Chapter 5

Studies in Simple Adaptive Dynamical Systems

This chapter introduces some specific examples of simple adaptive dynamical systems. The term 'simple adaptive dynamical system' is far from elegant but aims to mark out a class of system that exhibit characteristics for generative and interactive activities. The *dynamical* approach was introduced in Chapter 2 where it was suggested that it was highly suited to subserving a rich continuous flow between human and machine components of a performance network. Reviewing algorithmic composition in Chapter 3, it was also suggested that formal dynamical systems could be successfully applied to create a strong sense of linear impulse and, at least at the timbral level, a perceptually strong dynamic morphology, making them useful as compositional algorithms. Consideration of the characteristics necessary to realise the richer conversational model of interaction also ear-marked the importance of *adaptation*, in providing a coherent, internally generated response. Following the examination of Alife installation art in Chapter 2, it was suggested that the many layers of adaptation in the evolving ecosystems created a system too complex and unresponsive for live music purposes, and it was proposed that the simpler, single agent systems offered a more suitable model. The *simple* qualifier here then aims to reign in the boundaries, these will become clearer with the illustrations presented below.

- The audio examples discussed in the text can be found on the accompanying DVD, tracks 1- 17.
- Max/MSP of objects for most of these models are also available on the DVD, along with help files that illustrate the basic mappings described here.

5.0.1 Models

One of the propositions of this thesis is that adaptation is not only something to be considered in an interactive context, but that adaptive responses present a useful device in generative *composition* practice. Generative art is typically discussed in terms of designing a process. Processes like L-systems or CAs can be specified, and unfold to generate a particular structure. At the other extreme, part of the fascination with evolutionary processes, is that they can create outcomes which exceed the expectations of the artist, surging off into the computational sublime. This project aims to carve a middleground, retaining the coherent unity of the L-system as it develops through time, but introducing multiple sets of parallel processes which influence each others' path. The focus of interest lies between the generation of 'structures' (as in Xenakis' interest in 'out of time structures') and 'composing interactions' with a sonic by-product (as in Di Scipio's AESI) and pushes toward composing interdependent, reactive structures ‘in time’, or more appropriately, ‘behaviours’.

A handful of models have been selected, each of which are dynamical systems of some form, and each of which exhibit some level of adaptation to environmental input which is observable as a characteristic pattern of behavioural response. These have been appropriated directly from, or inspired by, cybernetics, Alife and ecology. The systems include models of homeostasis, entrainment, pattern propagation and population distributions. Similar techniques may have application for learning, searching, or problem solving within engineering AI or AI music approaches. Here however their potential as adaptive pattern generators is explored.

Models are considered individually in this chapter and a range of mappings are explored. Consideration is given to both specific compositional goals, and the ways in which their adaptive characteristics can be employed in both compositional and performance situations. In all instances, the aim was not to replicate any specific musical style or idiom directly, but to attempt to create a sense of musical life and coherence by using systems which can be seen to exhibit some degree of goal-directedness and/or adaptation to their environment. This goal-directedness seems a first important step in creating digital systems which exhibit some degree of independence, and ultimately a musical ‘personality’.

5.0.2 Mappings

In previous chapters it has been suggested that the responsive characteristics of such models make them attractive as interactive mechanisms in live generative performance and may also provide rich dynamics that are potentially capable of generative interesting musical material. In other words that adaptation has potential as a compositional as well as an interactive mechanism. In order to explore whether this is true, a number of different mappings were examined for each model. This project departs from the approach taken by champions of algorithmic composition such as Xenakis or Roads.

In some of the most successful examples of algorithmic composition (e.g. Xenakis (1971b), Roads (2001)) formal processes were developed for specific compositional situations. In these cases the algorithm and the mapping are tightly intertwined. We could almost say that with stochastic systems as Xenakis’ *GENDYN*, what we hear is the process itself: a direct sonification of the stochastic models. In a very different way Di Scipio’s AESI also presents the process itself, although in this case it makes little sense to talk of a process distinct from its sonification. In both cases, the process has been designed with a very specific compositional aim and this aim defines the mapping.

The current project is motivated by a broader aesthetic aim: a desire to create a form of *behavioural* generative system for performance and composition. The proposal is that the dynamics of simple adaptive systems are capable of evoking a minimal sense of agency or goal directedness that invites an attribution of intentionality, or personality. As discussed in Chapter 4, certain algorithms may be more or less suited to structuring particular levels of musical material. The aim of this chapter then is to explore some different ways of mapping a range of models in order to ascertain primarily, whether any of the models are effective at all.

In order to structure the explorations, mappings were explored at different levels of complexity and at varying degrees of remove. These are summarised in Figure 5.1. In the simplest case (Figure 5.1.a), the numerical outputs are directly sonified, for example being used to specify the pitch of an oscillator. In this case the model is used directly to *generate* musical material. This approach tarries with that of a data visualisation exercise and allows immediate appreciation of the basic form of a model’s dynamics. Rather than mapping every data point, certain characteristic features can be used to generate short events. Combinations of continuous and feature-based mappings can be also be used in conjunction to create *multiple mappings* (Figure 5.1.b). This is an effective way of producing multiple parts that are closely related. These two approaches can also be applied to sample-based sonification: features can be used to *trigger* short samples, or the full data stream can be used to continuously manipulate some aspect of pre-specified sound material – for example continuously altering the playback speed of a sample (Figure 5.1.c). Alternatively the outputs can be used to *control* some other audio process acting on existing (or generated) audio, such as a filter or other effect (Figure 5.1.d).

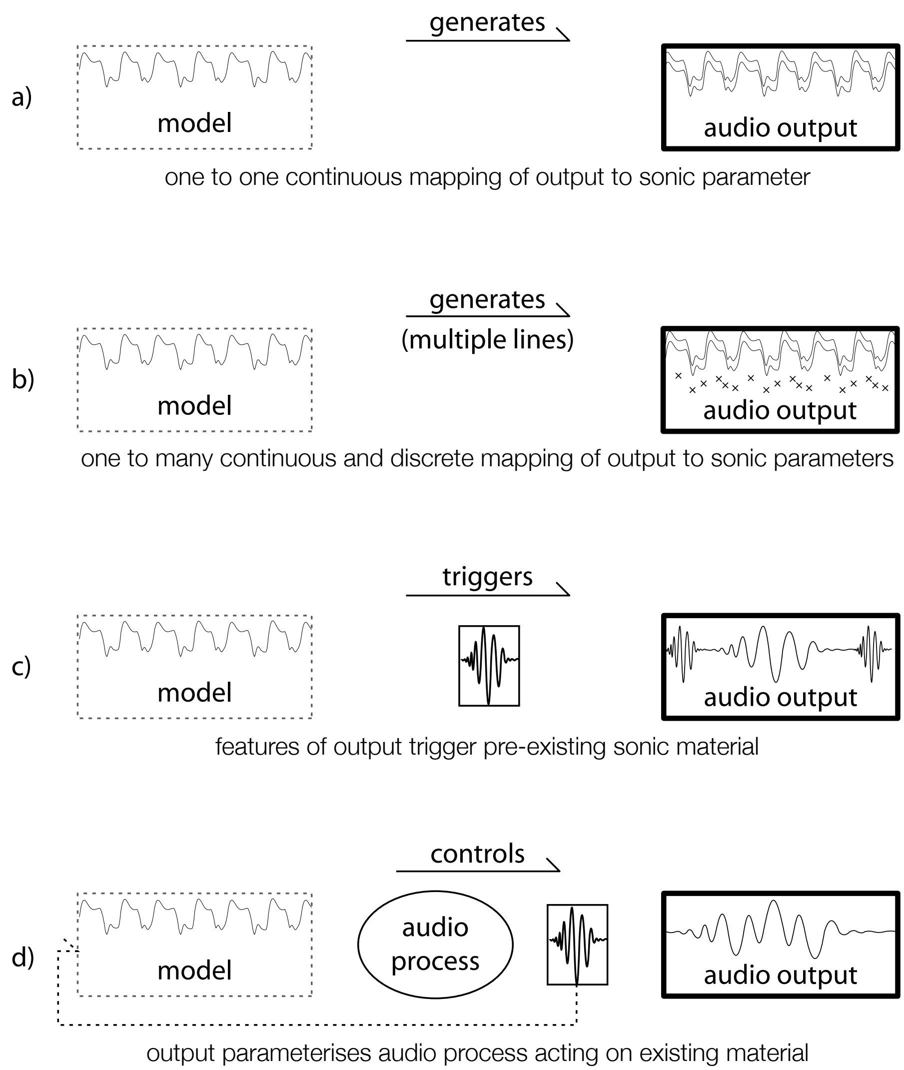

Figure 5.1: Outline of the mapping techniques explored. Outputs of the model are used to: generate material directly creating either single lines (a), multiple different but related lines (b); trigger pre-existing sonic material (c); or to parameterise some other DSP process (d).

5.1 Adaptation and Homeostasis

In the 1950s, Cybernetician Ross Ashby built an electro-mechanical machine called the homeostat. By all accounts the thing itself was an engaging machine, but its notoriety in certain circles is due to the theoretical ideas that it incarnated. One of the conundrums that preoccupied Ashby, was how a system (biological or mechanical) could be at once state determined, and yet adapt to a changing environment and learn. Ashby (1952) proposed that one of the key mechanisms underlying adaptive behaviour is homeostasis, and like all good cyberneticians, provided a concrete, physical device to demonstrate his theoretical notion of *ultrastability*.

Adaptive behaviour is a major research topic in contemporary cognitive science, and indeed the importance of homeostatic adaptation is re-emerging in philosophical circles as a key aspect for understanding of life, mind, autonomy etc. (e.g. Di Paolo (2005)). Basic homeostatic adaptation is the starting point for the current exploration of adaptive systems for interactive and generative music. Iconically and practically then, Ashby's homeostat provides inspiration for one of the central conceptual and algorithmic devices used throughout the projects presented here. The term homeostasis was coined by Canon

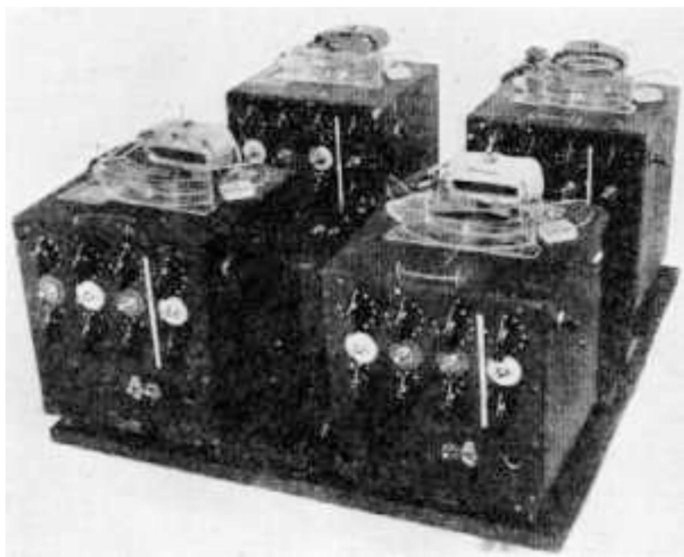

Figure 5.2: Ashby's electro-mechanical homeostat.

to describe the internal self-regulating mechanisms of biological organisms which maintain essential variables such as blood temperature, pressure and sugar levels in a dynamic balance. Cyberneticians such as Wiener (1948) and Rosenbleuth (1943) provided us with a systemic understanding of the patterns of organisation subserving adaptation and homeostasis- i.e. self-correcting negative feedback loops. The process is illustrated by every day examples such as thermostatically controlled heating systems or lavatory stopcocks and was expressed by Wiener (1948) in a characteristically wordy statement:

> "When we desire a motion to follow a given pattern the difference between this pattern and the actually performed motion is used as a new input to cause the part regulated to move in such a way as to bring its motion closer to that given by the pattern" - Wiener (1948), p.6

Ashby advanced the concept of a self-correcting feedback system in his theory of self-regulating *ultrastability*. He defines an ultrastable systems as one that is able to reconfigure plastically in response to any of its essential variables going out of bounds. In a self-correcting system the relation between the input carrying the signal error and the regulation device is fixed (the ballcock in a cistern is attached to a stiff rod connected securely to the valve). An ultrastable system exhibits a higher order stability which allows self-regulation of the regulatory mechanism itself (a cistern which could change the position of the nut on its rod, or even invert the relationship between the angle of the ballcock and the valve). Ashby illustrates the difference by inviting us to consider the mechanisms controlling an autopilot. A standard autopilot might consist of a gyroscope connected to the airelons on the aircraft wing: if the craft banks in one direction, the gyroscope measurement induces the necessary change in the airelons to roll the craft back to horizontal. If the connections between the gyroscope and airelon were reversed, the smallest bank in either direction would be amplified: the autopilot would implement positive rather than negative feedback and this would continue until the craft crashed.

The higher-order stability central to Ashby's concept of ultrastability refers to a system which would be able to adapt to, and compensate for, this reversal of connections. In this case, once the roll reached a certain critical magnitude, the connections between gyroscope and airelon would themselves invert until the roll was corrected and the aircraft restabilised. In order to achieve this Ashby argued that a system *necessarily* requires a mechanism consisting of a primary direct feedback between sensorimotor system and the environment, and a *secondary* feedback, operating intermittently at a longer timescale, between the essential variables and the sensorimotor system. It is this secondary feedback system which reconfigures the sensorimotor connections when the essential variables exceed their limits. Ashby's mechanical homeostat was a physical proof of concept for this theory of ultrastability.

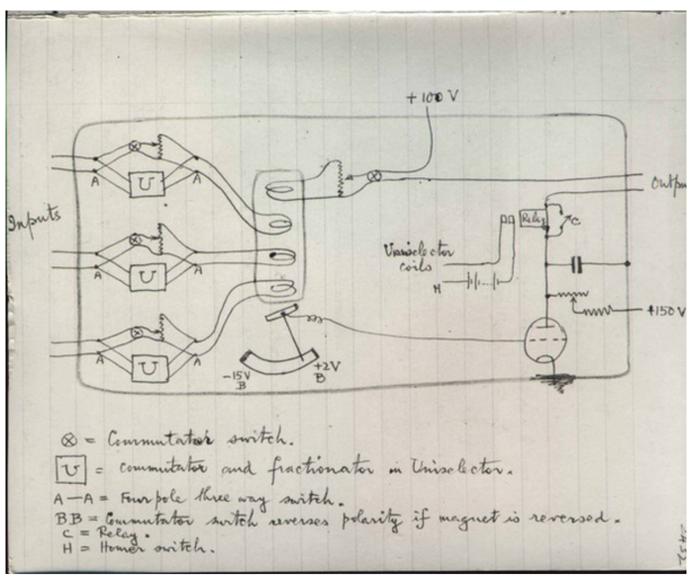

Figure 5.3: Diagram of part of the homeostat circuitry from Ashby's notebook.

The machine consisted of four units with a pivoted magnet on top of each. The angular deviation of each magnet's position representing the essential variables which were to be maintained within 45°. Each unit sends a current proportional to the deviation of its magnet from the centre (no current being sent when it is centred). This was achieved by dropping a wire from each magnet into a trough of liquid with electrodes at each end, so providing a potential gradient. The wire therefore picks up a graded potential depending upon the position of the magnet. The viscosity of the liquid in the troughs affects the behaviour of the homeostat: highly viscous liquids creating a turgid, stable system, more fluid liquids producing wilder, more fluctuating behaviours which take longer, if at all to stabilse. These electrical connections model the primary feedback, where any one unit can be conceptualised (arbitrarily) as representing either the environment, or the sensorimotor system of an agent in that environment.

The units were joined with connections between each magnet. The connections operated via coils where the torque on each was proportional to the sum of the currents in connected units. Each unit also had a recurrent connection. The current on each was modified by passing it through a commutator and potentiometer which determined the polarity and proportion of each input which is passed. These act as parameters to the system, implementing a secondary feedback which was controlled by a uniselector on each unit. The uniselector has twenty-five discrete states, each consisting of a three random values derived from a standard statistical table. Each uniselector checks the value of the outputs of its daughter unit, assigning new values to the commutator and potentiometer if the magnet's angle of deviation exceeds the critical value of 45°. The new values affect the movement of a unit's magnet, and so change the potential that is passed to connected magnets.

> "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" - Ashby (1952), pp.102-103

By a process of trial and error, the machine is able to maintain its essential variables within specified limits. Ashby also demonstrated that the machine could exhibit basic reinforcement learning, adapting to alternate environments and presented it as an example of basic self-organisation.

Wiener (1967) described the homeostat as "one of the greatest philosophical contributions of the present day" (p.54), but it was not without critics. Grey Walter (1953) dubbed it the *Machina Sopora*, suggesting that if it were to be judged entirely by its behavior, the naturalist would classify it as a plant (p.124). Another fair criticism which has been raised is that the mechanism used to achieve homeostasis (i.e. random search) is incredibly inefficient and unpredictable. In Ashby's electro-mechanical device, there are 25^4 $(390,625)$ different combinations of uniselector parameter values that a four unit homeostat can randomly explore in order to find a combination that leads to stablility. This prompted Singh's (1966) critical description of the homeostat as a 'permutational orgy'. As well as taking an incalculable length of time to stabilise, the system is incapable of accumulating adaptations, i.e. once it has achieved a certain behaviour, the stochastic nature of adaptation makes it likely to be lost irretrievably as Ashby puts it:

> "In general, if the Homeostat is given a problem A, then a problem B, and then A again, it treats A as if it had never encountered A before; the activities during the adaptation to B have totally destroyed the previous adaptation to A." Ashby (1952)

The ultimate goal of the device was to maintain consistency in the face of change, which may not seem like a very interesting musical attribute. Its indeterminacies revoke consideration of its employment as a robust learning device, particularly in a real-time situation. But the basic adaptive and dynamical process by which it achieves homeostasis is appealing.

The system illustrates the appearance of unpredictably complex behaviour arising from the interactions of simple devices. The internal adjustments made provide a minimal form of goal directed behaviour: the homeostat behaves *as though* it were seeking to keep its magnets in central positions. Despite its basic mechanical, deterministic substrate, the system exhibits open ended and unpredictable, yet coherent behaviour.

> " . . . but what strikes me about them is their singular liveliness. I can't actually think of any prior example of a real machine that would randomly - open-endedly as I would say - reconfigure itself in response to its inputs. When I think of 1950s machines, I think of lathes, drilling machines and whatever - deterministic devices that either respond predictably to commands or just break down and never work again. It seems reasonable, then, to speak of the homeostat as having a kind of agency - it did things in the world that sprang, as it were, from inside itself, rather than having to be fully specified from outside in advance." - Pickering (2002)

From a practical creative perspective, the system offers an attractive balance of autonomy and controllability. System behaviour arises from an internally controlled, open-ended configuration, but is parameterised by the degree of viscosity. Although it is 'doing its own thing', we can induce it to operate within a given field. The characteristically different responses to different forms of input displayed also provide a form of global control. Finally as will be discussed below, as a modular system, the size of the network and degree of interconnectivity have significant impact on its behaviour, and can be engineered for specific tasks.

5.1.1 A Model of the Homeostat

The key aspects of the machine were simulated in a neural-network style model. The machine is conceived as a network of I units, each connected to J other units (shown schematically in Figure 5.4) where the output of each unit is updated according to the weighted sum of the output of all other nodes as shown in Equation 5.1 (these weights modelling the potentiometers and commutators described by Ashby). In this simulation if the output of any node exceeds a prespecified value, weights connecting units in the network are re-randomised, simulating the role of the uniselectors in assigning the system parameters. As in Ashby's machine, the recurrent connection is held constant. Investigation showed that the frequency of uniselector action (i.e. testing outputs) did not have any effect on the major properties so it was held constant and outputs were checked at every iteration. Viscosity was implemented by constraining the amount by which any one unit could move between iterations.

$$\mathrm{O}_{i(t+1)} = \sum_{j=0}^{j} I_{ij(t)} \times W_{ij(t)} \quad \text{where} \quad I_{ij}(t) = \sum_{j=0}^{j-1} O_{j(t-1)} + O_{i(t-1)} \quad . \tag{5.1}$$

Where $O_{i(t+1)}$ is the Output of the i_{th} unit at time $_{t+1}$, $I_{ij(t)}$is the input to the i_{th} unit from the j_{th} and $W_{ij(t)}$ is the weight from unit j to unit i.

5.1.2 Homeostat Behaviour

This basic model is capable of replicating the principle characteristics of Ashby's homeostat. Primarily, once stable it will actively resist small interferences (the primary feedback mechanism bringing all outputs back into line), large perturbations trigger weight changes representing the secondary feedback mechanism which reconfigures the uniselector action in Ashby's machine. This is shown in Figure 5.5. Once stable, the system

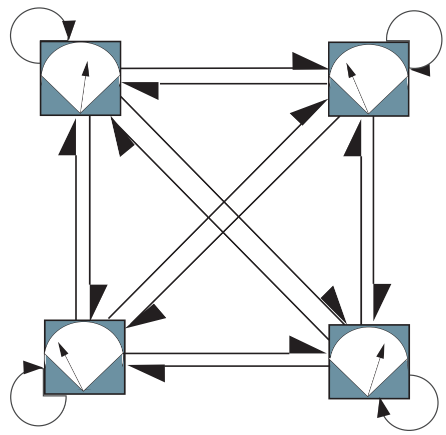

Figure 5.4: Schematic of a fully connected four unit homeostat. Each unit is represented as a square box, its output being the deviation of the small arrow from the centre. Weighted connections between units are represented by the uni-directional arrows which link each unit.

exhibits a minor transient response to perturbation below critical limits (marked **a** in Figure 5.5). At point **A**, the output of unit one was forced *outside* its critical limit. This causes weights changes, and the system enters a different stable state. Note also that when stable the system sometimes converges to a point attractor (iterations 0 - 500 in Figure 5.5), or oscillates in limit cycles, (as in iterations 500 - 1000 after critical perturbation) often each node entering a cycle of different lengths. This can be used to generate basic poly-rhythmic patterns.

In the original physical machine, the degree to which the system state was historically 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. 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. This is demonstrated in Figure 5.6 (left) which shows stability as a function of viscosity. Here stability is measured as the time taken for all units to stabilise from an initially random weight selection. In a later paper, Gardner and Ashby (1970) also discussed the effect of network size and connectivity on the stability. Figure 5.6 (right) replicates his results, showing the inverse relationship between stability and either size or connectivity of network.

5.1.3 Example Mappings from the Homeostat

Simple pitch control

The basic behaviour of the homeostat can be heard clearly if the outputs are mapped directly into pitch deviations as in Figure 5.1.a.

- In Track 1 the outputs of a ten unit homeostat control the frequency of ten sine wave oscillators, offset by a small amount to increase clarity. The initially unstable network settles with each input entering a limit cycle of a different length. This produces a minimal poly-rhythmic loop.

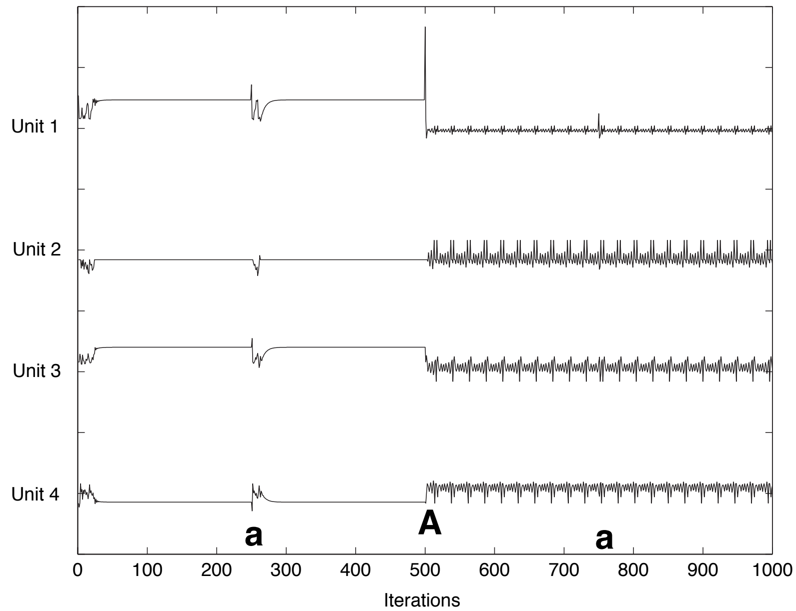

Figure 5.5: Outputs of a four unit homeostat demonstrating stability to minor perturbation (**a**) and re-stability after critical perturbation (**A**).

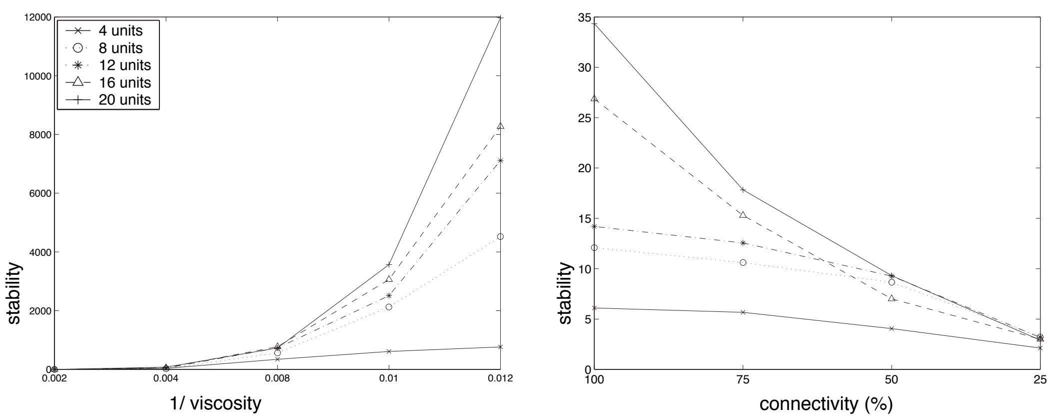

Figure 5.6: Change in stability as a function of number of units and 1/viscosity (left) and as a function of connectivity (right). Stability is taken as the point at which all units remain inside limits (and therefore weights remain constant), and measured as the number of iterations taken to achieve this state, averaged over 200 runs.

- Track 2 illustrates a similar mapping made using MIDI. Here the outputs of a four unit network are mapped to pitch bend, producing microtones of $\frac{1}{32_{th}}$ tone. The effect of applying a small input to a stable network can be heard: at around 10″, the regular pattern deviates for a few cycles and is then reinstated. This track also illustrates the effect of employing multiple mappings. As well as mapping outputs to a continuous pitch variation, a ‘melody’ line is created by using the output of each unit to specify the pitch of a percussion instrument. The timing for each unit is determined by selecting random number N for each unit in the range (2,10) and voicing its pitch every n_i beats. This creates a strange harmonised melody line.

This approach was used in the *AdSyMII* installation described in Chapter 6, and also forms the basis of a track *Sines* which was commissioned by generative film makes Iain Helliwell for the LUX Open 2002, a festival of experimental film at the Royal Art College. Here the homeostat is used to simply control the pitch of a set of four sine oscillators, giving an organic feel. This is given on Track 3.

Splicing and remixing audio samples

As well as determining low level musical attributes, the homeostat works well as a method of re-mixing existing audio material. In this example, the output range of the homeostat is scaled to the length of an audio track. The original piece is *Planting trees, creating beauty* by Norweigan trumpeter Arve Henriksen, an excerpt of which is given on Track 4. At each update, a short sample from the source material is triggered by each unit, the position being determined according to the value of each output. Rather than each output playing its selection on every beat, the sound is thinned out by specifying that some outputs only play when negative, some positive. This creates changes in density as well as changes in content. A similar mapping process was used in the Self-karaoke system and is described in more detail in Chapter 8, Section 8.2.2.

- Track 5 gives an example. The network is initially stable, each unit in a fixed limit cycle. This causes each to repeatedly play back the same few sections of the original file. At 15″ a large input is applied to unit one, triggering weight changes and causing the network to rapidly settle into a fixed point attractor. The network is perturbed again, once more settling to a limit cycle. Over the next minute, a series of small perturbations cause a sequence of deviations from a repetitive cycle until around 1′15 the viscosity is turned right up. This causes all units to rest at a similar value, all triggering the same quiet section of bowed metal. The network is perturbed once more, and the viscosity turned down, making the system more excitable, and causing it to take longer to stabilise. At 1′45 you can hear the units converge, this time reiterating a vocal sample, until the network is perturbed a final time just before the end. The homeostat is iterated at 160ms intervals giving the rhythmic pulse which can be heard.

Spectral Filter Automation

Even when in 'Machina Sopora' mode when the homeostat settles quickly to a point attractor its dynamic response can be put to good effect. Track 6 gives an example where the outputs are used as an 'automated effects' device. Here the change in the outputs of the four units are scaled, and used to control the amplitude of the first 30 bins of a spectral filter, the remaining set at zero. The filter works by performing a Fast Fourier Transform (FFT) on an incoming audio signal and splitting the signal into a number of bins. The amplitude of each can be individually controlled[1]. In this example, the audio input is Morton Feldman's *Piano Piece for Three Hands*, which is provided dry in the example along with the filtered output. Rather than applying an input by hand as in the case of the examples above, the amplitude of each attack in the piano part is analysed and used as the input to unit one of the homeostat.

- In track 6, the viscosity is set high so the system settles quickly whenever perturbed. Each attack therefore triggers a very brief period of oscillation, heard here as spectral fluctuations after each note which die out between chords as the homeostat settles. Once settled, the entire spectrum of the filter is at zero, meaning that just the dry signal is heard. Notice also that quieter notes are insufficient to trigger

[1] This is essentially a fine grained graphic equaliser like you might have on your stereo to allow you to boost or cut bass, treble, mid etc.

the homeostat: as there is no change, the amplitude of each filter bin remains at zero here also, giving no wet signal. At around 1′10 the viscosity of the system is increased, resulting in larger spectral fluctuations which continue between notes. The increased activity of the homeostat means more bins have higher values at any one time, also increasing the overall volume level.

Rhythmic generation
The incommensurate lengths of the cycles into which the outputs often settle can be used to generate regular, if lopsided, rhythms.

- Tracks 7 and 8 give examples of a rhythm generator made for a live performance at Wrong Music[2], an organisation dedicated to experimental noise music. Here each output of an eight unit homeostat is used to trigger a different drum sample. The actual value of the output being used to determine the playback speed, giving rise to the variation in pitch which can be heard. Throughout both tracks, the system was repeatedly perturbed, producing both small variations from the principle beat, and larger changes in texture. In track 8, the update time was also manipulated creating gaps and dense sputters.

5.1.4 Summary of Homeostat Features

The homeostat exhibits a number of behaviours and features which make it attractive as a generative system for music. It exhibits a range of dynamics and characteristics which can be used to generate novel but arguably evocative material. As can be seen from Figure 5.5, when it stabilises it either converges to a single point, or to limit cycles, with each output often settling of a different length cycle. This in itself can be used to create complex polyrhythms. As also shown in Figure 5.5 it exhibits different responses to perturbation: small changes causing a temporary deviation from the current attractor, which is usually returned to after brief deviation, large inputs triggering weight changes which invariably lead to the system settling on a new attractor. The viscosity variable also enables global control over the nature of its dynamics: high values creating turgid, repetitive systems, and low values creating wild searching behaviour.

The exact output therefore can never be known, but the behavioural dynamics can be controlled on a qualitative level. Weights on the recurrent connections have a strong effect on the nature of the system's response to perturbation, and general behaviour. In this implementation, these are set when a new instance of the object is made. Again, although the effect of any one set of weights cannot be predicted, the idiosynchracies of any one configuration can be learnt in a more performative way by playing with the system. As these are randomised on initialisation, the random number generator seed can be saved so that 'favourite' configurations can be returned to. Despite Grey Walter's suggestion that this machina sopora is closer to plant than animal life, these characteristics provide a balance of autonomy and responsiveness which seems appropriate for the development of interactive and generative music systems. This basic homeostat is explored within a generative music system in Chapter 6, in an interactive installation in Chapter 7 and in a performance system in Chapter 8.

5.2 Entrainment in Neural Oscillators

Since the early 1980s, neural networks have been used in algorithmic composition, but invariably employed as pattern matching or learning mechanisms. Here a continuous time model was developed and used for the generation of musical material.

[2] `http://www.wrongmusic.co.uk/`

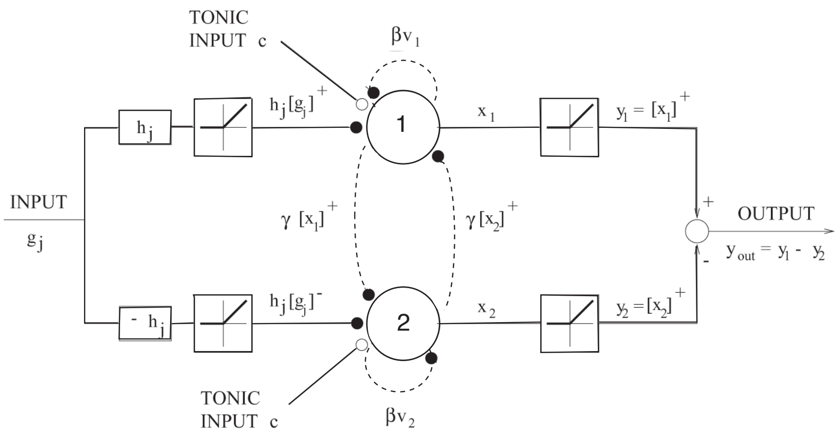

Figure 5.7: Schematic of a neural oscillator node. The oscillator equations simulate two neurons in mutual inhibition as shown here. Black circles correspond to inhibitory connections, open to excitatory. The mutual inhibition is through the $\gamma[x_i]^+$ connections ($[x]^+ = \max(x, 0)$), and the βv_i connections correspond to self-inhibition. The input g_j is weighted by a gain h_j, and then split into positive and negative parts. The positive part inhibits neuron 1, and the negative part neuron 2. The output of each neuron y_i is taken to be the positive part of the firing rate x_i, and the output of the oscillator as a whole is the difference of the two outputs.

5.2.1 A Neural Oscillator Model

Neural oscillators are continuous time, real valued neuron models, arranged in pairs such that the output of one inhibits the activity of the other, creating an oscillatory output at a fundamental frequency. If a periodic input signal is applied to the pair, it will entrain the input frequency. When nodes are arranged such that the output of one node acts as input for other nodes, the frequency of oscillation across the network will be identical, although the phase and exact shape may vary. Using simple mappings into sound, this produces musical material that shares a common pulse or metre, but varies rhythmically. This property also means that the basic pulse can be set by an external (user controlled) input signal.

Neural Oscillators have been used in robotics tasks that require rhythmic movement such as sawing (Williamson (2002)), and drumming (Kotosaka and Schaal (2001)), and in models of rhythmic entrainment (Thaut (2003)). Here, a small network of simple neural oscillators was built, based on the model described by Matsuoko (1985).

The oscillator system consists of two simulated neurons arranged in mutual inhibition, as shown in Figure 5.7. The time evolution of the oscillator is given by equations 5.2 to 5.6, where $[x]^+ = \max(x, 0)$. The output of the oscillator is y_{out}, β and γ are constants (here set to 2.5). c is a constant that determines the amplitude of the oscillation and τ_1 and τ_2 are the time constants that determine the natural frequency (in the absence of input), and shape of the output signal. Inputs (g_j) to the oscillator are weighted by gains h_j.

$$\tau_1 \dot{x}_1 = c - x_1 - \beta v_1 - \gamma[x_2]^+ - \sum_{j=0}^{j} h_j [g_j]^+ \tag{5.2}$$

$$\tau_2 \dot{v}_1 = [x_1]^+ - v_1 \tag{5.3}$$

$$\tau_1 \dot{x}_2 = c - x_2 - \beta v_2 - \gamma[x_1]^+ - \sum_{j=0}^{j} h_j [g_j]^+ \tag{5.4}$$

$$\tau_2 \dot{v}_2 = [x_2]^+ - v_2 \tag{5.5}$$

$$y_{out} = [x_1]^+ - [x_2]^+ \tag{5.6}$$

5.2.2 Neural Oscillator Behaviour

If an oscillatory input is applied, the node will entrain the input frequency i.e. it will produce an output of equal frequency, but not necessarily the same phase, as the input. This can be shown to be true over a wide range of input amplitudes and frequencies (see Appendix A, Figure A.1 for an illustration).

The fundamentally dynamic nature and specific behaviours associated with its entrainment properties make this model an attractive resource. The input driving signal can be given either by an external source, or from another software system, making it a useful component for the modular approach adopted here. Networks of oscillators exhibit a range of musically-relevant behaviours which are parameterised by a handful of variables. The sonic effects of changing these parameters is of course determined in part by the mapping and is discussed below in a simple case. In general, the fundamental frequency and form of the output can be controlled by the two time constants (τ_1 and τ_2). If run at audio rate, this can be used to generate audio signals directly. Iterating the model at slower speeds enables the generation of either melodic or rhythmic lines according to mapping scheme adopted. The entrainment property means that networks of these modules can create material of chosen degree of density, where each part bears a global relation to the whole. This creates parallel streams of data which retain their individual identity over time, but move in relation to each other.

5.2.3 Example Mappings from the Neural Oscillator

Pitch control

One of the simplest, and perhaps most effective, methods of sonifying this system is to simply map the output value of each unit onto a pitch value. When the bias of each oscillator node is between zero and one, the output will always be in the range (-1,1). This means the output can be easily mapped onto pitches in a chosen audible range. Figure 5.8 shows an example where the pitch has been quantised to semitones. The scored notes represent the waveform within the dotted box above.

The periodic oscillation of the node produces a basic arpeggiated effect. Under this mapping, changing the constant c varies the amplitude, and so pitch range of the line. Quantising the continuous output means that small changes in output, as well as fixed values result in a constant pitch. In the example shown in Figure 5.8 these repeated values were excluded, automatically introducing some rhythmic variation. The time constants affect the fundamental frequency of oscillation as well as its form, so can be used to alter the melodic contour of the output. Changing the absolute value of the weight between nodes as well as its sign determines the extent and nature of the influence of each node on connected nodes, changing the relations between parts.

Figure 5.8: Mapping from a continuous output to quantised notes. The section of score represents the graphed output within the box and *above* the horizontal line only. Notes are only re-voiced if they have changed by more than a semi-tone across timesteps, creating the spaces shown here as rests.

- Track 9 gives an example of the basic arpeggiated line generated as well as the effect of inverting the weights between nodes. In the example here, the outputs of two nodes with the same bias but slightly different time constants are played on two pianos. Initially node two is played alone after five cycles (20″) the second piano enters. The weights are negative, causing the outputs to be in opposite phase, creating a sense of turn taking. The weights are then inverted at 55″ causing both parts to play in unison.

.

Applying an external input can have several musically useful effects. Primarily of course, if above a certain amplitude, it will determine the overall frequency of the system output. Continuous periodic input (such as a sinusoidal function) of low frequencies clamps the outputs of strongly connected nodes during the positive or negative parts (depending on the polarity of the weight). This causes the output to freeze at a particular value, being 'released' when the amplitude of the input drops. Sonically this creates the effect of a line pausing, or resting on a pitch, then 'coming back to life'. Finally although the external input entrains the overall frequency of output, characteristics of the fundamental oscillation are preserved. This produces an inner pattern which is modulated at the period marked by the main input.

- Examples of these effects can be heard on track 10. Again there are two voices here, a piano and a sustained synth sound. Initially the synth is clamped, repeating the same note. Once it comes in it takes a simple descending four note motif, which is modulated by the input frequency, altering the pitch of some of the notes in the internal structure. Here the synth sound is triggered only at local minima rather than continously, creating a bass line feel.

- Track 11 gives an example with four parts playing and demonstrates the effect of altering the input amplitude and frequency. At the start, four nodes are connected with different time constants and biases, giving each a characteristic shape. There is no input signal, so the frequency is determined internally by the nodes. From 20″ - 60″ the amplitude of the input signal is gradually increased. This has a differential effect on individual units depending upon how closely they are connected to it, and how strong their weights are. At 1′10, the frequency of the input signal is decreased, the longer period clamping the outputs. Here repeated notes are omitted so this audibly this thins out the parts. Finally at 1′50, the input is removed and the ensemble returns to its initial repetitive cycle.

This melodic mapping was used in the installation *Organised Entry* which is described in Chapter 6.

Rhythmic mappings
The network can also be used to generate rhythmic patterns by defining certain points in each oscillation to trigger a percussive voice. If a number of nodes are arranged in series, with an external input, the frequency of oscillation is constant, but the oscillations may vary in shape or phase. This provides a means of generating layers of rhythmic patterns with a constant metre, or pulse, but with a much greater freedom in the placement of individual beats than is common in most computer music. These discrepancies in timings can bring a human feel to the output, akin to expressive deviations from the beat. Equally however, deviations can make the output simply sound ‘out of time’. In these examples, standard GM percussion instruments are triggered at either local minima, local maxima or at zero-crossings. An example is shown in Figure 5.9.

- A simple rhythmic example is given on track 12 with successive nodes voiced, demonstrating the possibility for generating conventional rhythmic patterns.

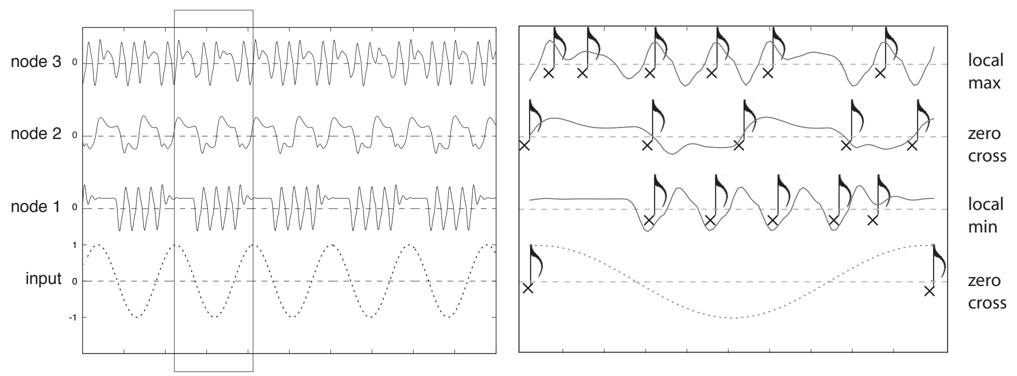

Figure 5.9: Outputs of three nodes in series (left) and detail, showing beats triggered (right): sinusoidal input and node three are triggered at local maxima, node two at zero-crossings (falling and rising) and node one at local minima.

- A more interesting example is given on track 13. There are four nodes in series, weights between each node, and biases are held constant. The different rhythms are produced by changing the time constants of individual nodes. A change is made half way through to one of the connecting weights, demonstrating how the basic beat is preserved, whilst varying the ornamentation and altering the stress.

5.2.4 Summary of Neural Oscillator Features

The basic oscillatory patterns of the nodes mimic the wave-like structures of many melodic and phrasal structures in instrumental music. The continuous nature of the outputs provides scope for mapping to a range of musical domains. Control of individual parts is made possible by altering the gains and time constants, creating variations across components in a network operating at a unified frequency. Altering the weights between the nodes obviously also gives control over the relationship between constituent parts. The same mechanism could be applied in interactive system with a performer - positive weights on the input causing the system to spring into life when the performer plays, negative weights causing their playing to inhibit the system which would only play when they are silent. In artificial neural networks, these weights model basic mechanisms of inhibition and excitation which are fundamental to neuronal communication. Musically this process can be used to mimic basic modes of interaction between musical parts or players i.e. unison or contrary motion

In the examples above, the entrainment property of the neural oscillator is utilised to provide a metrical unity across each rhythmic part. This property also provides an implicit beat detection mechanism that can be used to set the network outputs to a user-defined pulse. A beat interval supplied via a MIDI (or ASCII) keyboard, or analysis of audio signal can be used to set the frequency of the input signal, to which the rest of the network entrains. For certain settings, changes in the input frequency change the shape of the output signal. The result is a system that can keep time with a human player, but will produce novel, unpredictable rhythmic variation.

This simple network, even with hand set parameters can be used to generate intruigingly musical outputs, with connected nodes creating a sense of ensemble. One of the immediate draws backs of this implementation is the incessant nature of the output. The melodic mapping described above was used for the installation *Organised Entry*, presented in Chapter 6, but combined with another system which acted as a mixer, controlling the entries of individual units in the network.

5.3 Pattern Propagation in Cellular Automata

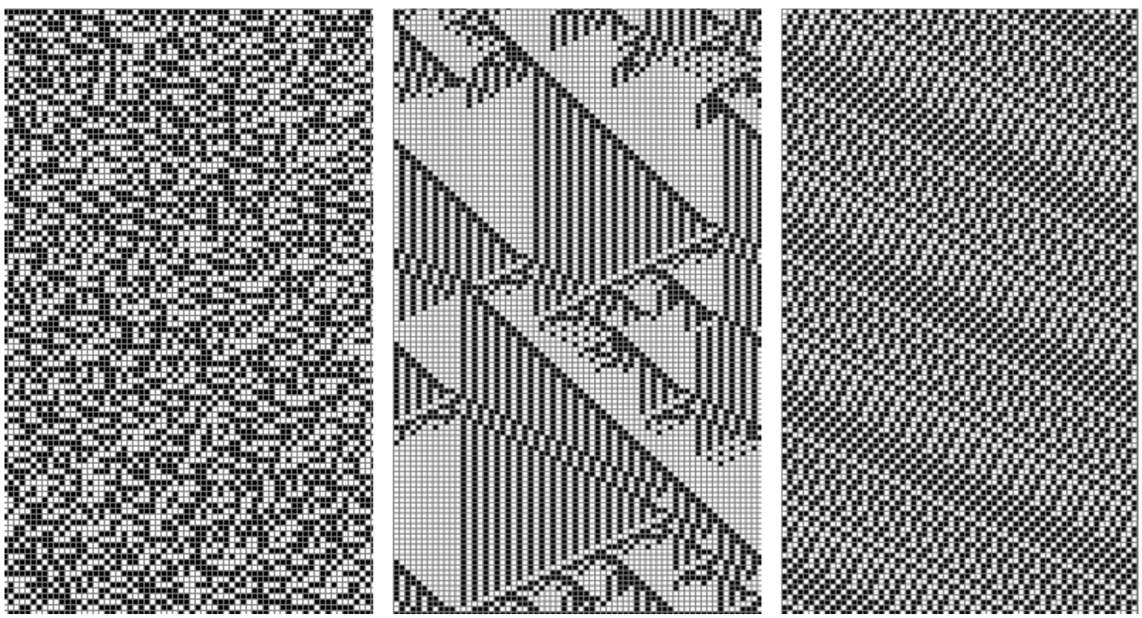

Figure 5.10: Graphical representation of 1D CAs: chaotic (left), complex (middle) and ordered (right).

Cellular automata are amongst the most used Alife system, their pattern propagation properties being an attractive means of generating low level structures. As described elsewhere, they have been used to specify pitch information, as well as to control signal-level parameters for sound synthesis. A basic description of CA is given in Chapter 3, Section 3.1.2. In the current project, the different forms of patterns generated by CAs have been used to generate different rhythmic textures. The main mapping employed is that described in Chapter 4, Section 4.3.4 which creates pitched rhythmic patterns from their output. Further examples are given on tracks 14, 15 and 16 which correspond to the graphical representations of the rules (chaotic, complex and ordered) shown in Figure 5.10.

- Track 14 gives an example of the chaotic rule set shown in Figure 5.10 (left). The random distribution of black and white cells creates an almost continuous spattering rhythmic pattern. Recall that the mean and variance of the frequency distribution of the look up table are used to determine the root pitch, and the size of the intervals of the triad above this pitch respectively. In chaotic rules, each configuration is equally likely, meaning that the frequency distribution of the look up table is flat. This creates small intervals both between iterations and a low variance across each iteration. This is heard as sets of close chords which vary minimally

- Track 15 gives an example of the complex rule set shown Figure 5.10 (middle). Here the localised patterns evident in Figure 5.10 can be heard as broken blocks of regular pulses. The mixture of local areas of order and higher level more complex forms mean that the same rules are used repeatedly for a period, then change. This means that the frequency distribution of the use of rules is skewed in any one iteration, and varies over iterations. These high variances result in the wider chords which can be heard, as well as the larger changes in root note which signifies the start of each line.

- Track 16 gives an example of the ordered rule set shown Figure 5.10 (right). The mapping employed creates short rising phrases from the diagonal stripes with a regular rhythm. As an ordered rule set, the same individual rules are used over and over again, preserving the same set of pitches. Here only two different rules are used on alternate iterations, producing the alternating root note which can be heard.

5.3.1 Summary of CA Features

The discrete patterns formed by 1D CAs provide a mechanism for generating strongly rhythmic patterns. Under the mapping used here, although there is no metre imposed, the patterns propagated produce distinct patterns of stress which punctuate the low level events with structured accentuations. From a compositional perspective, the distinct rule classes provide a means of varying the rhythmic complexity or accessibility. Rather than mapping only the immediate state changes onto musical events, the use of changes in statistical properties of the process can be used to relate different musical dimensions.

A CA was used in conjunction with a homeostatic network in *AdSyMII* described in Chapter 6. CAs are usually seeded and left to run, but it is also possible to interfere with the state flow by changing the state of cells in current influential neighbourhood. This could for example cause an ordered rule to diverge, an interruption from which it may or may not recover. A similar principle was explored in the individual based ecology model described below.

5.4 Ecology Models

The models presented above predominantly have application in generating material. Other models taken directly from, or inspired by population modelling have also been explored as a means of controlling parameters in, or orchestrating, other systems. Two different classes of model were examined, an individual-based model and a set of coupled differential equations.

5.4.1 N-species Lotka-Volterra Model

The Lotka-Volterra model (Lotka (1925), Volterra (1926)) appears in all undergraduate textbooks as the simplest ecology model that describes predator-prey relationships. It consists of two coupled differential equations as shown in Equation 5.7

$$\frac{dF}{dt} = F(a - bS) \text{ and } \frac{dS}{dt} = S(cF - d) \tag{5.7}$$

Where F is the number of prey (rabbits, small fish, flies, etc.) and S is the number of predators (foxes, sharks, spiders etc.), a = reproduction rate of prey, b = predation rate, c = reproduction rate of predators (per prey eaten) and d = death rate of the predator.

For any positive values of a, b, c and d the system oscillates in a limit cycle. In ecological terms this is incredibly over simplistic, as no ecology consists of only two species, but is made up of numerous trophic levels connected in a complex food web. For the current purposes the ecological validity can be ignored and the potential dynamics of the system increased by creating a model for N species.

There are many ways of generalising the basic Lotka-Volterra equation. The one employed here was developed by Arneodo et al. (1980). In contrast to the simple limit cycle exhibited by the two species Lotka-Volterra model, the n-species model used here exhibits a broader range of dynamics for a larger number of species which is readily parameterised. The system of ordinary differential equations for n-species can be re-written as:

$$\frac{dx_i}{dt} = x_i \sum_{j=1}^{n} A_{ij}(1 - x_j) \tag{5.8}$$

where x_i represents the ith species and A_{ij} represents the effect that species j has on species i. The A_{ij} terms can then be represented as a matrix. For three species the values can be defined as:

$$A = \begin{bmatrix} A_{11} & A_{12} & A_{13} \\ A_{21} & A_{22} & A_{23} \\ A_{31} & A_{32} & A_{33} \end{bmatrix} = \begin{bmatrix} 0.5 & 0.5 & 0.1 \\ -0.5 & -0.1 & 0.1 \\ \alpha & 0.1 & 0.1 \end{bmatrix} \tag{5.9}$$

where α parameterises the whole system.

As shown in Figure 5.11, the system exhibits a range of dynamics which are controlled by the α-value. Low values ($\alpha \leq 0.75$) cause the system to converge on a fixed point attractor, at higher values simple periodic behaviour emerges. Increasing the value beyond this causes period doublings until at around $\alpha = 1.5$, the system exhibits chaotic dynamics.

.

- Track 17 provides a simple example where each of the three outputs are mapped to the playback speed of three different versions of the same sample. The recordings are of Inuit caribou ladies babbling. At an initial α value of 0.75 the system is converged on a point attractor, and each sample plays back at normal speed. As α is increased to 1.2 you can hear the simple periodic behaviour emerge as uniform oscillations in the playback speed; the period doubling evoked at $\alpha = 1.4$ gives a double loop, and at $\alpha = 1.5$ the chaotic dynamics create chaotic pitch changes.

In itself the behavioural repertoire of the GLV model is perhaps a little limited, but mechanisms like this are a useful addition to the compendium of objects. This model was used as a mixing device in the *Organised Entry* installation described in Chapter 6.

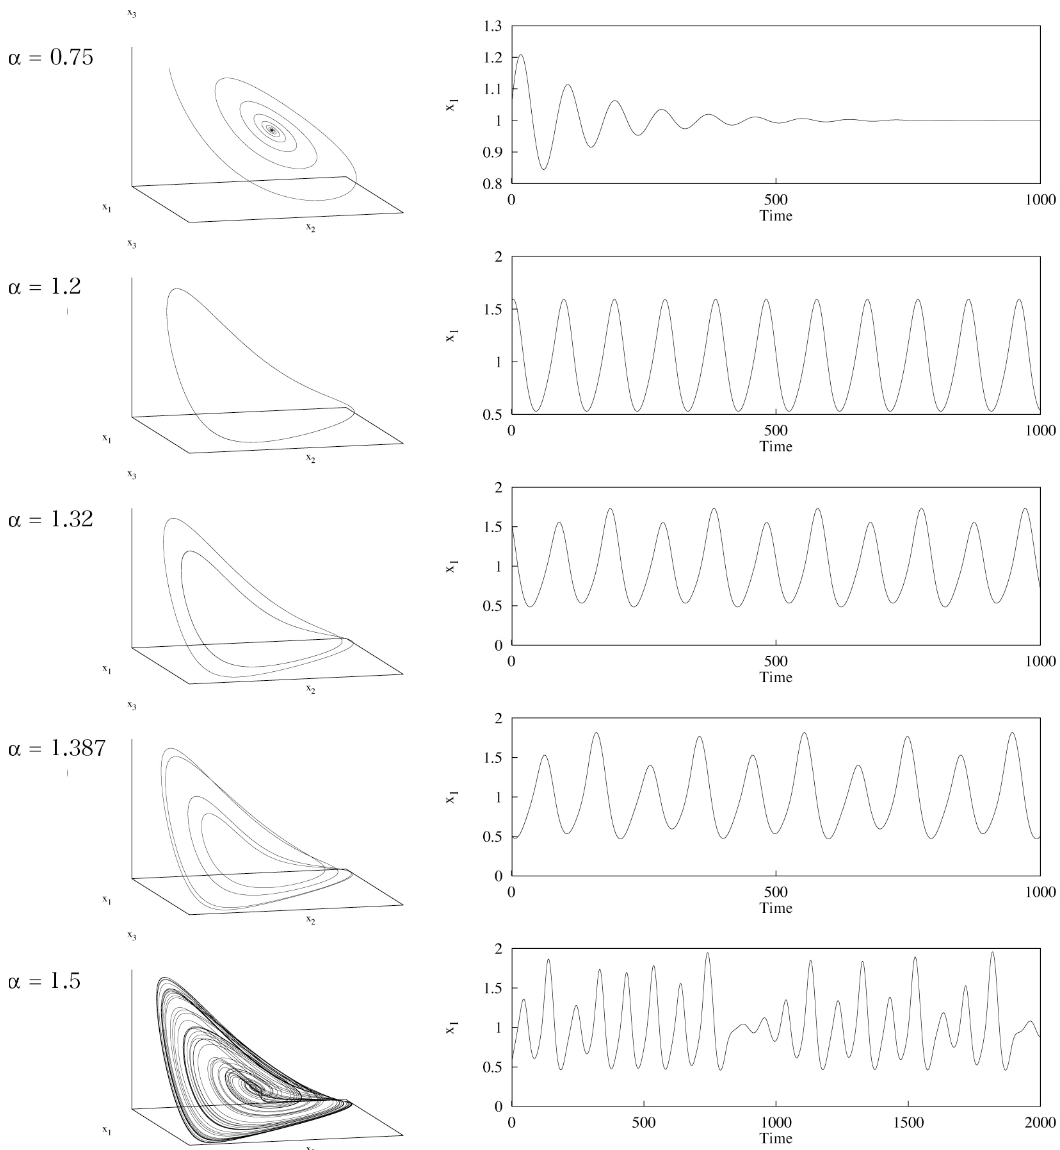

Figure 5.11: Period doublings in a three-species Lotka-Volterra system: phase space on the left and x_i on the right.

5.4.2 A Simple Agent-based Model of Spectral-temporal Organisation

Evolutionary agent-based models are often used in Alife art and music as a means of generating diversity, and exploring more open ended behaviours. This simple model serves to illustrate how systems can be contrived to fulfil specific methods of control and organisation. Specifically, the model aims to provide a mechanism for distributing spectral and temporal features of events into unique niches. In contrast to the traditional application of GAs as a means of achieving a single honed individual, this model aims to achieve specific properties at the population level.

Within music, whether a symphony score, R and B track or improvisation group, the function, value and significance of each part or player only makes sense relative to every other voice. A central aspect of composition is in balancing the various lines such that they each occupy their own unique space. Many composers, particularly of the acousmatic tradition, draw inspiration from the organisation of sound in the natural world. Bio-acoustic studies of natural habitats suggest that each organism occupies its own sonic niche both in frequency/ spectral domain and in time. Results from studies performed in Sequoia national park (Krause (1993)) support this hypothesis [3]. If one creature stops vocalising, another joins the chorus, keeping the bio-spectrum intact. This idea is supported by reports that point to the disruptive effect of human industrial noise on populations of local wildlife. For example the population decline of birds living in areas of motorway development has been attributed to the noise of the traffic preventing communication and therefore mating (Barot (1999)).

Drawing from these observations a simple model was implemented to investigate whether a self-organising mechanism, based on the premise that sound objects could only persist if they occupied a unique spectral/temporal niche, could be used to organise a randomised set of pitch-time values into unique and stable spectro-temporal niches.

The Model

The model is a simplified version of those used in individual-based ecology models (e.g. Epstein (1996), Forrest and Jones (1994)). The system consists of a population of agents which are defined by their pitch (P) and vocalisation time (VT) values. Based on the premise that individuals in any one species can only reproduce if they can hear each others' mating calls, reproduction can only occur between individuals of the same pitch if they share the same VT value AND no other individuals of any other species hold this value.

A population of agents is initialised with pitch and VT values selected from a uniform random distribution over the intervals $[1, 10]$ and $[1, 100]$ respectively, and an energy level. There are currently no spatial dimensions, and no resources. Time is discrete, and each iteration consists of N timeslots, during which individuals vocalise. At each timeslot t_n any agent with vt value n, produces note p. If that timeslot is uniquely occupied by agents with the same pitch value, reproduction occurs. Half the number of agents with coincidental values are produced. Offspring inherit the parental P value which remains fixed. VT values are inherited and mutated, using creep mutation with wrap around, with a probability of 0.1. Energy levels are reduced for all agents on every iteration according to whether or not they reproduced: taxes for those that did not reproduce are twice those that did. When energy levels reach zero, the agent dies.

This mechanism alone was sufficient to produce populations which inhabited unique pitch-time spaces, but in the absence of any external resources, additional factors were required to curb the population and introduce novelty. A global population maximum

[3] The team suggest that the biophonies of natural habitats can be used as a measure of the health, or stability of an environment: the more clearly demarcated each species is in spectro-temporal map, the more stable the system.

was set. When this is reached, no reproduction can occur until some agents die out. This produces periods of stasis. A maximum is also set for each pitch class. When this is reached, the number of agents of that pitch is reduced to X% of the remaining population by randomly culling individuals.

In the absence of any external resources, or reproductive mutation of pitch values, the system is extremely sensitive to the initial distribution in terms of the number of agents that can reproduce. Once a pitch class dies out, there is no possibility for it to re-enter the population. An extreme example, with only one pitch class is shown in Figure 5.12. Drawing from the observation that in natural acoustic ecologies, when one spectral niche is freed, another organism adapts its call to fill the gap, here when the number of pitch classes (species) drops below a certain threshold, pitch values of the remaining members of the population are mutated with a low probability at each iteration for the remainder of their lives, introducing life-time variability in pitch.

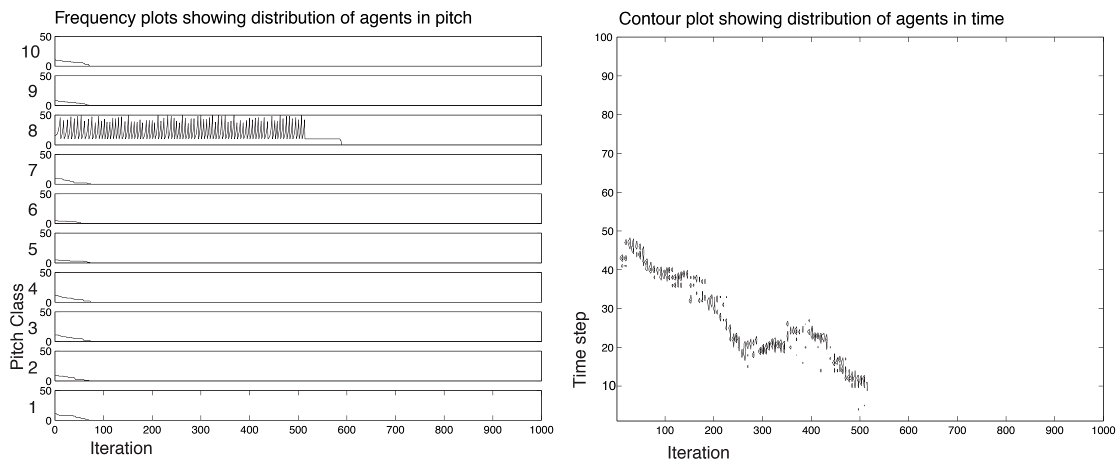

Figure 5.12: Figures showing distribution of agents in each pitch class (left) and across time slots (right) for system with no lifetime pitch variation.

The system described above is capable of organising an initial random population into subgroups that occupy unique areas of pitch-time space. Figure 5.13 shows the initial and final distributions of a population in pitch-time space in a typical run. Figure 5.14 shows the movement of the population over 200 iterations in pitch (left) and time (right). Even in this simple model, it seems that the reproduction criterion (in conjunction with the restraint thresholds) is sufficient to produce populations that are stable - in terms of neither dying out nor overcrowding - yet dynamic in producing movement of sub-groups through pitch-time space.

The model also enables external manipulation of the population dynamics. In Figure 5.15, four agents of pitch class five, onset time 80 were introduced are iteration 250. The system was started from the same initial seed as that shown in Figure 5.14, demonstrating the potential for a user to change the course of the evolution of the system.

Summary

The model presented here is extremely simple, and in its current form, the sounds produced are far from interesting musically. However, it suggests that population distributions can be controlled according to simple reproduction restrictions. The reproductive success of each agent is a function of the global environment, which comprises the behaviour of every other agent. This produces a unity between musical output (which is

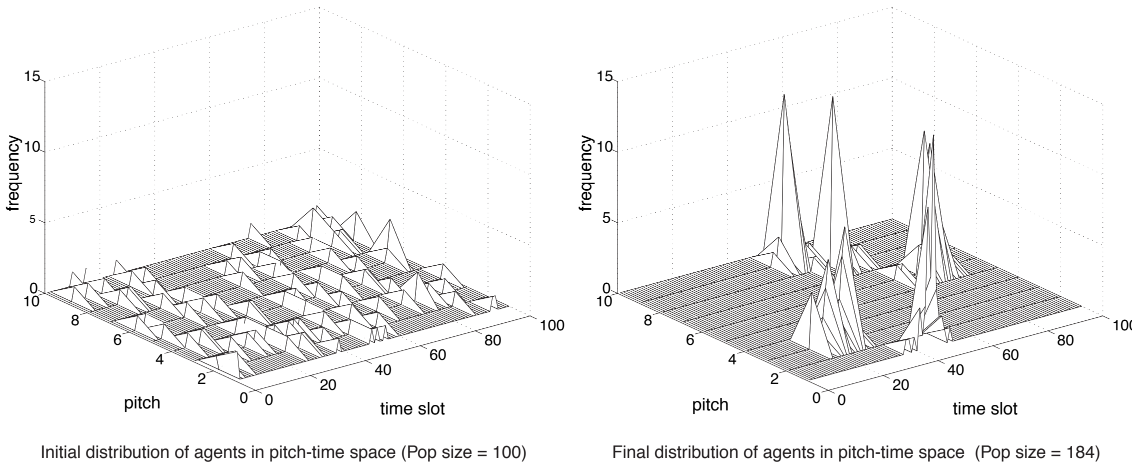

Figure 5.13: Figures showing initial distribution of agents in pitch-time space (left) and final distribution (right).

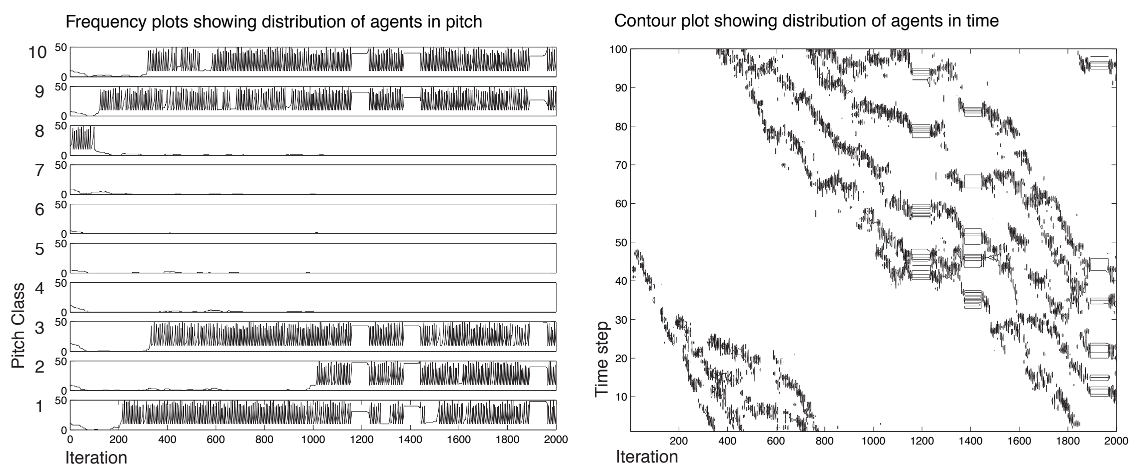

Figure 5.14: Figures showing distribution of agents in each pitch class (left) and across time slots (right) for system with lifetime pitch variation.

the collective behaviors of all elements), and system state. The evaluation of all parts is inherently a dynamic function. This possibility for establishing a coherence *across* a population offers an interesting approach to generative music systems that contrasts with existing evolutionary approaches where the focus is on getting a small subpopulation to achieve a certain criterion, or the pairwise testing of coevolutionary models.

5.5 Implementation

All these models were first developed in C++, and their basic behaviours examined. Where appropriate their response was compared with previous implementations. The homeostat, neural oscillator and GLV equation were then developed as Max/MSP externals so that they could be used within this environment. This makes the exploration of different mappings very swift compared to coding the equivalent DSP or MIDI mappings from scratch. The CA and agent-based model were developed as stand-alone Windows applications.

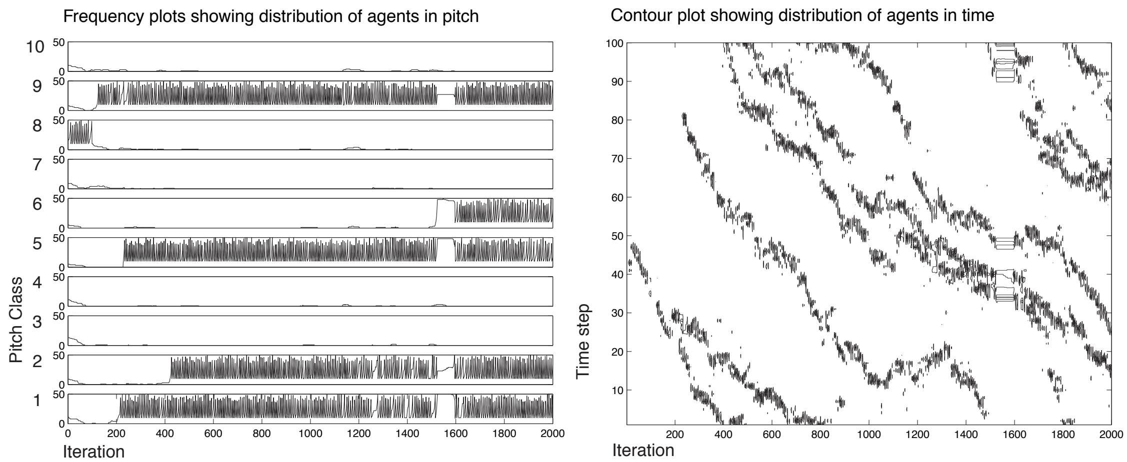

Figure 5.15: Figures showing distribution of agents in each pitch slot (left) and across time slots (right) for system with lifetime pitch variation. At iteration 250, four agents of pitch class five, onset time 80 were introduced

5.6 Discussion

Although very simple, these studies demonstrate the compositional potential of simple adaptive systems in terms of both the sonic effect of their dynamics under different types of mappings, and the practical impact of adopting algorithms that we can only influence rather than directly control.

The homeostat and neural oscillator models in particular both generate a set of evocative behaviours. The interdependencies between the outputs of separate nodes in the networks arguably creates a sense of dynamic cohesion under continuous and multiple mappings in which each part is audibly related and influence is decentralised. Under quantised, rhythmic mappings textures are created that have a strong sense of pulse in the absence of any rigid metrical constraints.

Both systems are parameterised by a handful of variables which allow the user to shape their behaviours whilst retaining the generative independence of the model. This is useful in both compositional and live situations. In addition, both respond to external influence which can be applied manually or algorithmically, evoking a contingency which goes beyond button pressing. The response of the homeostat to perturbation provides an interesting form of control by which we can suggest that 'something' happens, leaving the details of what that 'something' is to the algorithm. When mappings are designed to take this into account, this can create some enthrallingly organic deviations from, and recapitulations to, previous material.

This rhythmic interpretation of the CA takes advantage of its inherent pattern propagation properties, and the use of multiple mappings here, as well as in track 2 of the homeostat demonstrates how single models can be used to generate sets of independent but related musical lines.

Many of the mappings used here act to quantise the continuous outputs of some of the algorithms. For example in using the NOSC outputs to trigger MIDI notes, much information is being thrown away. In some respects, the true, continuous, dynamic nature of the models is only preserved under mappings such as that used for the Lotka-Volterra system. This is an example of the *control* mapping outlined in Figure 5.1.d – in this case the algorithm's three outputs were used to continuously alter the playback speed of three versions of the same sample. These sorts of mappings are perhaps most typical in Sonic

Arts domain where algorithmic composition is popular. The use of neural models to generate arpeggiated forms typical of classical or early electronic music may seem to be a mixing of worlds, but it is precisely this synthesis of difference that characterises the cyber-nature aesthetic of Alife visual art, a synthesis which I am interested to evoke in the sonic domain.

From a systemic perspective, perhaps the most interesting mappings are those that go beyond a simple one-way number-to-note formula and feedback into the system. This was demonstrated in a very simple case in Section 5.1.3. In this case the algorithm is controlling a filter process operating on an existing sample, but the sample is itself affecting the homeostat. This stitching-together of algorithm and implementation is a promising direction for a more collaborative approach to interactive and generative composition and performance and will be pursued a little further in Chapter 8.

The mappings used here have been developed for illustrative purposes. The main thrust of this thesis is to lay the ground work for a more collaborative form of man-machine musicianship, a collaboration in both systemic terms – such that human and algorithm are mutually influential – and a collaboration in aesthetic terms – such that the vagaries of algorithmic composition play out alongside the established acoustic traditions. The implementation of the algorithms in the form of software objects that do *not* impose any restrictions on the way in which they are mapped is also quite intentional. This is in line with the modular approach central to the musical communities that are developing around software such as Max/MSP. In developing algorithms in this way, it is hoped that other musicians can adopt these context-free algorithms for their own compositional ends.