHz are perceived as pitch. At the lowest level the specific pattern of the wave form - the period and amplitude - determine our perception of pitch and volume.

Within Western music, pitches of proportional notes operate in proportional cycles, namely octaves. Octaves are logarithmically separated, a frequency doubling resulting in the perception of the same absolute note an octave higher. For example the A above middle C is standardized as 440Hz, the A an octave above this at 220Hz etc. The scale of equal-temperament, adopted by western tonal music, splits the octave according to whole-number ratios: 2:1 the octave, 3:2 the perfect fifth, 4:3 the perfect fourth, 5:4 the major third etc., producing just 12 notes.

The combination of different pitches produce harmony, and the changes of these relations between pitches over time creates harmonic structure. It is these changes in harmonic structure that evoke emotions such as happiness, sadness, or anticipation and satisfaction. By affecting our emotions in more subtle ways, harmony arguably plays a major role in determining the type of atmosphere music evokes.

In most styles of music, the selection of allowed note frequencies (ie scales), determines the possible harmonies that can arise, and invariably dictates their progression and thus the harmonic structure. In attempting to create an abstract rather than specific style of music, the number of allowed frequencies will be broadened and a micro-tonal approach adopted. The difficulty then lies in finding an algorithm capable of defining notes that are meaningfully related so as to produce an harmonic structure that bears the emotive qualities of music.

Pulse. The domain of pulse covers frequencies of approximately 10Hz to 1/8Hz (8secs). Patterned sounds occurring within this time frame are experienced as rhythm.

Rhythm was arguably the first and is the most fundamental element of music. It has been suggested its repetitive nature has ties to diurnal cycles or circadian rhythms [Simon, 1968]. Perhaps it is no coincidence that we seem to appreciate music with a

---

[4] Stockhausen claimed that the most exciting musical advancements could be made by collapsing these boundaries. Although fixed in acoustic music, electronic techniques mean that the discrete domains become a continuum: events can be time-stretched or compressed such that timbre becomes rhythm or even form, or vice versa. He provides a nice thought experiment comparing the compression of a Beethoven and Schoenberg symphony into a fraction of a second: the aperiodicity of Schoenberg resulting in a crunchy noise rather than a note, the periodicities of the Beethoven producing a vowel like tone.

pulse rate lying in the range of the human heart beat (0.5 - 4Hz, or 240-30 beats per minute), indeed most rhythms fall in this bandwidth.

The term *beat* in music is used to refer to a regular stress or accent. In western music, consistent numbers of beats are grouped, the number and length of these beats being described as 'meter', and signified by the time signature at the start of a written score. Rhythms are then based around beats of these temporal intervals. Although rhythmic patterns are inherently repetitive, particular periods of repetition and note placement relative to the beat are specific to certain styles.

If the musical qualities of rhythm are linked with rhythmic patterns in the natural world, it seems possible that these qualities could be captured by an iterative algorithm without imposing top-down constraints of specific period.

Form. Any musical structures with a period above 8 seconds can be described as form. As with the domains of frequency and pulse, the structures of form are also comprised of patterns or cyclic successions of events. These structures occur at many levels, from harmonic and rhythmic patterns over a number of bars, to the organisation of an the entire musical work[5]. Often the high level patterns are embedded at lower levels, akin to fractal self-similarity. The most simple forms consist of simple repetition: harmonically standard blues music cycles through just three chords; The high-level structure of Baroque music is a ternary form (A-B-A) containing identical repetition of the first section.

Structures and patterns at the level of form enable the listener to make sense of the music. It has been suggested that when listening to sophisticated music, we employ cognitive strategies similar to those used to comprehend life events [Miranda, 2001]. These same processes are evoked when reading literary texts[6]. It is of no surprise then that music and literature seem to share similar underlying structural principles [Jordan and Kafalenos, 1997]. Larivaille [Larivaille, 1974] has proposed a five-stage elementary scheme for narrative :

1. initial equilibrium
2. disturbance
3. reaction
4. consequence.
5. final equilibrium

This scheme is evident in classic texts such as Homer's epic *The Odyssey.* Similar patterns between sections in classical music can also be seen such as in the exposition-development-recapitulation of the Sonata form. In fact the scheme describes the basic balance of novelty and repetition or familiarity that encourages cognitive engagement, enables comprehension through recognition and securing attention through stimulation.

The generation of structure at the level of form represents the biggest challenge for generative music. Even the use of GAs imbued with musical knowledge in the form

---

[5] One of the longest traditional musical forms must be the Omizutori ceremony of Japan, in the Temple of Nara, which lasts three continuous days and nights.

[6] This is certainly true of Western music and literature if not other cultures.

of meaningful encoding and reproduction operators, have failed to create harmonic structures even when provided with existing melody lines [Phon-Amnuaisuk and Wiggins, 1998]. Other EAs used to generate new music have been successful only in creating sequences of notes. Whilst this seems like a hard problem, it seems possible that certain abstract models of complex systems may have low-level properties capable of generating a structured form at a higher level.

The division of music into constituent elements according to time-scale on which they occur provides the basis for the selection of 'fundamental elements' of music and separates the quintessential patterns from those specific to certain styles. Pitch and rhythm are seen here as the building blocks of music, with multiple simultaneous notes producing harmony. Melody, often considered basic to music can be seen simply as the interaction of rhythm (in time) and harmony (in pitch space).

In seeking an fundamental basis for musical patterns, one possibility is to return to the proposed links between music and physiological processes and patterns in the natural world. There exist several algorithms in the field of Adaptive Systems that aim to capture these very processes.

1.5 Cellular Automata as rhythm generators

Previous researchers attempting to use CAs in the generation of new music have concluded that they are more suitable for sound synthesis and are not capable of generating large scale structures [Miranda, 1995]. However, this work (see section 1.3) focused on the production of pitch values from static states. The inherent pattern propagation properties of CAs suggests that if interpreted dynamically, they may in fact provide a *very* suitable basis for the generation of rhythmic structure.

CAs have been used extensively in the modeling of morphogenesis eg [Prusinkiewicz, 1994], as well as more specific pattern formation in living organisms [Markus and Kusch, 1995]. In these models, it is primarily the end-state or states that are of interest. Indeed these ultimate states form the basis of Wolfram's classification [Wolfram, 1984]:

- Class 1: Evolve to homogeneous state.
- Class 2: Evolves to simple separated periodic structures.
- Class 3: Yield chaotic aperiodic patterns.
- Class 4: Yields complex pattern of localized structures.

Although all these classes have potential to produce some interesting musical effects, the periodic and patterned structures of classes 2 and 4 suggest a suitable starting point for the generation of rhythmic patterns.

In addition to the qualitative final states determined by the rule set, CAs can be conceptually divided according to the number of dimensions in which they are represented (typically 1, 2 or 3 for visualisation purposes), as well as the number of possible states of each cell, (2+), and the type of initialisation or seeding (single, random etc.). Multidimensional states and grids hold promise for future investigation, but the simplicity of a binary 1D CA is perhaps conceptually closest to our perception of musical rhythm. A 1D CA of Class 4 will be used to derive rhythmic patterns.

1.6 Homeostasis

> "The fixity of the internal environment is the condition for free life" [Bernard, 1870]

Rather than imposing high level rules regarding the structure of the musical form, the desired approach is the development of a system that generates music at the level of individual notes. Further, that there is some relation between these notes such that a higher level harmonic structure is created in which the whole is somehow 'balanced' musically.

Perhaps the most ubiquitous process of balance, or maintenance of constancy is homeostasis[7]. Homeostasis is the central unify concept in modern physiology, and is also the root of our current understanding of feedback, and so of Cybernetics .

1.6.1 Biological Homeostasis

The capacity to maintain homeostasis is seen as a fundamental of life: the gradual decline in homeostatic potential being a key characteristic of the onset of the aging process [Goya and Bolagnani, 1999].

Homeostasis typically depends on the action and interaction of a number of body systems to maintain a range of conditions within which the body can operate most efficiently. The larger functional unit of the homeostatic network is made up of often anatomically discrete systems: the idea of an immune-neuroendocrine homeostatic network in higher mammals has recently been proposed. [Goya and Bolagnani, 1999]

Homeostatic control is typically achieved by negative feedback, for example blood levels of Thyroid Stimulating Hormone (TSH) serve as feedback for the production of TSH. Positive feedback is also used: during uterine contractions, oxytocin is produced which triggers an increase in frequency and strength of uterine contractions, producing more oxytocin etc.

In many cases, the critical levels being maintained, and even the source of triggering stimuli are regulated and adapted by the network. For example, the ventilation system regulates $C0_2$, and pH as well as O_2 levels (complementary pH regulation occurring in the kidney system via control of bicarbonate retention.). Chemoreceptors in the carotid and aortic bodies provide information to the integrating centers of the brainstem regarding peripheral levels of O_2, CO_2 and pH level. Changes in arterial $C0_2$ translate into changes in cerebral spinal fluid pH. The central chemoreceptors respond to changes

---

[7] Bernard was the first to recognise the importance of constancy in the internal environment of multicellular organism. The concept was developed by John Barcroft (England, 1872-1947) and W.B. Cannon (America, 1871-1945) who coined the term 'homeostasis'.

in a drop in pH by increasing the breathing rate. In healthy persons, this mechanism maintains $C0_2$ and pH within critical levels. However, when the lung performance is below optimal (ie diseased or damaged), diffusion efficiency is reduced. This causes an immediate increase in breathing rate due to 0_2 deficiency, causing and a drop in pH. In this state pH and $C0_2$ levels no longer serve as reliable triggering stimuli, due to the changes in diffusion efficiency. Therefore O_2 levels are used as a trigger stimuli, and the kidney system up-regulates, changing bicarbonate retention to regulate pH levels. [8]

In abstract terms this process, based on interaction and multiple feedback suggests a mechanism by which individual notes could be 'meaningfully' related, the process of feedback between parts being paralleled with harmonically resonant relations in pitch. However the complexity of the physiological reality of this process is daunting.

1.6.2 Ashby's Homeostat

An abstract formulation of the homeostatic process was developed by Cybernetician Ross Ashby [Ashby, 1956] to explain the origins of adaptation in living systems. His central concern was the reconciliation of adaptive and spontaneous behaviour in the state determined system. He presents his solution in the form of an abstract description of an 'ultrastable system':

> *Two systems of continuous variables (that we call 'environment' and 'reacting part', interact, so that a primary feedback (through complex sensory and motor channels) exists between them. Another feedback, working intermittently and at a much slower order of speed, goes from the environment to certain continuous variables which in their turn affect some step-mechanisms, the effect being that the step-mechanisms change value when and only when these variables pass outside given limits. The step mechanisms affect the reacting part; by acting as parameters to it they determine how it shall react to the environment.* [Ashby, 1956, 7/26]

In order to clarify this abstract statement, Ashby built an electro-mechanical device which displayed the fundamental properties of homeostasis: the maintenance of variables within certain limits in the face of external perturbation. His abstract presentation of the homeostatic process provides the basis of the current system for the generation of musical harmony.

Ashby's 'homeostat' consisted of 4 interconnected units, each received input from every other, as well as itself. The output of each was proportional to the sum of the (weighted) inputs, and was maintained within certain limits. If any output transgressed its limit, changes were made in the connection strengths of the connected units until stability was regained.[9] The basic homeostatic behaviour is essentially achieved through feedback and 'trial and error' of connection strengths.

Despite displaying the basic characteristics of homeostasis (as well as the ability for training [Ashby, 1956, 8\9] and see section 2.2.1), the simple 4 unit system is a far cry

---

[8] Don't give oxygen to an emphysema sufferer!

[9] Ashby's original description of the machine is presented in appendix A

from the complexity of physiological homeostatic systems. The obvious discrepancy in size was noted by Ashby, and he later presented findings regarding the effect of both size and connectivity on stability [Ashby, 1970]. In *Design for a Brain*, larger networks comprising sparsely connected parts are also considered, bringing the system more into line with current modular conceptions of neural and physiological architecture.

Self-regulation However, as discussed above, in natural systems homeostatic networks are self-regulating. In Ashby's model, the value of the critical boundary is set by hand, the viscosity of the liquid in the troughs which determines the degree to which the system is state determined is similarly hand-set. Self-regulating systems of this type have not received as much attention within the field of Adaptive Systems as their ubiquity in nature dictates. In this project then, possibilites and properties of self-regulating networks will be investigated, and deployed as the basis for generating musical harmonies.

Chapter 2

Methods 1 : Low level Algorithms

2.1 Pilot Work: exploring methods of audio generation

In order to gain a feeling for the effects of simple processes in computer composition, a number of simple rules based on stochastic selection of values were explored.

2.1.1 A stochastic method for Rhythm production

A simple stochastic method was developed that utilises random selection whilst preserving a sense of melody or rhythm. By exploiting the fundamental periodicity of music, and pattern recognition biases in human perception [Meyer, 1956], randomly selected timings or note values can be made to sound rhythmic by simply looping the selection.

N numbers $n_0, n_1 ... n_n$ are randomly selected. N defines the 'density', and can be any number, but somewhere between 3 and 10 works well. The value of n determines the frequency of each note ie it will be played every n beats. For example if three numbers 3, 5 and 9 are selected, notes will be played on beats 3, 5, 6, 9, 10, 12, 15 etc. The period of the whole loop is therefore the lowest common multiple of the entire set. In this case the loop will repeat every 45 beats. Where one value of n is a multiple of another, or a common multiple is shared by a subset, an accent is created, or if different pitch values are used, a chord.

A very simple example using pitched notes can be heard on track 1 of the cd. A transcript of the loop is shown in fig 2.1. In this example, the selected numbers are 4, 4, 2, and 11. These are played on the notes G♯, E, A and G respectively with the semiquaver as the smallest unit of time[1]. For readers unaquainted with musical notation, the figure can be viewed as a graph of pitch space against time, a dot denoting the occurence of an event in the respective dimensions. It is sufficient to notice that the combination of 2 fours and a 2 gives a regular repetition of alternately accented beats (ie. more simultaneous events occur in pitch space when the periods coincide), the addition of one larger odd number changes the pattern from period 4 to period 44, and occurs on unaccented beats, creating a more interesting syncopated rhythm. This demonstrates

---

[1] The regular 2 and 4 beat notes are shown as quavers for clarity.

Figure 2.1: Transcript of track 1: a simple random loop

how non-trivial musical effects of accent and syncopation arise from the ratios between the periods of notes.

2.1.2 Stochastic harmony generation

This loop generation method was combined with another simple stochastic algorithm for changing the pitch. Here the loop was made up of 7 randomly selected numbers of different pitch values and played in conjunction with a simple 'bass line' consisting of just two alternating notes. The pitch changes of both parts are controlled by another set of random numbers: Every b beats, the pitch of the loop notes are shifted by an amount bl. Every c times the loop pitches change, the pitches of each of the notes in the bass part are similarly shifted by a random amount cl, where b and c are in the range $(10 : 30)$ and bl and cl are in the range $(-6 : 6)$

Tracks 2 and 3 are examples of this algorithm. In track 2, the random loop is played using short notes (plucked strings), producing a rhythmic effect. In track 3 the notes are sustained, producing dense chords. This demonstrates the enormous effect that changes in instrumentation can have.[2].

These audio examples show that interesting musical fragments, and even harmonic changes can be produced with very simple stochastic algorithms. However, inevitably, such processes rely upon chance, and their percentage success rate in producing enjoyable, or even interesting music is very low. When using such stochastic processes, the only possibilities are to run the processes many times and make a selection (using a GA, KBS or probabilistic methods), or constrain the selection as here. Thus although these examples demonstrate the potential for algorithmic composition, they also reinforce the need to examine other generative processes that need neither constraining nor training.

---

[2] The point is made here, but the main focus of this project is the generation of structure rather than effect. Explorations of instrumentation and timbre are left for future work...

2.2 Investigating Homeostatic Systems

Prior to the development of any audio systems, a simple simulation of Ashby's homeostat [Ashby, 1956] was implemented and thoroughly tested to ensure its functional equivalence. This basic system was then developed into a network and possibilities of self-regulation of critical parameters were explored. This section describes the development and testing of several different homeostatic systems.

2.2.1 A simple Homeostat

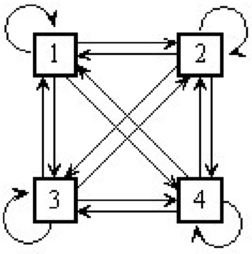

Figure 2.2: Diagram of connections for a fully connected 4 unit homeostat.

A homeostat is conceived as a system comprising N units each connected to M other units, including a recurrent connection (see fig. 2.2 for the diagram of connections for a fully-connected 4 unit homeostat ($N = M = 4$)) Thus for a fully connected network, the input to unit N at time t is the weighted output of units $1-_{(N-1)}$ at time $t-1$, plus its own (weighted) output at $t-1$. The output of each unit is proportional to the sum of these weighted inputs (see equation 2.1). These weights are initially randomised, and are re-randomised whenever the output of the unit exceeds a prespecified 'critical value'.

In simulation, these relations inevitably operate in discrete, rather than continuous time, affecting the simulation of certain physical parameters (see section **viscosity** below). The commutators and potentiometers described by Ashby [Ashby, 1956, 8\2] which controlled the amount fraction and polarity of the passing current, are modelled as weights and their signs.

Algorithm and Assumptions

Critical Constants Whilst developing the simulation, two constants were initially included in the output calculation as Ashby writes: "the torque on the magnet is *approximately* equal to the algebraic sum of the currents in A, B, C". [Ashby, 1956, (8\2)]. Similarly, the output is described as "approximately proportional to the deviation from zero". (ibid.) In the physical homeostat, these would have been determined by properties of the materials, and although constant, their values unknown. In this simulation however, experimentation proved that once all other variables were set to suitable values, these two constants were redundant.

Uniselector Action In Ashby's machine, the frequency of uniselector action (ie weight changes) could be varied, and was held at three seconds. The size of this interval did not seem to critically affect the performance of the simulated homeostat and is held at 1, allowing weights change at every iteration.

Viscosity In the original physical machine, the degree to which the system was state determined was controlled by the viscosity of the liquid in the troughs in which the outputs trailed. This damping effect was modeled by restricting the variation in outputs in any one unit from one iteration to the next.[3]

The effect of changing the value of this variable proved similar to the assumed effect of varying the viscosity of a liquid: low values (representing high viscosity) result in turgid, stable behaviour; high values produce more exploratory 'run-away' behaviour as each unit does not have time to achieve stable parameter settings before other units transgress the critical limits (see section 2.2.1 and figure 2.4 .

Main Algorithm and Procedure The main output algorithm used was as follows:

$$\mathrm{O}_{i(t+1)} = \left[\sum_{j=0}^{j} \mathrm{input}_{ij(t)} \times \mathrm{weights}_{ij(t)} \times \mathrm{On}_{i,j} \right] - \mathrm{Target} \tag{2.1}$$

where

$$\mathrm{input}_{ij}(t) = \left[\sum_{j=0}^{j-1} \mathrm{output}(t-1) \right] + \mathrm{output}_i(t-1)$$

$\mathrm{O}_{i(t+1)}$ is the Output of the i_{th} unit at time $_{t+1}$
Target = 0.0 for all units.
The 'On' array is used to implement different levels of connectivity thus the *ith* element is always set to 1 for the *ith* unit and the remaining connections are randomised according to the specified percentage connectivity.

Procedure. At the start of each run, all weights and outputs are initialised. Weights are selected randomly from the range -1:1, outputs are initially set within the critical limit (-0.05 : 0.05 when constant) to allow the weights to stabalise. At each iteration, the output (ie deviation from target value) of each unit is checked: if outside the critical range, weights are randomised to all connected units. The recurrent connection weight remain constant throughout each run.[4].

---

[3] This is in keeping with the fact that in simulation the activity of the homeostat is discrete rather than continuous.

[4] Ashby's ' homeostat did not have a uniselector on the recurrent connection

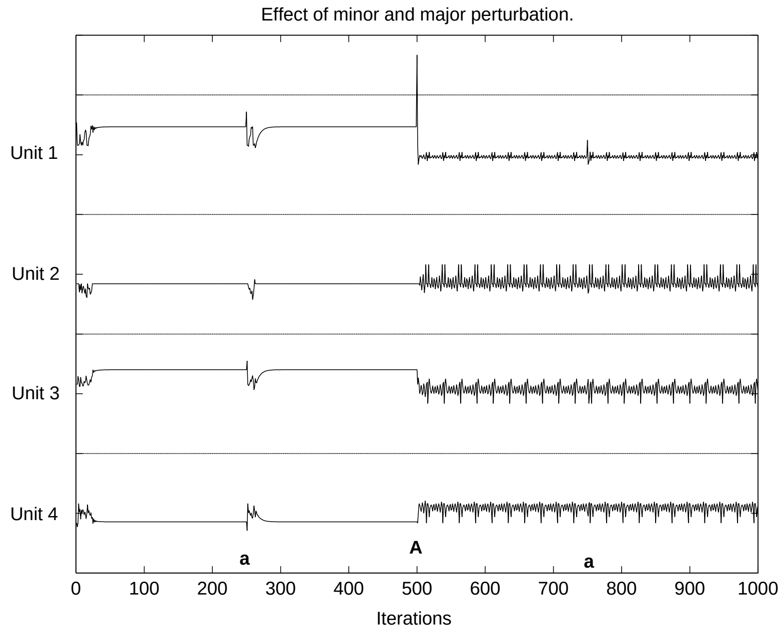

Figure 2.3: Outputs of a 4 unit homeostat demonstrating stability to minor perturbation and re-stability after critical interference.

Ascertaining Homeostatic Properties

In order to ensure the simulation mimicked Ashby's machine functionally, several of the experiments he reports were carried out.

Essential Characteristic Properties As shown in figure 2.3, the system - once stable - exhibits robustness against minor perturbation. At point **a** (iteration 250), the output of unit 1 was manually increased by a fraction of the critical deviation. A transient effect can be seen in the other 3 units, after which all outputs return to their original value. At point **A**, the output of unit 1 was forced *outside* its critical limit. This causes weights changes, and the system enters a different stable state. The example illustrates the existence of converged as well as oscillatory stable states.

- This simple 4 unit homeostat is used in the AdSyM Interactive Loop Generator, where this oscillatory behaviour produces interesting harmonic loops (see section 3.2).

Factors affecting Stability As mentioned above, the value of the *maxchange* variable, used to simulate viscosity, critically affects the behaviour of the homeostat. In accordance with Ashby's reports, stability is also affected by the number of units in the system and the degree of connectivity. Figure 2.4 shows the effect of increasing

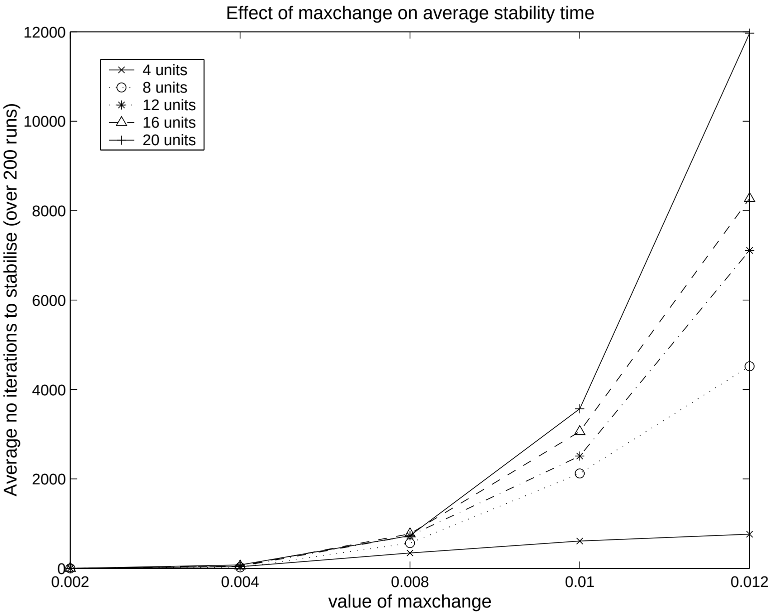

Figure 2.4: Change in stability as a function of number of units and value of maxchange

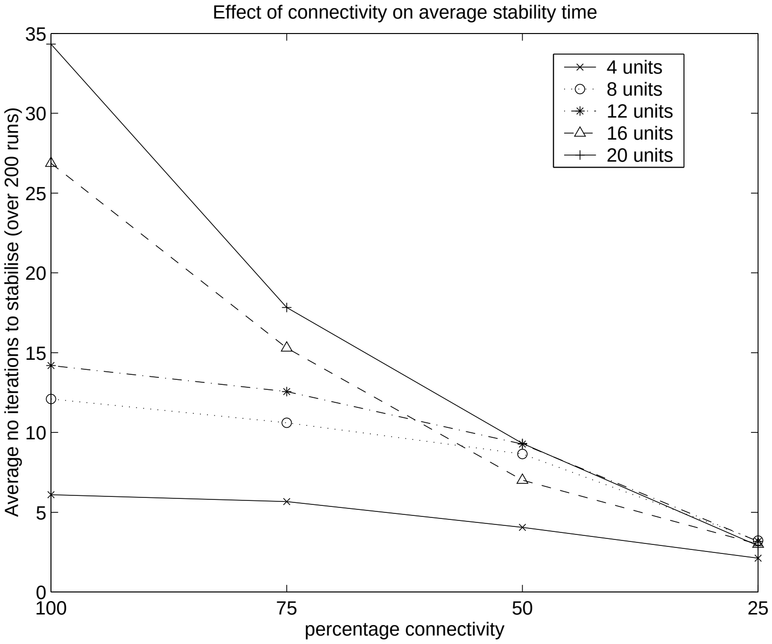

Figure 2.5: Change in stability as a function of connectivity

maxchange for different numbers of units and demonstrates the inverse relationship of stability with both these factors. Figure 2.5 shows a similar inverse relationship between stability and degree of connectivity.

In all experiments, stability is taken as the point at which all units remain inside limits (and therefore weights remain constant), and measured as the number of iterations taken to achieve this state, averaged over 200 runs.

'Training' the Homeostat Ashby also demonstrates the possibility of simple 'training' in a modified network. If three units are unidirectionally (rather than fully) connected, the (polarity of the) response of unit 2 to an induced movement in unit 1 can be controlled by forcing the output of the third unit outside limits until suitable connection weights are established. Figure 2.6 shows the connectivity (or 'diagram of immediate effects' in Ashbian terms) for this system (bold arrows), including influence of the user (dashed arrows).

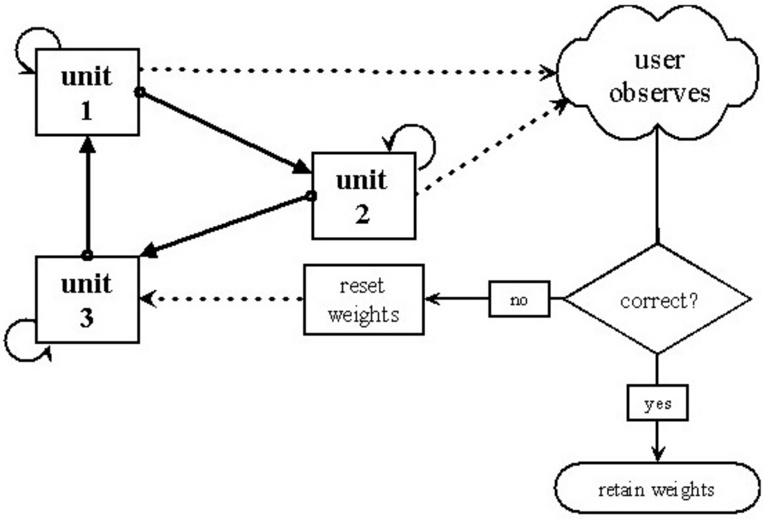

Figure 2.6: Diagram of immediate effects for a 3 unit unidirectionally connected homeostat including conditional user interference.

A graph of the outputs over time under these conditions is shown in figure 2.7. At points **a**, the output of unit 1 is minimally perturbed. The desired response is that unit 2 reacts in the opposite direction to unit 1. On the first trial (iteration 250), unit 2 moves similarly upwards. As this is the 'incorrect' response, weight changes are induced by forcing the output of unit 3 outside the critical limits at **A**. These changes are sufficient to produce the desired, opposite, response of unit 2 to the second trial at iteration 810.

This is effectively learning by trial and error. Although the number of trials required to change an association (same or opposite movement) is indeterminate, once the settings are found, the system is robust and will continue to display the appropriate behaviour.

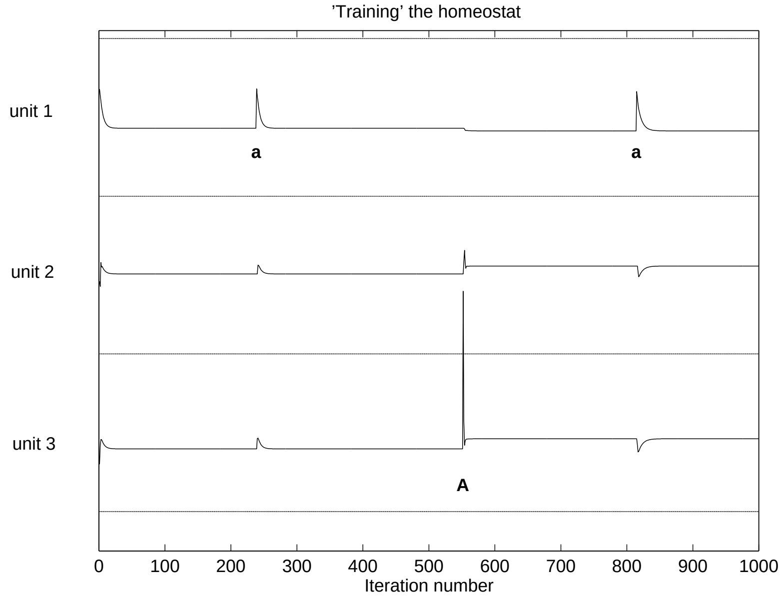

Figure 2.7: Outputs of a 3 unit unidirectionally connected homeostat demonstrating 'training'

- This functional property is used as the principal mechanism in the AdSyM Adaptive Sequencer Module.

2.2.2 A Self-regulating Homeostatic Network

Although the simple 4 unit homeostat exhibits some impressive properties, such simple systems are rare in nature. In an attempt to closer approximate the dynamics of living systems, and to produce behaviour suitable for the generation of musical structures, a small modular network was constructed. The network can be conceived as two 4 unit homeostats and a 2 unit homeostat[5] (see fig. 2.8). Each are fully *intra*connected but the *inter*connections can be varied randomly according to a user specified percentage connectivity.

Various methods of network self-regulation were explored. The two most influential parameters that were fixed (by hand or circumstance) in Ashby's model were the critical deviation and viscosity. Various methods of direct and proportional control of these parameters by the outputs of certain units were investigated.

Defining Control Initially, the two parameters were simply set to the output values of two different units. This was rejected as it meant that as soon as *maxchange* was

---

[5] The number of parts and their sizes were selected arbitrarily: the effects of these variables and their ratios deserves investigation.

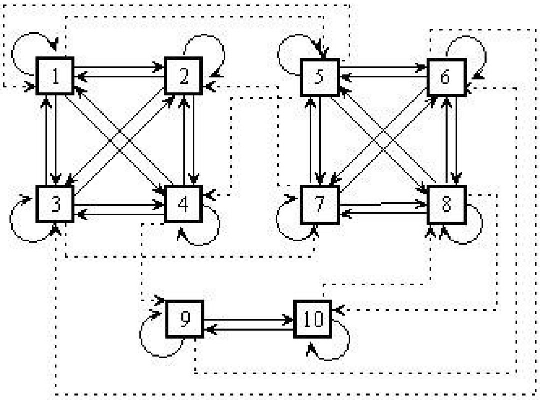

Figure 2.8: Schematic Diagram of Network showing Full interconnections and 10% interconnectivity

set to a low value, the system froze, with many units stuck on this value. Aesthetically, this produced uninteresting behaviour. More seriously if the unit controlling the value of the critical deviation was stuck on the same or lower value, all other units at or above this value would be out of limits. Thus continual weight changes occur, but the system cannot move to self-correct. Homeostasis cannot occur.

After further experimentation, self-regulation of the critical deviation was rejected, and proportional control of just the *maxchange* parameter was implemented using equation 2.2. The equation was derived to invoke a relationship, in discrete time, analogous to the differential equation governing the damping effect of viscosity on a moving body.

$$\text{maxchange}_{(t+1)} = \frac{\text{output unit}_{i(t)} + \text{critical deviation}}{2} \tag{2.2}$$

To ensure that the essential qualities were preserved in the modified network, resistance to perturbation and effect of connectivity on stability were examined. Figure 2.9 shows that the system returns to the same stable state after minor perturbation (marked **a** at iterations 250 and 525), and re-stabilises following critical perturbation (marked **A** at iteration 350). As we would expect, the increase in size means the network takes longer to stabilise.

Figure 2.10 shows the increase in the average time to stabilise with an increase in percentage connectivity, demonstrating the same inverse relationship between connectivity and stability that is observed in the simple homeostat. Here connectivity refers to the degree of *inter*connectivity (connections between distinct modules), each module being fully *intra*connected.

- This auto-regulated network provides the harmonic structure for the AdSyM Autonomous Homeostatic Music System.

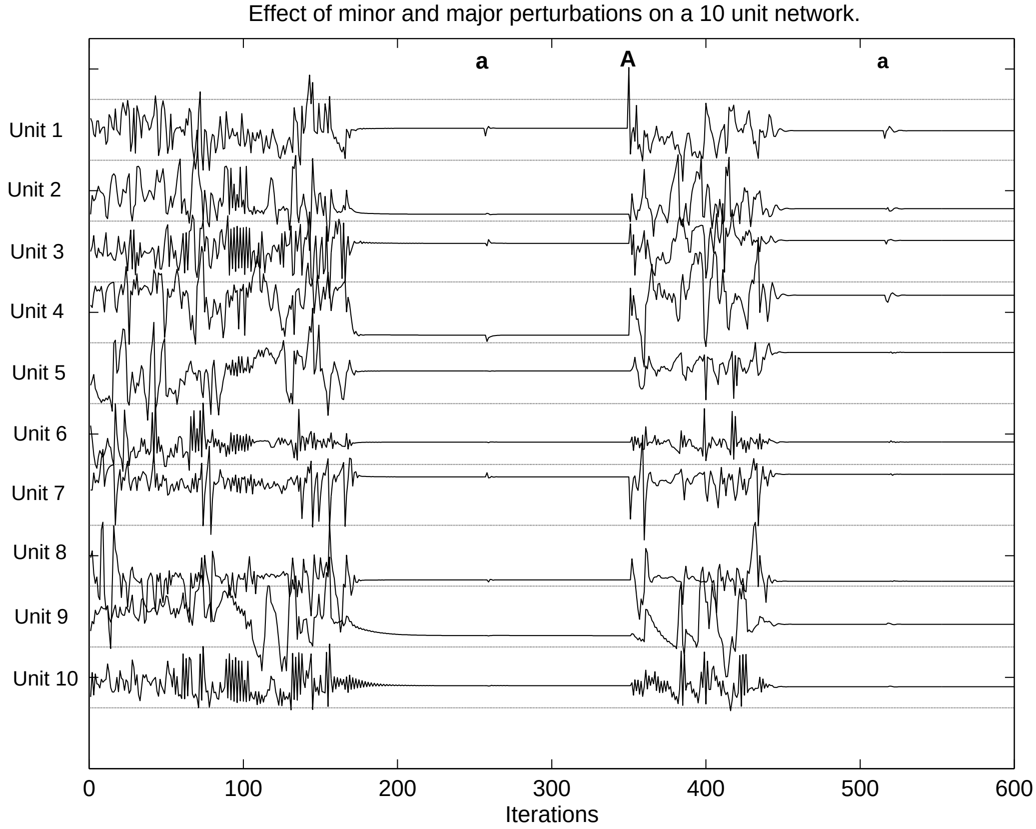

Figure 2.9: Outputs of a 10 unit self-regulated homeostatic network demonstrating stability to minor perturbation and re-stability after critical interference

2.3 Cellular Automata

In both music generation systems, the *harmonic* structure is generated using versions of the homeostatic systems described above. The *rhythmic* structure in both is defined by a cellular automata (CA). Both systems employ the same centrally seeded 1D binary CA[6], depicted graphically in figure 2.11. This CA was selected as its emergent pattern bears a good mixture of repetition, novelty and variation. When seeded centrally (so that only a few cells in the center are alive), the pattern extends vertically at each iteration. At each time step, there is a change in the proportion of cell states (alive vs dead). These binary states are interpreted musically as 'play' and 'rest', so the triangular shape produces an increasingly dense rhythm.

---

[6] This particular CA has been classified as a 1D totalistic CA. http://www.mirwoj.opus.chelm.pl./ca/index.html

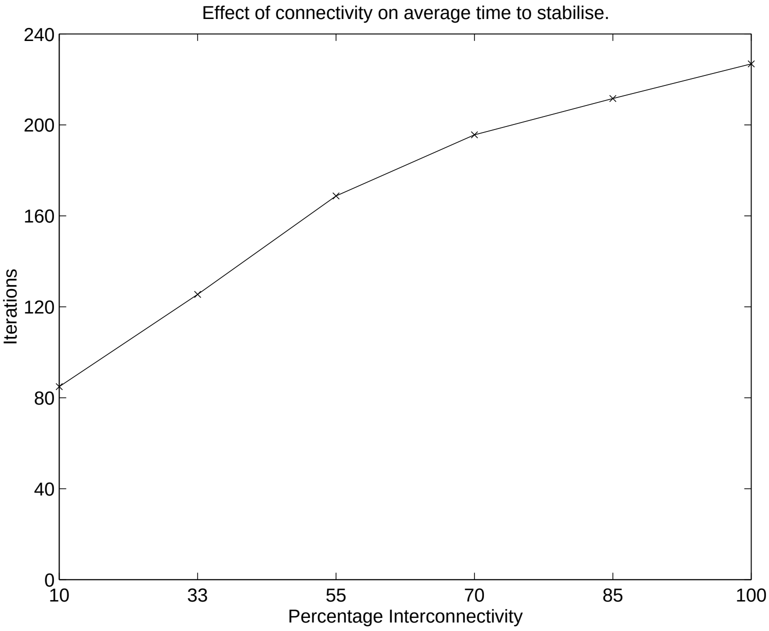

Figure 2.10: Change in stability as a function of interconnectivity for an auto-regulated network

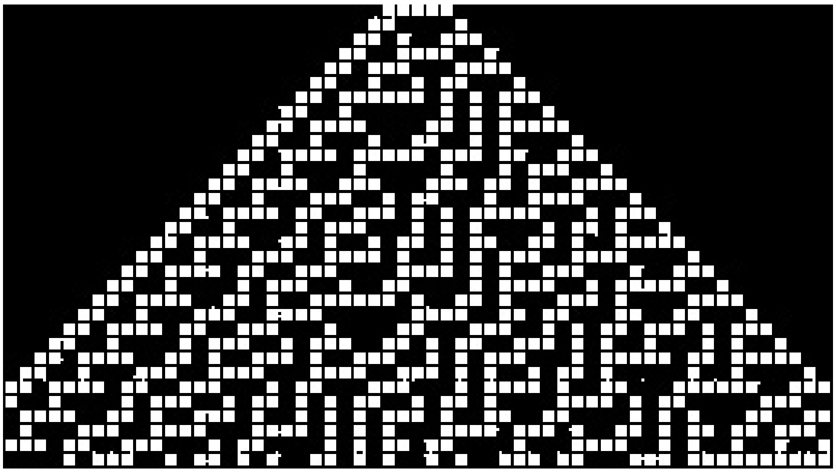

Figure 2.11: Graphical representation of totalistic CA seeded centrally

Production rules A 1D CA can be conceived as successive rows of cells that are updated synchronously at each iteration. The states of a particular cell at time t is determined by the states of the cells in its neighbourhood at time $t-1$, geometrically, those immediately above. The 'neighbourhood' describes the degree of influence to the left and right, thus for a neighbourhood of 1, a cell will be influenced by three cells: the one immediately above it and the two either side of this. The production rules for the CA depicted in figure 2.11 (which has a neighbourhood of 1) are given below [7].

All possible states:	111	110	101	100	011	010	001	000
specific rule:	0	1	1	0	1	1	1	0

- The pattern propagated by these CA production rules is used to generate rhythm in both generative music systems.

---

[7] These rules can be expressed more succinctly using Wolfram's hexadecimal notation: the above rule (01101110 in decimal form) can be written '$6E$'.

Chapter 3

Methods 2: Music Systems

3.1 Outline of the systems

Three different homeostatic systems were built.

AdSyM I: Interactive Homeostatic Loop Generator

- A simple 4 unit homeostat generates harmonic loops, accompanied by a CA-defined percussion part.
- Aimed at producing interesting harmonic loops that can be saved and used as samples.
- Enables user control of: tempo, volumes of separate parts, degree of pitch bend, harmonic root, degree of stability (through changing viscosity) and minor and critical interference of homeostat output.

AdSyM II: Autonomous Homeostatic Music System

- A self-regulated homeostatic network produces harmonies and controls the dynamics of individual instrumental parts as well as transposition of the chords it generates. A CA defines the rhythmic part, and the timing of the melody is generated using a simple stochastic loop.
- Aimed at generating longer pieces of music.
- Designed to run autonomously.

AdSyM III: Trainable Sequencer Module (prototype)

- A modified 3 unit homeostat can be 'trained' by the user to associate different sets of preset loops.
- Produced as a prototype (proof of concept) for next-generation digital music production software.

3.2 AdSyM I: Interactive Homeostatic Loop Generator

The Interactive Homeostatic loop generator was designed to enable user-controlled production of harmonic loops. A screen shot of the GUI is given in appendix B, figure B.1. The harmonies are produced by 4 sustained notes, the pitch of each being controlled by the outputs of a fully connected 4 unit homeostat. These notes are also used to create a 'melody' line, the timing of which is determined according to the stochastic method described in section 2.1.1. The rhythm is played by three percussion parts based on the the CA described above. Each part can be played independently, or in unison. When the rhythm part is turned on, there is an 'auto-perturb' option which triggers chord changes when the rhythm repeats.

3.2.1 Homeostatic Harmony

Pitched notes are produced by transforming the outputs of each of the units of the homeostat described in section 2.2.1 into a MIDI signal. This operates dynamically. Because of the desire to use note values other than those of the even-tempered scale, micro-tonality was implemented using MIDI pitch bend. This decrements or increments the pitch of a given note by a fraction of a semi-tone (see section 4.1). An output to pitch bend mapping was devised such that the min-max critical deviation of the output values produces a difference of 1 whole tone at base level. This is shown in equation 3.1.

$$\text{pitch bend}_i = (\text{output}_i \times \frac{\text{pb intervals}}{\text{critical deviation}} \times \text{Pb Factor}) + 64; (3.1)$$

where
output_i is the output of the i_{th} unit,
pb intervals is the number of intervals in a semi-tone (32),
Pb Factor is the user defined multiplying factor.
critical deviation $= 0.05$
64 is added as the MIDI pitch bend command is centralised, so that values $0-63$ cause a decrease the pitch and values $65 - 127$ an increase (see section 4.1).

3.2.2 Melody line

The note timings of the melody line are produced using the method described in section 2.1.1 using just 4 random numbers. These are selected at initialisation and remain constant throughout. Each number is associated with one of the outputs of the homeostat, playing whatever note is defined at that time step according to equation 3.1.

When the outputs are stable and all converged on single values, the melody is made up of just 4 notes. During unstable or oscillatory periods however, more complex melodies are created, as the period of oscillation rarely coincides with value of the random number with which it is associated.

3.2.3 Rhythmic Percussion part

The CA described in section 2.3 (shown in figure 2.11), is used to produce a rhythm, played by 3 different percussion instruments. The rule set given in fig 2.3 is used to generate the CA pattern in a 13 x 22 grid[1]. This is done in order to preserve the changes in rhythmic density (arising from the triangular shape). The states of individual cells in the CA are interpreted very simply as a rhythmic score: on = play, off = rest. Each instrument then 'reads' along each successive row, the last cell in row N being followed by the first cell in row 1. This produces a long rhythmic loop. The changes in state of cells at discrete time steps is represented as the 3 instruments play consecutive rows, thus at any one time, the future, present and past states of any cell are simultaneously presented on different instruments.

3.2.4 Procedure

When the GUI form is created, the homeostat algorithm is initialised, MIDI messages are sent to set the necessary channels to the relevant patches (synth instruments), the random values defining the melody timings are assigned and the structure of the CA is generated and stored in a 2 dimensional array. The central 'target' value of each unit is set to a particular pitch value (here the chord of C, but the system is surprisingly insensitive), from which it deviates according to the pitch bend amount determined by the outputs.

At each iteration, outputs of each homeostat unit are calculated and transformed into pitch according to equation 3.1. These are all played over the 4 'harmony' channels as continuous notes. Any number of these notes may also be played over one of the 4 'melody' channels if so defined. If the CA cell currently defining any of the percussion instruments is on, the relevant percussion part is also played. The time ratios of these 3 parts can be adjusted to allow changes in relative tempi of each part.

Each iteration is separated by a user defined time step (measured in ms). At the end of each iteration, checks are made for changes in user controls which are implemented in the next iteration. The system continues until the user presses 'stop'.

User Controls The system provides user control of the most influential parameters, both of the homeostat and the audio realisation.

- **Volume control:** The volumes of the melody and rhythm parts can be independently varied using sliders. Due to MIDI restrictions, the harmony can only be set on or off. The other two parts can also be turned on or off independently.
- **Tempo:** The tempo can be set in an edit box, and is defined as the time interval (in ms) between each iteration of the homeostat. This can take any value, although the percussion part loses definition at values below about 160 (this is used to create an interesting effect in track 5).

---

[1] Row length was set to the prime no. 13 to avoid any factors which may constrain the meter. The number of rows is essentially arbitrary.

- **Harmonic Root:** The chord on which the homeostat is seeded can be transposed up or down (semi-tones).
- **PB Factor:** The output-pitch bend mapping can be increased by a factor (always relative to 1 rather than cumulative) entered in an edit-box. This means that a given change in output value will cause a greater difference in pitch.
- **Homeostat Interference:** Three buttons on the GUI allow the user to change the output value of unit 1 by a preset amount: *min up* and *min down* in/decrease the output by 1/4 of the critical deviation; *Max Dev* increases the output by twice the critical deviation. These changes are effected in the following iteration.
- **Viscosity:** The user can also change the value of the 'max change' variable, affecting the stability of the system. Low values produce repetitive and monotonous (in the truest sense of the word) harmonies, higher values more random loops.
- **Auto-perturb:** When the rhythm part is set to 'on', this causes an automatic increase of unit 1 output beyond critical limits, inducing a change in the harmony at each repetition of the percussion part. This enables the system to be used autonomously as well as interactively.

Any loops the user likes can be recorded straight into a digital sound editor.

3.3 AdSyM II Autonomous Homeostatic Music Generator

The Autonomous is essentially a development of the Loop generator in automatic mode. A screen shot of the GUI is given in appendix B, figure B.2. Here the GUI acts primarily as a display rather than controller.

Rather than a simple 4 unit homeostat, the central component is the self-regulated homeostatic network described in section 2.2.2. Outputs of this network are used to control not only the pitch of the 4 chords, but the volumes of the melody and harmony parts, the transposition of the chord, and the value of maxchange (as described above). Although the user can still perturb the outputs at any point to induce a harmony change, this is done automatically whenever the homeostat stabilises. The time step is also automatically controlled.

The melody line is produced in exactly the same way as in version 1. However, the increase in complexity of the homeostatic network, means that it takes longer to stabilise, producing more interesting melody lines. The rhythm is similarly defined by the states of the same CA, but here is played by synth 'plucked strings' - essentially a pitched percussive sound.

The outputs of units 1-4 control the pitch bend amount as in equation 3.1. In addition, units 9 and 10 determine the volume of the melody and rhythm parts respectively, and unit 8 controls the shift in root value of base of the chord (transposition). Output unit 7 controls the value of maxchange as described in 2.2.2).

These updates are made at the repetition of the rhythmic loop. (this point does not trigger a homeostat interference as above). At the end of the CA loop, the output values of units 8,9 and 10 are converted to a new transposition value and volume settings for the instrumental parts according to equations 3.2 - 3.4. The change in root values affect only the sustained chords and not the melody and rhythm parts.

$$\text{Root value} = \text{output}_{unit_8} \text{ x } rc \tag{3.2}$$

$$\text{Melody Volume} = (\text{output}_{unit_9} + \text{ critical deviation }) \text{ x } mvc \tag{3.3}$$

$$\text{Rhythm volume} = (\text{output}_{unit_10} \text{ critical deviation }) \text{ x } rvc \tag{3.4}$$

where
$rc = 140$.
This value is not critical, but as the outputs will generally be in the range $(-0.05 : 0.05)$, this converts the shift amount to $(-7 : 7)$ semi-tones.
mvc = rvc = 1270.
The addition of a constant equal to the critical range, multiplied by this factor produces values within the MIDI velocity range of 0-127. Values outside this range (ie when outputs are beyond the critical limits) are curtailed.

As well as periodic updates of volume, gradual changes (crescendi and diminuendi) are implemented according to the update value as follows:

```
if (current velocity $\geq \frac{1}{2}$ top velocity )
then (reduce velocity each iteration)
else
(increase velocity each iteration)
```

In addition, a gradual accelerando is implemented by reducing the time between iterations logarithmically (ie reducing the time step by 1 each iteration) to a minimum value of $180ms$. This reset to its original value ($300ms$) when the homeostat stabilises.

When stability is achieved, a critical perturbation of the output of unit 1 is induced, causing randomisation of connections and creating a new harmonic path. Because of the processing restrictions of the real time MIDI production, a 'cheap and dirty' algorithm is used. The system tends to oscillate as often as it converges in its stable state, but for the low values of interconnectivity used here (10%), this oscillation is typically between only two values. Thus rather than keeping a record of all outputs to date, and checking all outputs, only the last two outputs of unit 1 are retained and used to verify stability. The system is considered to have stabilised if outputs of two consecutive iterations, or pairs of outputs over the last 4 iterations differ by less than (0.00005).[2]

---

[2] This may seem a little *too* cheap, but because each unit has the power of veto over all others, even in

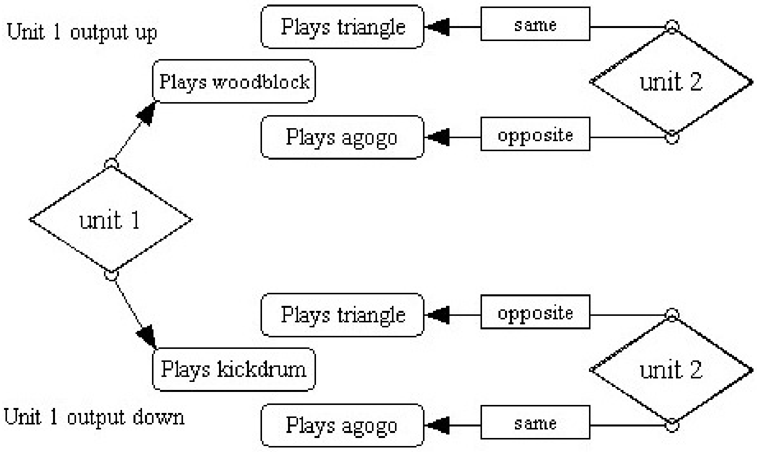

Figure 3.1: Schematic diagram of sample associations in the Adaptive Sequencer

Using the outputs of the units to control the musical parameters effectively normalises their distribution: outputs are most frequently within the limits of the critical deviation, but occasionally lie at extremes. Thus extreme behaviour in the system is reflected in the music produced.

3.4 AdSyM III Adaptive Sequencer Module

In both systems described above, the homeostatic *process* is deployed to create musical structure. This system utilises a *functional* properties of the homeostat, namely the capacity for 'training' described in 2.2.1. A screen shot of the GUI is given in appendix B, figure B.3.

Whilst there is a range of digital sequencers available, the vast majority are essentially models of analogue multi-track recorders. There is not currently anything on the market that can be updated dynamically without having to reassign channels. The ability for 'learning' in the homeostat suggests its use as a mechanism in the next generation of more adaptive music software.

This implementation is developed as a prototype. The system allows the user to alter the association between two sets of preset loops.

Procedure The homeostat is configured as described in section 2.2.1. Rather than generating audio from the values of the outputs, their relative behaviour is used to trigger

---

a sparsely connected system, referencing only one unit is sufficient. In a more rigorous evaluation of the system, identity, rather than rough equality between outputs would be necessary. However, generally, outputs only converge to this extent when stability is imminent. Change become logarithmically smaller, thus waiting for the system, to converge entirely results in aesthetically boring behaviour. In trials, the algorithm never made any incorrect classifications ie. outputs of *all* units were always within limits and close to stability).

preset samples. Just as in the experiment described, the output of unit 1 is perturbed, and the relative behaviour (similar or opposite movement) of unit 2 monitored. Here however, the direction of perturbation triggers one of two loops and the direction of reaction of unit 2 triggers one of two other accompanying samples. The automatic pairings of samples, can then be reset by forcing the output of unit 3 outside the critical limit, re-randomising the connections. This is repeated until the desired samples play together. After this point, the pairings will persist until the system is reconfigured.

In the example implemented, the perturbation of unit 1's output upward triggers a woodblock loop, a kickdrum loop is played after downward perturbation. A shown in figure 3.1, if unit 1 is moved up, and 2 moves similarly, a triangle sample loops, if opposite, an agogo loop is triggered. The opposite associations are made for a downward movement of unit 1.

The samples are all synchronised as they are written using low -level MIDI signals being sent in the same iterative time-loop as the homeostat algorithm. If developed to use other samples an extra beat/tempo matching algorithm would need to be added.

Chapter 4

Implementation

The front-end of all systems is a Windows GUI, developed in Borland C++ builder. All interaction with the GUI was therefore written in OOP C++, but the main algorithms are in procedural C. All code was written by the author, except a MIDI component, which converts low-level MIDI messages to a MIDI synth using the MIDI API services[1] As the most complex and developed system, all descriptions relate to the Autonomous AdSyM unless otherwise stated.

Data Structures and Program organisation To enable the easy development of different versions of software, a modular approach to program structure was adopted. Therefore structures were used to store all variables, and code seperated into different units, containing the structure definitions and associated functions (akin to C++ objects) according to conceptual and functional use.

The principal divisions are the between homeostat algorithm and MIDI realisation functions, the GUI existing as another seperate object. MIDI handling functions are further seperated into melody and rhythm files. A schematic diagram of the structrual and functional divisions is shown in figure 4.1.

Specifically, 'Main' is an object that creates and controls the GUI, handling information flow between the GUI displays\interactive facilities and data held in structures. The constituent variables of the homeostat algorithm in each version are held in an embedded struct. All MIDI relevant variables for the harmony and melody parts are held in a seperate structure (HOMMIDI). The corresponding variables controlling these aspects of the rhythmic part, are represented as members of seperate instances of a third structure (PIZZRHYTHM). Each structure is declared in seperate header files along with the associated functions.

This modular organisation allows maximum re-usability in the development of different versions of homeostat-based software. The use of structures for data storage also facilitates easy interaction with the GUI: allowing user changes to directly update the

---

[1] The component was written by Sergio Fumana Grunwaldt and downloaded from http:\\www.bytamin-c.com\components\Component-MIDIsv121.htm. The source code is given in appendix ??

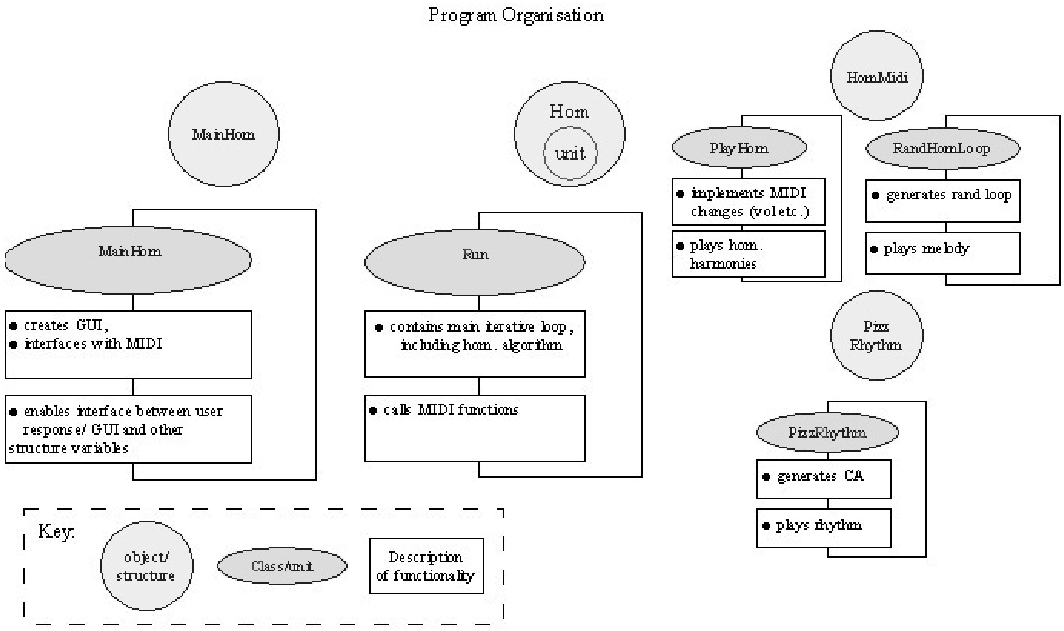

Figure 4.1: Diagram of Program functional and structural divisions.

members of the relevant structure in the Interactive AdSym. This removes the need for double representation, and extra 'check user update' functions. For similar reasons, the interface between GUI and data structures means that the GUI can easily be used to simply *display*, rather than receive variables as it does in the Autonomous AdSym, without making any radical changes to the program architecture.

Homeostat Implementation of the homoestat is fairly trivial. The use of seperate unit structures, within a larger Hom struct, means the size of system can be easily change. Implementing a network is then just a matter of setting the correct connections. So although the 3 programs contain conceptually disparate homeostatic systems practically, it is just a matter of changing the number of units, and using slightly modified connection functions.

Once embedded in the large program, processing restrictions on slower machines required a little code optimisation[2]. The only slight complexity here then, is the assignment of the same memory addresses to the pointers to outputs, and inputs, the latter being used in the subsequent iteration, obviating the need for another double update loop. This was sufficient to solve the interuption.

---

[2] Although the none of the programs are particulalry computationally expensive, because the system runs in real time, slight delays can be heard at a quick tempo when software synthesisers are used.

4.1 MIDI Realisation

The code for selecting the MIDI port and opening the channels can be found at the top of MainHom. Once this is set, audio is produced by sending low-level MIDI messages, which are realised using the selected synthesiser.

A brief MIDI primer. MIDI information is transmitted serially as 3 bytes. The first byte or *status byte* conveys the *type* of message, the following 1 or 2 *data bytes* elaborate, providing the relevant information. The component used sends these as hex chars. An example is given in below:

byte type:	Status byte	Data 1	Data 2
message description:	channel 1 note on	pitch	velocity
actual message:	play	middle C	very loud
Hex form:	0x90	0x3C	0x7B

The main message requiring explanation is the MIDI pitch bend signal. This operates on all notes played on the channel on which it is sent, and changes the pitch up or down by a fractional amount. The two data bytes are merge giving an MSB and a LSB enabling 16384 units of control. The control is centralised: values below 4096 lowering the pitch and those above, raising it. This provides control down to the nearest 128_{th} of a tone. However, such fine tuning is almost imperceptible, so in the current program only the MSB is used, giving 64_{th} tones (and thus the semi-tone is split into 32 as in section 3.1). The major shortcoming of the current approach is that conversion of the outputs (stored as doubles) to MIDI pitchbend (char) loses significant digits. However, the loss will be consistent, and so should not create any artefacts[3].

Channel Assignment As mentioned, the pitchbend operates over the whole channel. Obviously here, each note is individually 'bent', therefore each different pitch value is assigned a different channel. There are 16 MIDI channels, all are equivalent except channel 10 which is the percussion channel. On the other 15 channels, the synth instrument is selected by sending a 2 byte patch selection message, and can then be 'played' at different pitches using the 'note on' message detailed above. On channel 10 however, *all* instrument are untuned percussion instruments. This means, that patch selection is not required, and instead, the *pitch* (data 1) byte controls which instrument is played. Channel and patch assignment for the Autonomous AdSyM is as follows:

Channel nos	Part	Patch description	MIDI no
1 - 4	homeostat sustained chords	synth pad (warm)	0x5A
5 - 8	random melody line	vibraphone	0xC0
11-15	pitched rhythm part	plucked strings	0x2E

For the AdSym Interactive Loop Generator, the rhythmic part was set to channel 10.

---

[3] The interested or bemused reader is refered to **http://www.midiworld.com/basics.htm** for an introduction and **http://www.midi.com** for detailed information

4.2 Real time issues

The development of industry standard time threads was beyond the scope of the current project. For development purposes, the accuracy of the system clock is sufficient. Here then, a simple time thread is implemented using the windows VCL function GetTickCount() which returns the number of ms since windows was started.

All loops requiring synchronised calls to the MIDI component operate within a loop of the following form:

```
current_time = last_time = 0

begin outerloop
   current_time = GetTickCount()
   if(current_time >= (last_time + timestep))
      begin innerloop
         .
         put(MIDI note on messages)
         .
         last_time = current_time
      end innerloop
   .
   .
   .
end outerloop
```

This produces audio output that sounds synchronous to the human ear. Because all applications run in Windows, API function ApplicationProcessMessages() is included at intervals to allow windows system processes time on the CPU.

Chapter 5

Evaluation and Analysis

5.1 Interactive Loop Generator

Examples of music generated by the Interactive AdSym can be found on tracks 4-7 of the cd. For all tracks, maxchange = 0.008, critical deviation = 0.05 and the 4 unit homeostat is fully connected.

Track 4. Homeostat loops and random melody This track demonstrates the harmonic loops created by the homeostat in state of stable oscillation, as well as the effect of minor perturbation and the accompanying random melody loop.
Timestep = 40 ms, PB factor = 10.

At the start of the track, the homeostat plays unaccompanied: the system is in a stable state of oscillation, the different periods of oscillation of each unit creating an interesting loop. At about 6 secs, the output of unit 1 is forced downwards by $\frac{1}{2}$ the critical deviation. This causes a brief harmonic deviation. The system quickly restabilises and the original loop returns. At 17 seconds the ‘melody’ line is turned on, and a simple melody loop accompanies.

Track 5. High speed homeostat In this track, the harmonies are played with just the percussion part and timestep reduced to 1 ms. PB factor = 5.

At this speed, the individual notes of the rhythm part and any harmonic oscillations are almost imperceptible: the rhythm becomes almost a timbre, and the harmonies experienced as sustained chords. The chords are induced as above by forcing the output of unit 1 outside its critical limit. Given the microtonality of the system, the chords produced are surprisingly familiar. The progression contains tension and resolution similar to that experienced in tonal music, and even seems to return to the original chord.

Track 6. 80’s homeostat Although changes in instrumentation were not included on the GUI, an example of the effects of using different instruments is given here. The melody line is played on the piano and different percussion instruments used. The time step is similar to that of track 4, but the melody line contains a note that is played every iteration, giving a very 'driving' effect. The regular repetition of this note also makes the system states clearly audible: the example provides a good illustration of the progression through various states of stability and (user induced) perturbation.

Track 7: Autonomous mode In the above tracks, the system was to some extent 'played' like an instrument or sampler - different instruments turned on or off, parameters adjusted and the outputs interfered with to produce harmonic changes. This example demonstrates the system in automatic mode.

All three parts are played, the melody line using a 'vibes' voice, the rhythm part on triangle, conga and cabasa, and the harmony on a synth pad. Time step = 200 ms, PB factor = 1.

The homeostat quickly settles to a stable oscillation, the melody picking out a simple line from these notes. The changing proportion of states in the CA can be clearly heard: the percussion part initially including rests, becomes almost continous. At 42 secs, the percussion loop ends, triggering a critical perturbation of unit 1. After a brief period, the system restabilises, returning to repetitive harmonic loops.

General Comments These examples illustrate that the homeostat harmonies, random melody and CA rhythm produce some interesting musical fragments, both individually and in different combinations. The range of effects produced just by varying the combination of parts, tempo and pitch bend factor is quite impressive.

The system was presented at the Sussex University MSc poster presentation, as well as privately to a professional digital-music producer and a multi-media software engineer. All found the system 'intruiging', but only those who spent some time playing with the parameters showed a high level of interest. Several students of EASy commented that playing with a real-time audio version of a homeostat gave them their first real insight into 'what Ashby was talking about', suggesting the system also has some didactic value. One of the developers of V-Jam, a real-time MIDI triggered video mixer, commented that the patterns produced - particularly by the oscillatory states - may be interestingly used as video effects.[1]

5.2 Adaptive Sequencer Module

An audio example of the module controlling the pairings of 4 different preset samples can be found on track 8 of the cd. Once the system had stabilised, unit 1's output was decremented slightly triggering a woodblock sample, the weight configurations at this time cause unit 2's output to similarly increase: a triangle sample is played simultaneously. After 4 bars (7 secs), unit 1's output was decreased, triggering the kickdrum sample. Because the weight are the same, unit 2 moves similarly downwards and the agogo sample is played. Shortly after this point, the output of unit 3 is perturbed beyond its critical limit, and new weights are assigned. Thus when after another 4 bars

---

[1] This is easily implemented as V-Jam uses MIDI pitchbend signals to control framespeed.

(at 16 secs), the output of unit 1 is incremented and the woodblock plays again, the agogo sample continues, as the weight changes mean that unit 2 now moves opposite unit 1. To ensure that this is the case, unit 1 output is forced down a final time, and as expected, the kickdrum sample is played with the triangle.

The prototype was presented to several professional and semi-professional digital music producers and software designers/engineers. All commented on the current call for such adaptive mechanisms in digital audio engineering and confered it the system had potential in this field.

5.3 Autonomous Homeostatic Music System

The first 10 minutes of a sample presented both publicy and in the listener survey can be found on track 9 of the cd.

5.3.1 Public Appearances

The Autonomous AdSyM was very well received at two public appearances made at different stages of development. Early on in the project, the basic homeostat (essentially just the harmonies generated in AdSyM I) was used to open and close a live set of the headlining act 'Pthhh'[2] of the last night of the festival of Occulture 2002. (www.occulture.tv). Several audience members commented on the effectiveness of the homeostat as an outstanding atmospheric effect.

The full Autonomous AdSyM as presented here, was played for several hours in bar Sumo to an audience consisting of the members of *Blip* - a Science/Art discussion forum held at the Brighton Media Centre. The setting provided an ideal stage for the system, as listeners could experience the large scale structure of the system in terms of its' progression through various states over an extended time.

All those attending were 'very impressed' and commented that the system was 'exceptionally musical' compared to other, even knowledge based, generative systems. The Australian generative artist Paul Brown, the evenings' main speaker, showed similar enthsiasm, as did a representative from a renowned interactive multimedia company 'audiorom.com', who expressed an interest in commissioning the work for an installation project. John Petigrew, one of the founding members of SSeyo also expressed an interest in the system for use on his planned 'evolved art' cable channel.

In other private presentation, comments such as 'ooh is that you on the 'cello?' were made by naive listeners, who were all surprised to learn that what they were listening to was entirely computer generated. Although no comparison were made to any of the Great Composers, several people said it reminded them of Steve Reich's work, and many described it as 'film music', suggesting it has strong atmospheric value.

---

[2] The band is best described as electroacoustic-avant-garde-contemporary-classical. The four members play 'cello (myself), violins, guitar, flute, whistle and various electronic samplers

5.3.2 Listener Survey

Listener Evaluation: Method

Participants 10 male and 10 female participants between the ages of 21 and 35 took part in the listener evaluation. Participants were selected to cover a range of musical, and Adaptive Systems interest and knowledge. All participants were aware that the music was digitally based, but none knew the exact nature or details of the compositional process.

Design and Procedure The evalulation consisted of a mixture of structured and semi-structured feedback, made both during and after listening.

The audio was presented in real time from through headphones from a laptop. In order to enable comparison between responses within the small sample, the same version was presented to all listeners. Each were presented with a GUI consisting of just four buttons *start*, *stop*, and *like it* and *don't like it* (shown in appendix B, figure B.4). The full GUI was hidden to prevent distraction and discrete the system details from the more inquisitive listeners. Each listener was given the same set of written instructions (appendix D.1). These requested simply that they listen *as long as they wanted.* This was done to provide an implicit measure of interest. In addition, they were given the option to press *like it* or *don't like it* at any point during their listening. This was included for two reasons: the iteration number and type of response was recorded to investigate whether there were any universally dis/approved of points; the instruction was also designed to encourage a similar level of attentiveness and listening style in all users.

After listening, all participants filled out a short questionaire (appendixD.2) requiring selection of a scaled response. In addition, they were asked to justify their response to whether the music bore any qualities they *'normally associated with music'* and were prompted for 'further comments'. These requests and the 'don't/like it' responses were intentionally vague in keeping with the inherently subjective nature of musical appreciation.

The scaled questionnaire developed consisted of 10 items that were designed to gain direct evaluative, as well as more abstract subjective feedback (items 1,2,4,5,10 and 3,6,7,8,9 respectively). The evaluative items were designed to ascertain the degree to which it was deemed interesting, enjoyable and musical; the subjective items addressed structural aspects such as perception of degrees of repetition and novelty. Participants were given the choice of 5 response categories (Strongly Agree, Agree, Undecided, Disagree, Strongly Disagree). Although the omission of a central category theoretically forces particpant choice, the use of such scales increases the chances of non-response. Items were very short and simply worded and included a mixture of positive and negative statements.

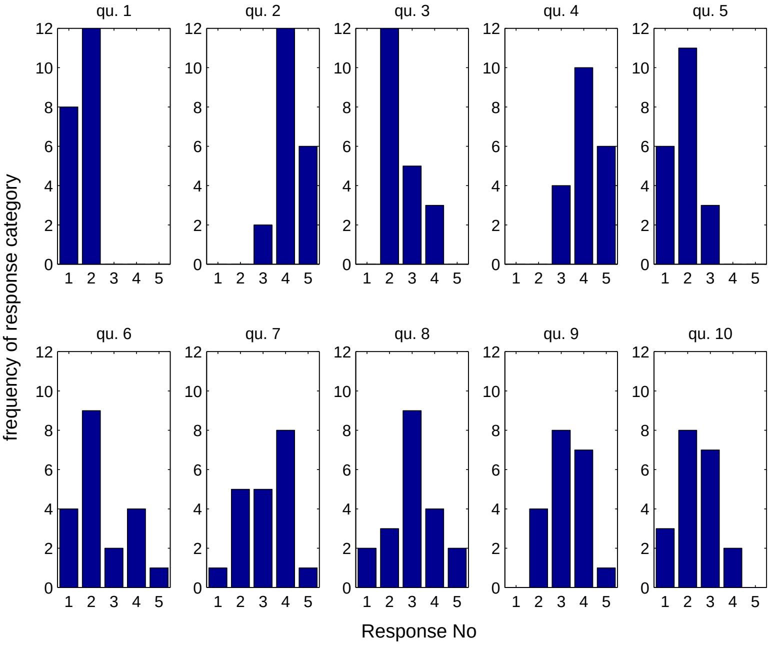

Figure 5.1: Frequency distribution of responses for each item.

Listener Evaluation: Analysis and Results

Analysis A Mann-Whitney U-test was performed, comparing the frequency of responses for each item of the questionnaire with the expected frequency (ie each of the 5 categories contain 20% of the responses)[3].

Results: Listening time Mean Listening time: 229 iterations (approx. 8 mins)
standard deviation: 127 (approx. 4 mins)
The average listening time (over twice your average pop song!) suggests the music was in some way engaging. The large sd suggests a large degree of variation of interest and/or enjoyment amongst listeners.

Results: Questionnaire response Figure 5.1 shows the frequency distributions of responses for each item. The modal category can be clearly seen in each case, 8 of which are valenced (ie not *undecided*). However, as the skew of the distributions suggest, only response frequencies for items 1,2,4 and 5 differ significantly from the even distribution. The mode, U-test value and p values (2-tail) are given in table 5.3.2, items for which responses differed sigificantly from the null case are shown in bold.

Analysis of the questionnaire responses implies that listeners:

---

[3] As there is no control group, this merely serves as a comparison and a means of ascertaining if there is significant agreement of response amongst listeners

Questionnaire Item	Mode response	U value	p value (2t)
1. **It was interesting**	**Agree**	5.5475	**0.0001**
2. **It was unmusical**	**Disagree**	-3.5023	**0.0002**
3. It was repetitive	Agree	0.8870	0.1875
4. **I would not describe it as music**	**Disagree**	3.1304	**0.0009**
5. **It was enjoyable**	**Agree**	2.3719	**0.0088**
6. I was sometimes surprised	Agree	1.1385	0.1275
7. It was not like anything I have heard before	Disagree	-0.3238	0.3730
8. It was directionless	Undecided	-0.1065	0.4576
9. It was predictable	Undecided	0.5370	0.2956
10. It could have been written by a person	Agree	1.2201	0.1112

Table 5.1: Mode response and Mann-Whitney test statistics

- Agreed that the music was interesting.
- Disgreed that it was unmusical (ie felt it was musical)
- Disagreed that they would not describe it was music (ie they would describe it as music).
- Agreed that it was enjoyable. (at p*leq* 0.01)

Only 2 people disagreed with the remaining evaluative item (*"It could have been written by a person"*). Perhaps not surprisingly, there was no significant consensus on the more subjective items. It is of interest however, that 60% of participants agreed that it was repetitive and 65% agreed that they were sometimes surprised. This suggests that the music achieved a mixture of repetition and novelty.

Real time judgements Figures 5.2 and 5.3 show the total instances of button presses of *like* and *dislikes*. Results are shown over just the first 300 iterations, as 80% of participants did not listen beyond this point. In total, 240 *likes* were made compared with 82 *dislikes*. There is not a great degree of absolute comorbidity in responses, although the time interval between iterations is approximately a second and there does seem to be clustering in both categories. The largest decisive gap in *like* responses is around iteration 225; this point also shows the clearest burst of *dislikes* (at this point less than half of participants are listenting). Interestingly, this coincides with a unique pattern of behaviour in the four units that define the pitch. Figure 5.4 shows the output of the 4 units that define the pitch (and thus describes the harmonic changes). Around iteration 225, output values change smoothly, in contrast to the previous oscilatory behaviour, producing a gradual contrary motion between the 4 parts in constrast to previous large jumps. This may be because the gradual change makes the microtonality more evident, listeners being unused to such continuous changes in pitch[4].

---

[4] It is also quite ironic, as one of the first laws in Western harmonisation is to avoid large jumps (in the bass line in particular) and encourage smooth progression!

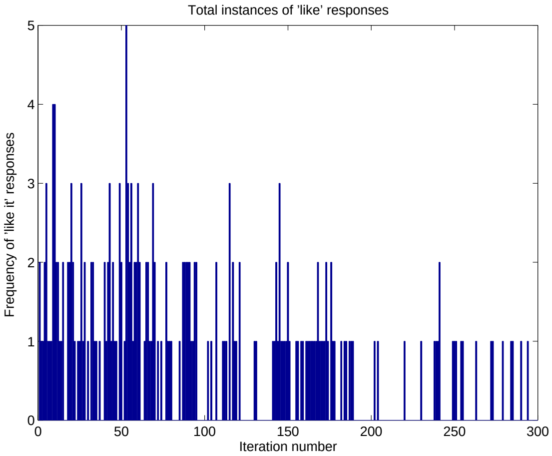

Figure 5.2: Bar chart showing occurences of 'likes' for all participants.

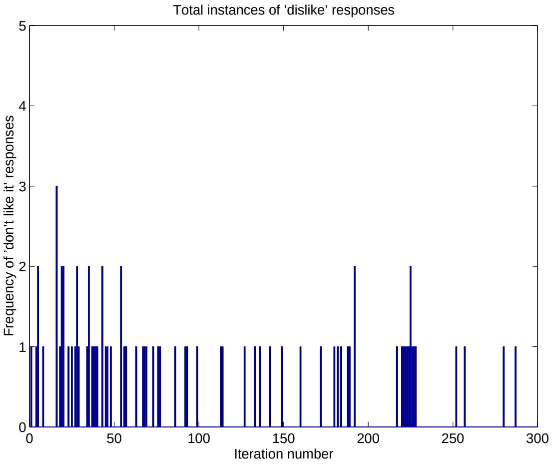

Figure 5.3: Bar chart showing occurences of 'dislikes' for all participants

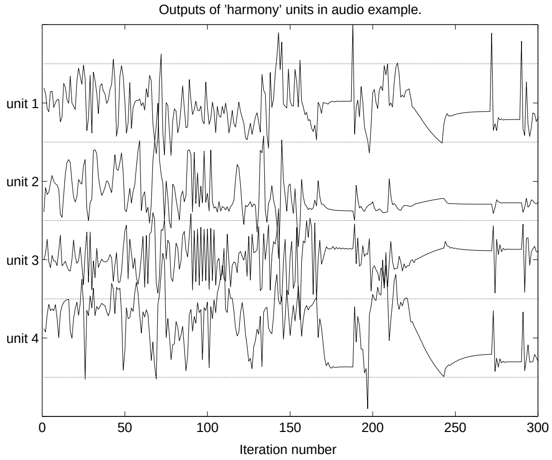

Figure 5.4: Output of units 1-4 for the first 300 iterations for the evaluated version.

open questions 19 of the 20 participants agreed that the audio example they had listened to bore qualities that they normally associated with music. In elaboration of their choice, several mentioned the presence of many of the normal elements of music: *'sense of melody ...driving sense of rhythm', 'there were definite harmonies if unusual at times', 'sense of harmonic and rhythmic structure and melodic progression'.* This reference to structure was made by several listeners: *'structure and development on different timescales/resolutions', 'certainly, if not composed by a person it must have been restricted in scale, structure etc.'.* Not surprisingly, *discordant* was frequently mentioned, but there were no negative comments regarding the microtonality.

The further comments of many listeners suggest that the music has emotive qualities: *'tension building and resolution of tension', 'It had the ability to generate mood...', 'It oozed atmosphere...* [5] .

Perhaps surprisingly, the periods of stability, leading to repetitive melodic loops, were not appreciated by anyone. In fact several listeners commented that they prefered the periods of change, and *'got bored in the samey bits'.*

5.3.3 System Behaviour

Figure 5.5 shows the outputs of all units over 2250 iterations (approximately 1 hour of listening time). The specific connections for this run are shown in appendix C. Examination of the outputs of each units suggest that the system passes through several qualitatively different states. These apparent phase changes are evident throughout the whole system. These are summarised below.

iteration	0-250	250-700	700-950	950-1250	1250-1650	1650-3500
state	runaway	stable	stable	stable	runaway	stable
description	random	conv/osc	oscillate	converge	random	static

This progression through and return to qualitatively different states provides a rudimentary high level structure.

Note that weight changes (induced by foring unit 1 output out of limits - evidenced by the sharp upwards lines in unit 1 output) are made at frequent intervals throughout most of these phases, but the qualitative behaviour persists. This suggests that the phase changes are not driven by something other than simple perturbation.

Although this stability is perhaps detrimental aesthetically, from an Adaptive Systems perspective, it is of further interest that the system finally stabilises to a point cycle that is seeemingly impervious to perturbation (fig. 5.6). From iteration 1650 onwards, the output of unit 1 frequently forced outside its limit, evoking weight changes. The effect of the movement can be seen in all units, but none transgress their critical boundaries. It seems that the system has reached an 'ultimately' stable state.

In the simple Ashbian homeostat, stability it achieved through random selection of suitable conenctions values. It could be that this self-regulating network is similarly achieving ultimate stability through the random selection of a suitable maxchange value, ie. as soon as a minimal value of maxchange is imposed, output variation is negligable,

---

[5] *(...of a crazy graveyard in Iceland'!)*

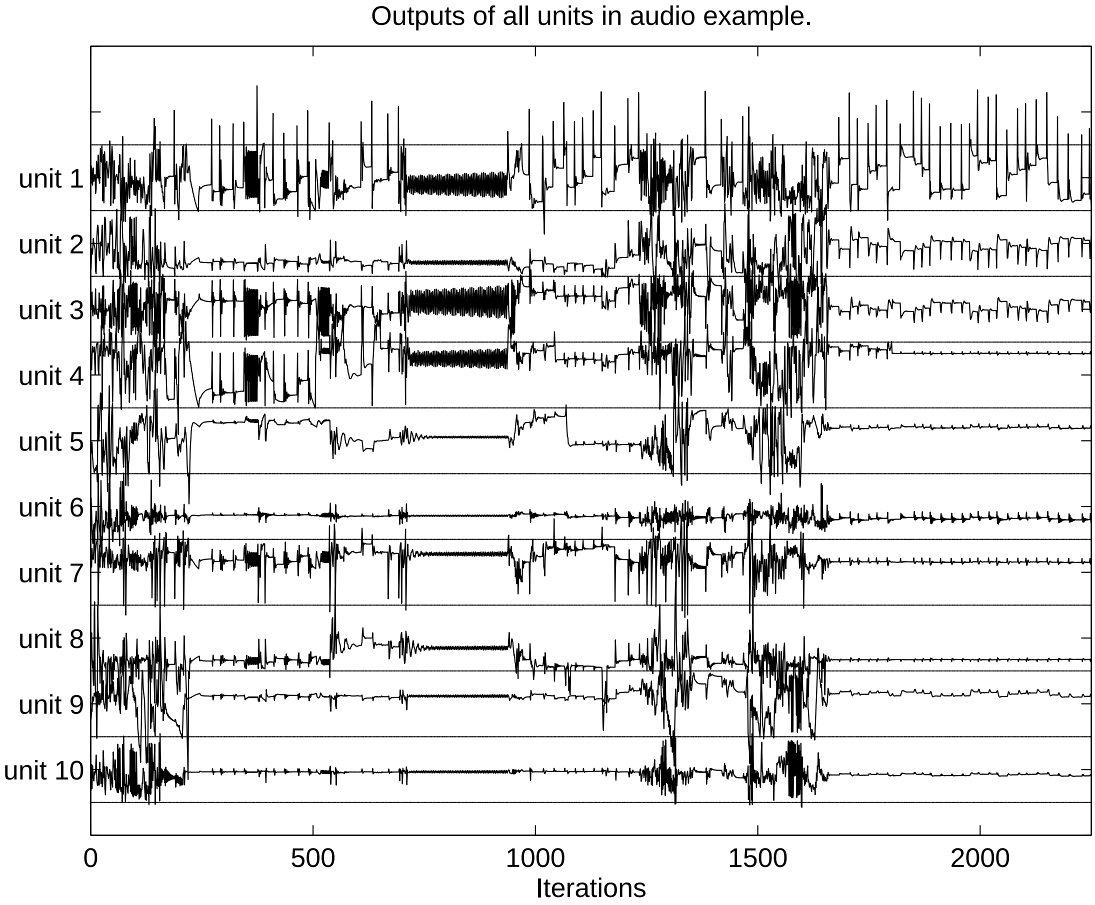

Figure 5.5: Graph of outputs of whole network over 2250 iterations

and the system stabilises. Figure 5.7 shows the value of maxchange over this same run. The final stable state is reached at around iteration 1650. As the figure shows, at this point maxchange oscillates (driven by the perturbations of unit 1) around 0.015. This is *not* the lowest value to date. This suggests that this ultimate stability is not simply a result of minising maxchange but perhaps a property of the dynamics of the network.

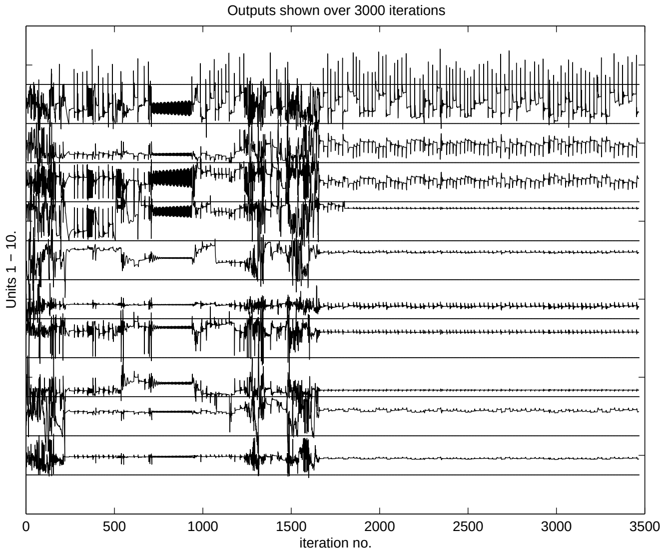

Figure 5.6: Graph of outputs of whole network over 3500 iterations

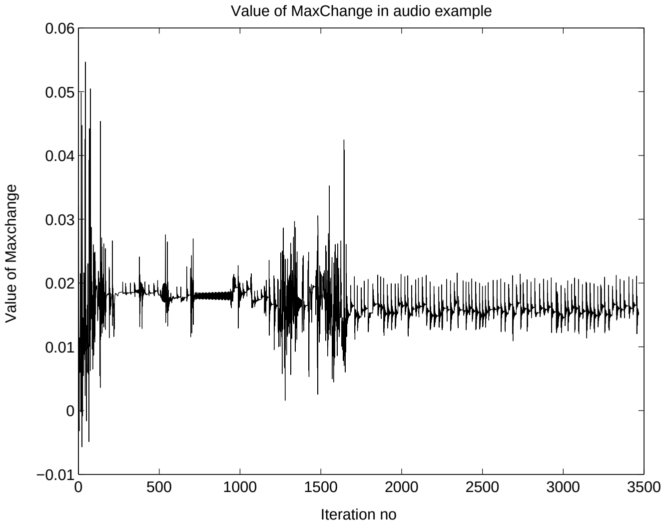

Figure 5.7: Value of Maxchange over 3500 iterations

Chapter 6

Discussion

The successes, shortcomings and possibilities for further development of each of the three AdSyM systems are reviewed. The constituent algorithms are then given seperate consideration both in terms of aesthetic worth, and possibilities of more theoretical application.

6.1 Adaptive Sequencer

The sequencer module as it stands is very simple, yet exists as a 'proof of concept' that homeostatic mechanisms can be deployed as control systems. An advantage of such a system is that changes can be made in real time - in the example given, the system was succesfully reconfigured in just 4 bars while it was playing. However, because weight change is random, the number of changes needed to ensure reversal of effect is infact indeterminate. In addition, once the system has been 'inverted', if the user wants to return to the original pairings, this will take just as long: in Ashby's terms, the system cannot accumulate adaptations.

If the module is to be used as an alternative to 'presets' in a traditional sequencer, this is a definite drawback, as it means the advantage of real time changes is undermined: real time changes can be made, but there is no guarantee *when* they will occur.

This indeterminacy however, suggests a new approach to traditional audio production techniques. Just as composers have been fascinated by the use of random processes in the generation of music, there seems to be a rising level of interest in 'stochastic' control in digital audio production. For example Matt Black of Cold Cut (Ninja Tunes) has recently commissioned a "Cold Cutter" plug-in for Cubase (personal communication). This effectively takes existing samples (eg simple drum loops as here), splices them into random lengths and re-organsises them randomly. Talking to others in the industry, as well as enthusiastic 'bedroom producers', such mechanisms are currently in demand. Those that viewed the system, conferred that such mechanims had potential in this area.

6.2 Interactive AdSyM

The Interactive loop generator demonstrated the potential of homeostatic processes and CAs in the generation of harmony. In combination these algorithms produced some interesting musical sequences that could conveivably be used as building blocks in a more convential digital composition process. Whilst the system was not capable of autonomously generating large scale musical forms, a large range of musical effects can be produced by changing just a few variables. This suggests it could form the basis of an adaptive parametric digital compositional tool, bridging the gap between sequencers and generative software.

6.3 Autonomous AdSyM

Response regarding the aesthetic worth of the music produced by the Autonomous AdSyM was goo. As well as complimentary remarks from all those attending public performances, the listener survey produced encouraging results. Despite the assumed large variation in interest, there was unanimous agreement that the music was interesting, and a general consensus that it was musical and enjoyable. Listener response suggests that the system succeeded in generating new music.

Several listeners who took part in the survey commented on the effects of the tempo and dynamic changes. Although the values of the latter were controlled by the outputs of the homeostat, the point at which they were referenced was triggered by the preset length of the CA repetition. Similarly a sudden decrease in tempo was triggered by stability in the homeostat, but incremental increases were hardwired. Given the importance of tempo and dynamics (speed and volume levels) in music, these effects should perhaps be more meaningfully controlled.

To this end, the apparent ability of the system to create a hi-level structure could be exploited. Rather than simply defining the harmonic structure, qualitative changes in behaviour could be used to determine dynamics and tempi. These mappings could be made in simple yet meaningful ways such as stable phase = quiet and slow, runaway phase = loud and fast.

The principle aim here was to produce an autonomous generative music system. Whilst here a feedback was implemented such that the network was perturbed when it reached stability, a fuller interaction between the homeostatic network and CA could be interesting. By generating the CA in realtime the achievement of the CA endstate (be it homogeneity, period repetition etc.) could be used to trigger changes in the homeostatic network. Possibilities such as the state of the homeostatic network affecting the rule base for the CA also seem worthy of exploration. The properties of the homeostat - ie responsiveness to external interference make it an exciting basis for the development of more complex interactive systems. The interactivity of the system could also be feasibly raised a level, such that the digital system interacted with live performers (via a MIDI digital interface). This would enable represetation of Ashby's original conceptual distinction between organism and environment. Such systems are currently under consideration.

6.4 Cellular Automata as the basis of rhythm

As hoped the pattern propagation properties of CAs transfer succesfully to the auditory domain as rhythm. The simultaneous presentation of several consecutive states using instruments with different timbres (in version 1) produced effective accents, the use of pitched instruments in version 2 producing another almost melodic line.

Although rhythmic, the repetition of the same initial states of the CA became boring. As well as aesthetically, it seems important to conceptually to employ CA's that evolve in real time. Only one CA rule base was used here, and the vast array of possible automata offer much scope for further investigation. Higher dimensional automaton, with multiple states warrant exploration; the multiple states being mapped onto variations in pitch or timbre. Comparisons between different classes would also be interesting. For example the effects of the evolution of classes 1 and 3 from initial random and regularly seeded states respectively. Almost universally (with the exception of some New York free jazz drummers!), musical rhythm is strictly repetitive and\or periodic. The effect of a transition from order to chaos or vice versa on the listener would be interesting to explore.

There is an increasing use of audio material in research into human pattern recognition and categoristation [Palmer and Krumhansl, 1990]. It has been proposed that listeners have internalized metrical hierarchies and use them to perceptually organize temporal patterns [Takeuchi, 1993]. In the survey presented here, several of the more experienced musicians commented on the '$\frac{4}{4}$ syncopated' rhythm. Given that the CA was 13 cells wide, no such regular period existed. This metre must have been infered. The derivation of patterns from formal rule sets such as CAs could provide an interesting new methodology in this area.

The success of the CA in producing musical material supports previous suggestions that CAs and music may share similar organisational principles [Miranda, 2002]. The potential for modelling musical processes using CAs could provide an in-road into the development of an abstract theory of music.

6.5 Homeostasis as the basis for harmony

The homeostatic control of pitch was succesful in a number of ways. On the lowest level, the use of pitchbend to implement microtonality introduced some intruiging, if unconventional chords. The frequent state of stable oscillations of different periods produced effective loops at suitable timesteps. In the simple 4 unit homeostat, the functional homeostatic property of resistance minor perturbation enabled the production of deviations from and returns to harmonic patterns seen in more traditional musics.

6.5.1 Self-regulating Homeostatic Networks

More interesting, however, is the top level structure produced by the regular perturbation of teh self-regulating network in the Autonomous version. The cycle through different phases of random-runaway and stable behaviour, bears resembelance to Larivaille's scheme for musical\literary narrative based on a cycle of equilibrium and disturbance, and arguably constitutes a structure of the highest-level of musical form. Certainly it seems to be a higher level of organisation than that achieved by any other generative systems to date. From a complex systems perspective, it is worthy of note that this high level behaviour is essentially produced from the interaction of the low-level rules governing the calculation of output in individual units.

The stable-runaway-stable pattern is characteristic of the behaviour of a simple homeostat at the level of response to individual perturbation, but here a similar pattern exists at a higher level: any one state persisting for a time despite frequent critical perturbation. Similar patterns of behaviour seems to exist at multiple levels. Given the non-linear nature of the system, and its sensitivity to intial conditions (The initial weight values impacting heavily on the state of the system), this self-similarity promotes consideration that the system may be chaotic. The possibility warrants further investigation.[1]

As well as wishing to produce interesting aesthetic behaviour, the development of the self-regulating network was motivated by a desire to closer approximate the dynamics of homeostatic networks in living organisms. The apparent state of 'ultimate' stability on which the system converges is arguably closer to the truly autoregulated physiological systems.

Although the simple 4 unit Ashbian homeostat invariably re-stabilises following the critical perturbation of one unit's output, this often takes several iterations and involves the transgression of one or more other units beyond the limit. In larger networks, such perturbations have an even greater tendency to cause runaway behaviour. In this case however, in the final phase, all other units stay within their limits despite repeated perturbation. The system achieves an ultimately stable state. The nature of the subserving mechanism is not clear from the investigations so far. That the system did not simply stabilise when a minimal value of maxchange was encountered, suggesting that it is a more complex than simple random search for an optimum value. If this can be shown to be a general property of self-regulating systems, the simplicity and robustness suggest application in engineering systems as well as in evolutionary robotics.

There has been a surprising lack of consideration of any similar homeostatic mechanisms in Adaptive Systems research (with one notable exception [Di Paolo, 2000]). Exciting advancements are being made in the complexity of behaviour shown by simulated agents, but these are almost exclusively based on evolved architectures [Slocum and Beer, 2000]. The use of homeostatic mechanisms enables ontogenetic rather than the phylogenetic learning of EAs and should be explored.

---

[1] Attempts were made to this end, to calculate the Lyapunov exponent from the time series data, but problems with the available software prevent the presentation of findings in the current paper.

6.6 Other Applications: Visualisation of complex systems

The audio realisation of both CA and homeostat provide an interesting new approach to visualisation of complex adaptive systems. With all multi-dimensional complex systems, graphical visualistation is inherently limited to 2, or at the maximum, 3 dimensions. Thus investigation of such systems, usually requires prior treatment with PCA or similar dimension-reducing processing. Variations in pitch, volume and timbre allow a vast number of discernable dimensions in the auditory domain. As evidenced here, audio realisation of dynamical systems provides a very effective means of visualising qualitative state changes as well as quantitative variation in particular variables. There exist many problems within Evolutionary Robotics and other simulations of adaptive behaviiour which would benefit from such an approach. Some difficult problems and impressive tasks have been solved using hi-dimensional DRNNS as agent architecture [Slocum and Beer, 2000]. Although these agents appear to solve difficult task, *how* they accomplish it remains a mystery. Being able to *listen* to the behaviour of such networks (eg weights, outputs etc), may give insight into the mechanisms subserving the high level behaviour. The avenue warrants further investigation.

6.7 Conclusions

> "...weird and surprising yet strangely familiar..." (listener 15)

The initial exploration of the deployment of homeostatic systems and cellular automata in the musical domain produced some promising results. Although not pertaining to any particular style, listeners agreed that the sounds produced were 'music'. The combination of homeostatic harmonies, CA rhythm and random melody placement appears to have gone beyond mere arrangements of sound in time, and captured a certain essential musical quality. This use adaptive algorithms, and in particular the exploitation of the *process* of adaptation constitutes a new approach to the generation of new music, and invites further study.

The self-regulating homeostatic network developed appears to show a final state of 'ultimate' stablity. The system demands further investigation and may prove useful in both theoretical and applied branches of Adaptive Systems.

The possibilty of using cellular automata to model natural music, and the deployment of physiologically inspired processes in the generation of digital music promises a much-talked-about and seldom-realised reciprocity between the sciences of the natural and artificial in the musical domain.

Bibliography

[Ashby, 1956] Ashby, W. (1956). *Design for a Brain.* Chapman and Hall, 2nd edition.

[Ashby, 1970] Ashby, W. (1970). Connectance of large dynamic (cybernetic) systems: Critical values for stability. *Nature*, 228(784).

[Bernard, 1870] Bernard, C. (1870). Lecons sur les phenomaenes de la vie communs aux animaux at aux vegetau. *American Lecture Series*, (900).

[Biles, 1994] Biles, J. (1994). Genjam: A genetic algorithm for generating jazz solos. *Proceedings of the 1994 International Computer Music Conference.*

[Boulez, 1963] Boulez, P. (1963). *Penser al Musique aujourd'hui.* Gonthier.

[Cope, 1991] Cope (1991). *Computers and Musical Style.* OUP.

[Di Paolo, 2000] Di Paolo, E. (2000). Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions. In *Proc. of SAB'2000.* MIT Press.

[Fux, 1965] Fux, J. (1965). *An Historical Study of Fugue.* Pergamonn Press.

[Gardner, 1970] Gardner, M. (1970). Mathematical games the fantastic combinations of john conway's new solitaire game "life". *Scientific American*, 223:120–123.

[Goya and Bolagnani, 1999] Goya, R. and Bolagnani, F. (1999). Homeostasis, thymic hormones and aging. *Gerentology*, 45(3).

[Jordan and Kafalenos, 1997] Jordan, R. and Kafalenos, E. (1997). Listening to music. semiotic and narratological models. *Musikometrika*, 6.

[Langton, 1992] Langton, C. (1992). *Artificial Life II*, volume X of *SFI Studies in the Sciences of Complexity.* Addison-Wesley, Redwood City CA.

[Larivaille, 1974] Larivaille, P. (1974). L'analyse (morpho)logique du recit. *Poetique*, 19:368–388.

[Markus and Kusch, 1995] Markus, M. and Kusch, I. (1995). Cellular automata for modelling the shell pigmentation of molluscs. *Journal of Biological Systems*, 3:999–1011.

[McAlpine and Hoggar, 1999] McAlpine, K. Miranda, E. and Hoggar, S. (1999). Making music with algorithms: a case-study system. *Computer Music Journal*, 23(2).

[McAlpine, 1999] McAlpine, K. (1999). *Applications of Dynamical Systems to Music Composition.* PhD thesis, Universtity of Glasgow.

[Meyer, 1956] Meyer, L. (1956). *Emotion and Meaning in Music.* Chicago University Press, Chicago.

[Miranda, 1995] Miranda, E. (1995). *Granular Synthesis of Sounds by Means of Cellular Automata*, volume 28. Leonardo.

[Miranda, 2001] Miranda, E. (2001). *Composing with Computers.* Focal Press.

[Miranda, 2002] Miranda, E. (2002). Evolving cellular automata music: From sound synthesis to composition. downloaded from www.csl.sony.fr.

[Orton, 1990] Orton, R. (1990). Music fundamentals. Unpublished lecture notes. University of York.

[Palmer and Krumhansl, 1990] Palmer, C. and Krumhansl, C. (1990). Mental representations of musical meter. *Journal of Experimental Psychology: Human Perception and performance*, 16(4):728–741.

[Phon-Amnuaisuk and Wiggins, 1998] Phon-Amnuaisuk, S. Tuson, A. and Wiggins, G. (1998). Evolving musical harmonisation. Technical Report 904, University of Edinburgh.

[Prusinkiewicz, 1994] Prusinkiewicz, P. (1994). Visual models of morphogenesis. *In Artificial Life*, 1(1/2):67–74.

[Revill, 1992] Revill, D. (1992). *The Roaring Silence.* Bloomsbury.

[Schoenberg, 1950] Schoenberg, A. (1950). *Style and Idea*, chapter Composition with 12 tones. Leo Black.

[Schottstaedt, 1989] Schottstaedt, W. (1989). Automatic counterpoint. *Current Directions in Computer Music Research.*

[Simon, 1968] Simon, H. (1968). *Formal Representation of human judgement*, chapter Pattern in Music, pages 219–250. Wiley.

[Slocum and Beer, 2000] Slocum, A.C. Downey, D. and Beer, R. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In J. Meyer, A. Berthoz, D. F. H. R. and Wilson, S., editors, *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior*, pages 430–439. MIT Press.

[Stockhausen, 1989] Stockhausen, K. (1989). *Stockhausen on Music.* Marion Boyars.

[Takeuchi, 1993] Takeuchi, A. (1993). Perceptual organization of temporal patterns in music. Technical Report 4aMU6, Research Lab of Electron., MIT.

[Wolfram, 1984] Wolfram, S. (1984). Universality and complexity in cellular automata. *Physica 10D.*

[Xenakis, 1971] Xenakis, I. (1971). *Formalised Music.* Indiana University Press.

Appendix A

Ashby's description of the Homeostat

"The Homeostat consists of four units, each of which carries on top a pivoted magnet. The angular deviations of the four magnets from the central position provide the four main variables.

Its construction will be described in stages. Each unit emits a D.C. output proportional to the deviation of its magnet from the central position. The output is controlled in the following way. In front of each magnet is a trough of water; electrodes at each end provide a potential gradient. The magnet carries a wire which dips into the water, picks up a potential depending on the position of the magnet, and sends it to the grid of the triode. J provides the anode potential at 150 V., while H is at 180V,; so E carries a constant current. If the grid-potential allows just this current to pass through the valve, then no current will flow through the output. But if the valve passes more or less, current than this, the output circuit will carry the difference in one direction or the other. So after E is adjusted, the output is approximately proportional to M's deviation from its central position.

Next, the units are joined together so that each sends its output to the other three. These inputs act on the unit's magnet through the coils A, B, C, so that the torque on the magnet is approximately proportional to the algebraic sum of the currents in A, B, and C. (D also affects M as self-feedback). But before each input current reaches its coil, it passes though a commutator (X) which determines the polarity of entry to the coil, and through a potentiometer (P), which determines what fraction of the input shall reach the coil.

As soon as the system is switched on, the magnets are moved by the currents from the other units, but these movements change the currents, which modify the movements, and so on. It may be shown that if there is sufficient viscosity in the troughs, the four-variable system of the magnet-positions is approximately state-determined. To this system the commutators and potentiometers act as parameters.

When these parameters are given a definite set of values, the magnets show some definite pattern of behaviour; for the parameters determine the field, and thus the lines of behaviour. If the field is stable, the four magnets move to the central position, where they actively resist any attempt to displace them. if displaced, *co-ordinated* activity brings them back to the centre. Other parameter-settings may, however, give instability; in which case, a 'runaway' occurs and the magnets diverge from the central positions with increasing velocity - till they hit the ends of the troughs."

Appendix B

Graphical User Interfaces

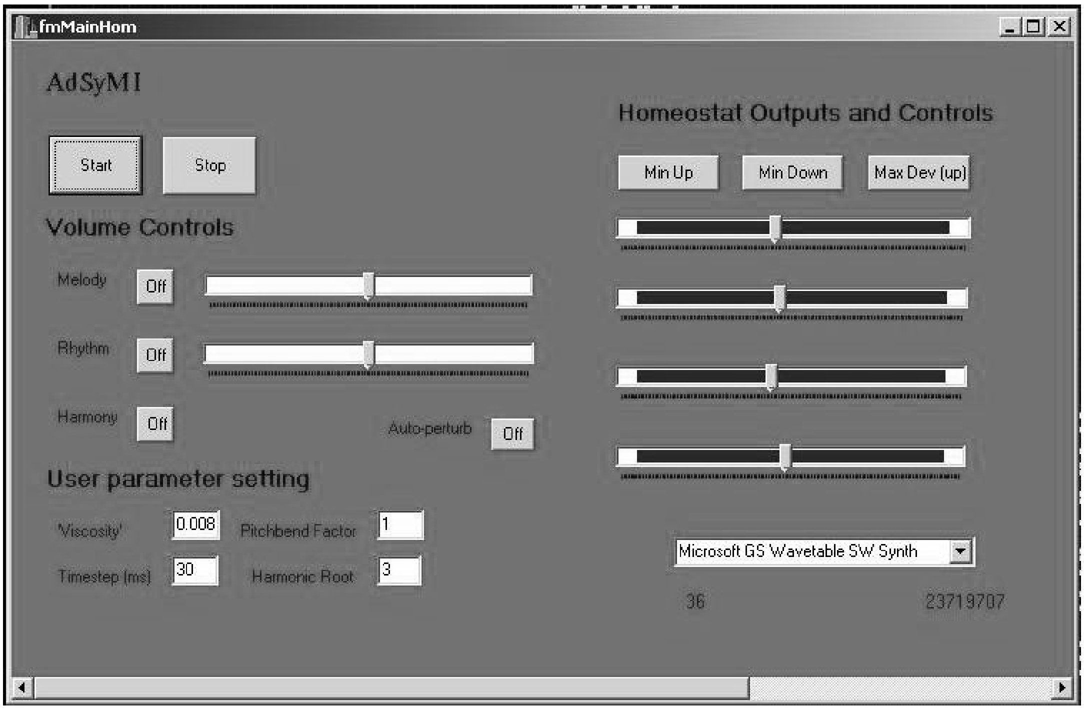

'Viscosity' 0.008

Pitchbend Factor 1

Timestep (ms) 30

Harmonic Root 3

Microsoft GS Wavetable SW Synth

Figure B.1: Screen shot of the Interactive AdSyM GUI

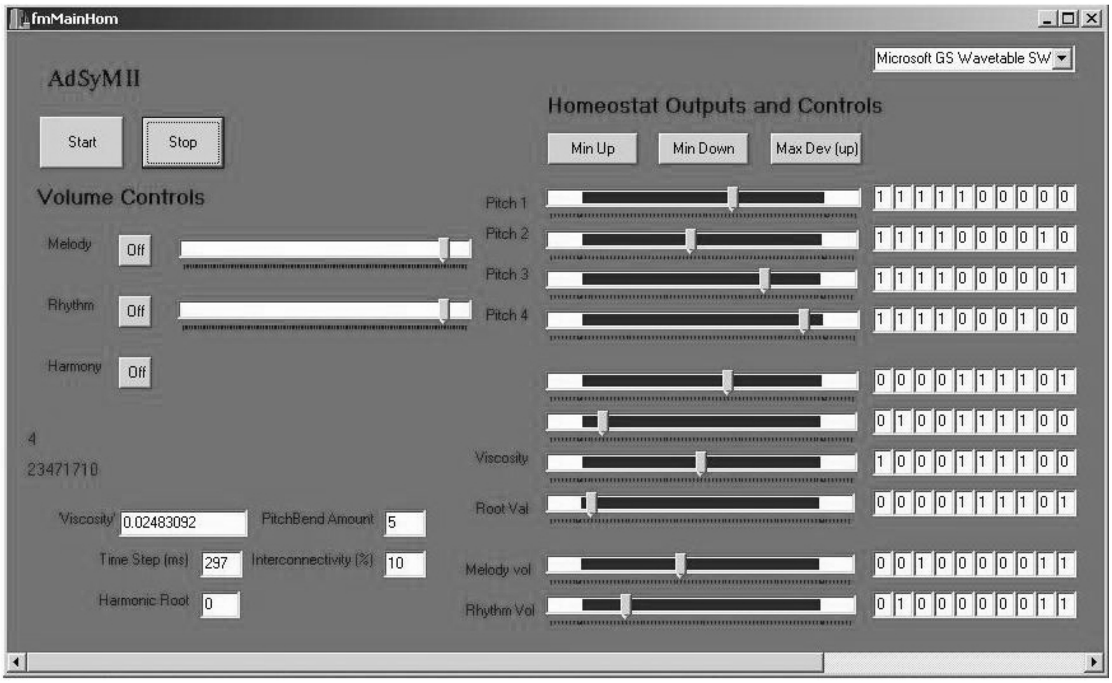

Microsoft GS Wavetable SW

4

23471710

'Viscosity' 0.02483092

PitchBend Amount 5

Time Step (ms) 297

Interconnectivity (%) 10

Harmonic Root 0

```
Pitch 1     1 1 1 1 1 0 0 0 0 0
Pitch 2     1 1 1 1 0 0 0 0 1 0
Pitch 3     1 1 1 1 0 0 0 0 0 1
Pitch 4     1 1 1 1 0 0 0 1 0 0

            0 0 0 0 1 1 1 1 0 1
            0 1 0 0 1 1 1 1 0 0
Viscosity   1 0 0 0 1 1 1 1 0 0
Root Val    0 0 0 0 1 1 1 1 0 1

Melody vol  0 0 1 0 0 0 0 0 1 1
Rhythm Vol  0 1 0 0 0 0 0 0 1 1
```

Figure B.2: Screen shot of the Autonomous AdSyM GUI

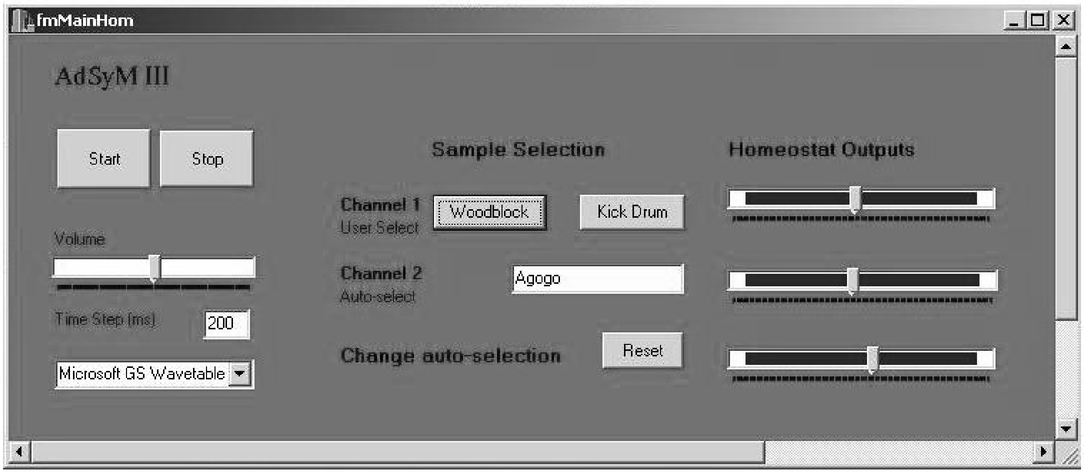

Time Step (ms) 200

Microsoft GS Wavetable

Channel 2 Auto-select: Agogo

Figure B.3: Screen shot of the AdSyM Adaptive Sequencer GUI

Figure B.4: Screen shot of the GUI presented to listeners in survey.

Appendix C

Connections for Network

```
unit no   0 1 2 3 4 5 6 7 8 9
       ------------------------
       0| 1 1 1 1 1 0 0 0 0 0
       1| 1 1 1 1 0 0 0 0 1 0
       2| 1 1 1 1 0 0 0 0 0 1
       3| 1 1 1 1 0 0 0 1 0 0
       4| 0 0 0 0 1 1 1 1 0 0
       5| 0 1 0 0 1 1 1 1 0 0
       6| 1 0 0 0 1 1 1 1 0 0
       7| 0 0 0 0 1 1 1 1 0 1
       8| 0 0 1 0 0 0 0 0 1 1
       9| 0 1 0 0 0 0 0 0 1 1
```

Units in each row receive input from units in each column set to 1.

Appendix D

Listener Survey

D.1 Instructions to Listeners

Instructions to Listeners

You will hear a continuous piece of audio. You may listen to it for as long as you wish. On the screen in front of you, you will see four buttons. While you are listening, if you particularly like or disliek a certain point, click on the “like it” or “dislike it ” buttons. These simply record hyour response, adn do not affect what you hear. You may do this as often or infrequently as you wish.
Please put the head phones on.

> Click ”start” to start listening. Click ”like it” or ”dislike it” whenever you wish. Click ”Stop” to finish.

When you have finished listening, please fill in the questionaire.

You will be given opportunity to ask questions at the end. Thankyou for your time.

D.2 Questionnaire

Questionnaire

No.__________ Age. __________ sex.__________

Please circle the number which best reflects your opinion.

	Strongly Agree	Agree	Undecided	Disagree	Strongly Disagree
1. It was interesting.	1	2	3	4	5
2. It was unmusical.	1	2	3	4	5
3. It was repetitive.	1	2	3	4	5
4. I would not describe it as music.	1	2	3	4	5
5. It was enjoyable.	1	2	3	4	5
6. I was sometimes surprised.	1	2	3	4	5
7. It was not like anything I have heard before.	1	2	3	4	5
8. It was directionless.	1	2	3	4	5
9. It was predictable.	1	2	3	4	5
10. It could have been written by a person.	1	2	3	4	5

Appendix E

Examples of Listener Comments

Listener 1.

> "In many ways reminds me of Asiatic styles of music. The repetitive nature of the piece along with the variations in the volume make it sound 'composed'"

Listener 2.

> "It was explorative and rhythmic... It oozed atmosphere (of a crazy graveyard in Iceland!)... It was playful and creative... It made me feel happy"

Listener 3.

> "It has the ability to generate mood, there seemed to be a structure."

Listener 7.

> "Tension building and resolving of tension."

Listener 8.

> "Better than any other non-human music I've heard"

Listener 15.

> "weird and surprising yet strangely familiar"

Listener 16.

> "...certainly, if not composed by a person, then must have been restricted in scale, type structure etc."

listener 20.

> "It sounded very 'soundtrack-ish', quite dramatic, I liked it..."

Appendix F

Track Listing

Examples given on the accompanying cd are as follows:

1. Simple random timing loop.
2. Random timing and pitch (pizz).
3. Random timing and pitch (sustained).
4. Homeostat loop and melody.
5. Fast homeostat and percussion.
6. Homeostat, percussion and melody (change in instrumentation).
7. Homeostat, percussion and melody operating in autonomous mode.
8. AdSyM Adaptive Sequencer in action.
9. Example of Autonomous AdSyM.

Pilot work: tracks 1-3.wav
AdSyM loop generator: tracks 4-7.wav
AdSyM Adaptive Sequencer: track 8.wav
Autonomous AdSyM: track 9.mp3