

Adaptive Systems Music: Musical Structures from Algorithmic Process.

A.C.Eldridge, BSc, MSc
*School of Cognitive and Computing Sciences,
University of Sussex, Sussex, UK
e-mail: alicee@cogs.susx.ac.uk*

Abstract

The generation of large scale structures at the level of musical form represents a major challenge for current approaches to generative music. This is especially true for approaches aimed at generating ‘new music’, which necessarily employ organisational principles other than those prescribed by traditional music theory. The Adaptive Systems Music project is an ongoing investigation of the use of biologically inspired self-regulating algorithms as compositional tools. This paper describes an autonomous algorithmic music system, designed to generate polyphonic audio. The main component is a self-regulating homeostatic network, based on a simulation of cybernetician Ross Ashby’s electro-mechanical homeostat. Part of this network is used to generate ‘homeostatic harmonies’, the system dynamics producing an interesting musical structure. A rhythmic part is derived from the states of Cellular Automata, and a ‘melody’ picked out of the harmonies using a simple stochastic technique. The music generated has received high commendation in public appearances; listeners in a more rigorous evaluation conclusively supported its musicality. Their comments suggest that this may be attributable to the internal structures generated. It is suggested that exploration of possibilities in the sonification of simulated systems offers a fulcrum for sci-art collaboration and has potential to aid advancements in theory and practice in both domains.

1 Introduction

It has been suggested that one of the great achievements of AI is the construction of machines that can create music of high quality [1],[2]. However, the most successful approaches are dedicated to replication of music of an established style. By training on a corpus of work, or embedding stiff constraints derived from traditional music theory, convincing music in the style of Mozart [3] or traditional Jazz [4] has been produced. Attempts to compose ‘new’ music by computational methods *without* conventional constraints have been less successful.

Part of the problem of composing new music by computational methods is the inherent lack of cultural reference [5]. Human music is infused with cultural conventions and traditions that create expectation on the part of the listener and imbue the music with certain forms of ‘deep’ structure. Some of these traditions have been formalised, as music theory: the rules of harmony and form which govern the structures within classical, liturgic and other complex musical styles. It is these formal descriptions that have been embedded as constraints in systems such as Cope’s EMI [3] and provide a certain ‘musicalness’ to the digital output. More primitive elements evade formalisation, being processed subconsciously. It is perhaps the absence of this level of cultural reference that differentiates Mozart’s works from EMI’s works-in-the-style-of-Mozart.

“...when music is primitively compelling it forces its way into the conscious; when music is structured in a way that engages analytical and predictive processes, it invites itself into the conscious.” Keane, [6] p.108

If these more primitive elements currently evade computational methods, and ‘new music’ necessarily does not subscribe to the constraints and structures of traditional music, are there any possibilities for generating audio that is sufficiently musical to engage the human mind?

One alternative to traditional methods that has been explored is the deployment of abstract models thought to embody some of the dynamics of the compositional process such as combinatorial systems, stochastic models, fractals and other iterative processes [7], [8], [9]. The potential for Artificial Life style models has also been recognised, and the use of cellular automata in the generation of musical structures has been explored eg [5]. Although these approaches have produced some interesting musical material, these consist largely of short fragments of monophonic note sequences. There are few examples of large-scale form, and nothing with any kind of ‘rich internal structure’.

Another alternative is to pursue a longstanding preoccupation with the emulation of natural structures. Since the time of Pythagoras, we have been creating rationales for musical structures based upon “facts of nature” [6]. The potential for simulation of natural processes and structures has increased with advances in computational power, and it has been suggested that it may now be possible to develop algorithms of sufficient sophistication to create resonant relationships, for example between harmony and form [2], the possibilities of rich internal structure instilling a sense of the ‘musical’ in the absence of cultural conventions.

1.1 Rationale of Current Approach

The idea behind the Adaptive Systems Music (AdSyM) approach is to use biologically inspired self-regulating algorithms to generate music in the belief that their intrinsic properties can create musically meaningful relations. A ‘bottom up’ approach is adopted: the algorithms used can be characterised by the specification of local rules which induce global organisation. These are used to organise similarly low level musical elements of pitch and rhythm. Rather than subscribing to particular musical conventions to impose harmonic, rhythmic and high level form, the organisational mechanisms of certain algorithms are exploited to generate structures in these dimensions.

Rhythmic Cellular Automata Some consider rhythm to be the first and most fundamental element of music. It has been suggested that its repetitive nature has ties to diurnal and/or biological rhythms [10]. If the musical qualities of rhythm are linked to rhythmic patterns in the natural world, it seems possible that these qualities could be captured by a biologically inspired iterative algorithm. Cellular Automata (CA) are dynamic, discrete, deterministic systems, defined by sets of local conditional rules that determine the state of a certain cell according to the states of its neighbouring cells. Operating iteratively, the local rules induce global organisation. CA rules can be classified according to the nature of this organisation into homogeneous, cyclic, chaotic aperiodic or complex local patterned. The self-organisation and pattern propagation properties of certain CAs in particular have been used extensively in theoretical biology [11], computer sciences [12] and the generative arts [1] in investigations of morphology and pattern

formation. Here it is hoped that the pattern propagation properties will transfer successfully to the auditory domain as a basis for rhythm.

Harmonic Homeostasis. In selecting a means of generating harmony, the role of traditional rules of harmony was considered, and a mechanism for creating meaningful relations between pitch values sought. Rather than considering ‘meaningful’ relations in terms of the natural harmonics of the 12 tone scale, a microtonal system is adopted, and harmonic relations determined according to the internal consistencies evoked by a homeostatic process. Homeostasis depends upon feedbacks, negative or positive, between a number of systems. For example during uterine contractions, oxytocin is produced which triggers an increase in frequency and strength of uterine contractions, producing more oxytocin etc. A homeostatic mechanism is adopted to generate harmonies, as the central feedback mechanism means that the value of each part is defined, not absolutely, but relative to all the other constituent parts, just as in a musical chord.

Higher Level Forms from System Dynamics Both these algorithms are essentially dynamic: the CA pattern, and homeostatic relations between parts develop in time (ie iteratively). By implementing the algorithm-audio mappings in real time, the algorithmic *process* is exploited to create rhythmic and harmonic structures. In biological systems, homeostatic processes are typically auto-regulating, for example critical triggering levels are adapted according to the state of the entire system. Biologically, this functions to maintain a range of conditions within which the body can operate most efficiently, even when there are radical changes in the external environment or internal system. More generally, the process achieves a robust stability within a multi-part system. Here, a self-regulating homeostatic network (see section 2.2) was developed with the intention that the system dynamics would produce long-term behaviour sufficiently complex to create musically interesting structures at the level of form. The successful generation of such structures may provide new possibilities for gaining insights into the underlying structural principles of traditional musical forms.

2 System Overview

The system presented here operates in real time, to generate polyphonic audio-output from separate algorithms operating within an iterative loop. Musically, the system produces three separate parts, voiced on different instruments. These can be conceived as a sustained four part harmony, a pitched rhythmic part, and a melodic part. The pitch values of all parts are derived from the outputs of a self-regulating homeostatic network, the timings of the rhythmic part are defined according to the states of a one-dimensional CA, and the note placements of the melodic part are determined using a simple stochastic method, all described below.

At each iteration, four pitch values are generated from part of the homeostat network, and the rhythmic and melodic parts voiced at these pitches according to the CA states and a set of randomly selected numbers respectively. The set of CA states and random values remain constant, and are repeated. Each part therefore repeats the same rhythmic loop with variation in pitch. When the homeostatic network stabilises (ie all outputs converge to a point or limit cycle) it is perturbed, pushing the system onto a new path, and creating new harmonies.

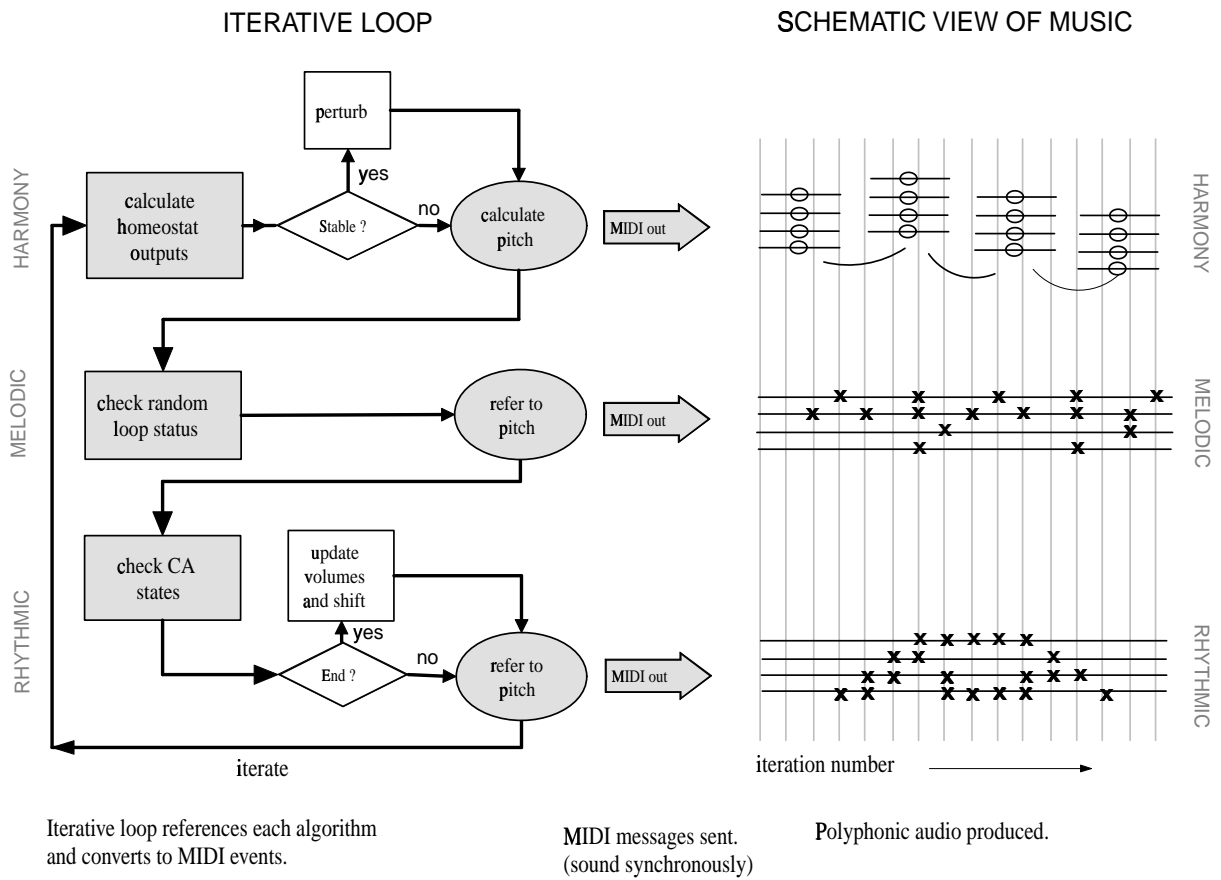


Figure 1: Schematic diagram of audio production from iterative loop.




Figure 2: Graphical representation of 1D CA seeded centrally

2.1 Component Algorithms and Sonification

2.1.1 Generating Rhythm from Cellular Automata.

A rhythmic pattern is generated from the states of a one-dimensional (1D) binary CA. Although CAs can be generated using multidimensional states and spaces, a 1D CA was adopted as this

is conceptually closest to our perception of rhythm. A 1D CA can be conceived as a row of cells that are updated synchronously at each iteration according to a conditional local rule set. A two dimensional representation can be created by simply presenting consecutive states as successive rows (see figure 2). The binary state of a particular cell at time t is determined by the states of the neighbouring cells at time $t - 1$, graphically, those immediately above. The extent of the influential neighbourhood can vary and is here set to 1. The rule set for the CA used here is given below and is implemented on a fixed size grid.

neighbourhood states:	111	110	101	100	011	010	001	000
cell next state:	0	1	1	0	1	1	1	0

Sonification A rhythmic part played by four pitched voices is derived from the states of this 1D binary CA. The rule set above is used to generate the CA pattern in a 19×22 grid (demarcated area in figure 2). The states of the CA are interpreted very simply as a rhythmic score: on = play, off = rest. A fixed size grid is used and repeated (rather than running the algorithm in real time) in order to preserve the triangular shape arising from the central seeding. This creates a changing ratio of 'on:off', producing differences in the length of rests when voiced and therefore textural variation. Four consecutive rows of the CA are voiced simultaneously at the four pitches defined by the homeostat, but transposed exactly one octave down. Each voice reads along successive rows, the last cell in row n being followed by the first in row $n + 1$. This produces a long rhythmic loop.

2.1.2 Simple loop melody

A third algorithm was developed to pick out a simple melody from the harmonies defined by the homeostat. The timings are determined by selecting four random numbers. These determine the period of the notes, so if 5,8,7 and 12 are selected, notes will be played on beats 1,6,8,9,11,13 etc. Each number is associated with one of the currently defined pitches, so notes are voiced at a *time* determined by the random value, at a *pitch* determined by the homeostat. In the example shown in figure 3, the notes G, A, B, and D are played every 5,8,7 and 12 beats respectively. This 'melody line' was included primarily as a means of highlighting the harmonic changes.



Figure 3: Transcription of example output from stochastic loop method

2.1.3 Generating Harmonies from a Homeostatic Process.

Dynamic harmonic relations are generated using an homeostatic process. An abstract formulation of the biological process of homeostasis was developed by Ross Ashby in a discussion of the

origins of adaptation in living systems [13] and embodied in an electro-mechanical machine - ‘the homeostat’. The current algorithm for producing harmony is developed from Ashby’s original formulation.




Figure 4: Diagram of connections for a fully connected 4 unit homeostat ($N = M = 4$).

A simple Homeostat A homeostat is conceived as a system comprising N units each connected to M other units, including a recurrent connection (see fig. 4). Operating in discrete time, the input to unit n at time t is the weighted output of all connected units (including its own recurrent connection) at time t (eq. 1). The output of each unit (at $t + 1$) is proportional to the sum of these weighted inputs (eq 1). These weights are randomly initialised in the range $-1:1$, and are re-randomised whenever the output of the unit goes outside a prespecified critical deviation (set here to $+/- 0.05$).

$$Output_{i(t+1)} = \sum_{j \in c} Output_{j(t)} \times Weight_{ij(t)} \quad (1)$$

where c is the set of connected units, with $i \in c$.

The system approaches one of two qualitative states: *runaway* where the outputs constantly exceed the critical range in wild oscillation, or *stability* where outputs converge to a point or limit cycle. The probability of stability is a function of system size (number of units), connectivity (total number of inputs in the system) and level of damping. In the original (continuous-time) physical machine, damping was controlled by the viscosity of the liquid in which the outputs trailed. This is simulated (in the discrete-time algorithm) using a ‘maximum change’ variable, which defines the maximum difference in output values for any one unit over consecutive iterations. Once stability is achieved, the system is robust to minor perturbation (ie a change in any one output within the critical range). Forced deviations beyond this limit induce random weight changes in connected units and a renewed search for stability. *Suitable* weight values are defined, not absolutely, but relatively: the multiple feedback creating an internally consistent relationship between parts.

Sonification. The relationships between these parts provide the basis for harmony. The output values of a 4 unit homeostat are transformed into pitch values. Rather than restricting pitch to the 12 semi-tones of the scale of equal temperament, the outputs are converted into MIDI pitchbend values according to eq. 2. Pitchbend alters the pitch of a given note by a fraction of a semi-tone, here providing microtonal resolution down to $\frac{1}{32}^{th}$ of a semi-tone. Each unit




Figure 5: Schematic Diagram of Network showing Full interconnections and 10% interconnectivity

is assigned an initial central pitch value from which it deviates according to the value of its current output. Equation 2 with Pb Factor set to one therefore produces a pitchbend of \pm one semi-tone when the output is at \pm one critical deviation. At each iteration, the combined output values produce a chord, dynamic changes in output values creating harmonies. Harmonic structure is therefore defined by the homeostatic process.

$$\text{pitch bend } (i) = (\text{output } (i) \times \frac{\text{pb intervals}}{\text{critical deviation}} \times \text{Pb Factor}) + 64 \quad (2)$$

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.

64 is added as the MIDI pitch bend command is centralised: values 0 – 63 cause a decrease in pitch, values 65 – 127 an increase.

2.2 Self-regulating homeostatic network

Whilst the process of homeostasis provides a new means of generating chords, a simple system stabilises quickly, producing uninteresting harmonies. In an attempt to create a system with more complex dynamics, as well as closer approximating the dynamics of biological homeostatic systems, a modular homeostatic network was developed in which the level of damping (**maxchange**) was self-regulated.

The network can be conceived as 2 four-unit homeostats and a two-unit homeostat (see fig. 5). Each are fully *intra*connected (bold lines) but the *inter*connections (dotted lines) can be varied randomly according to a user specified percentage connectivity. Rather than imposing a fixed damping constant, this value is derived from the outputs of one of the units according to equation 3¹.

$$\text{maxchange}_{(t+1)} = \frac{\text{output unit}_{i(t)} + \text{critical deviation}}{2} \quad (3)$$

where **maxchange** is the maximum amount by which the output of any one unit can vary over consecutive iterations.

¹The 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.

The system exhibits more complex dynamics, whilst preserving the essential characteristic properties of homeostasis (see [14] for a fuller investigation). This network forms the main component of the present AdSyM system which integrates the CA rhythm, stochastic melody and homeostat harmony.

An initial and somewhat incidental exploration of dynamic and tempi changes is also implemented. Each time the CA is repeated, changes in the volumes of the rhythmic and melodic lines, and a transposition of the harmony part are made according to the values of other parts of the homeostatic network (see equations 4-6). The time interval between iterations (in ms) is decremented by 1 each time step, to a preset minimum. This is reset when the homeostat stabilises.

$$\text{Transposition Amount} = \text{output}_{unit_8} \times tc \quad (4)$$

$$\text{Melody Volume} = (\text{output}_{unit_9} + \text{critical deviation}) \times mvc \quad (5)$$

$$\text{Rhythm volume} = (\text{output}_{unit_{10}} + \text{critical deviation}) \times rvc \quad (6)$$

where

$tc = 140$. This value is not critical, but as the outputs will generally be in the range $(-0.05 : 0.05)$, this creates a transposition in the range $(-7 : 7)$ 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.

3 System Output

Audio examples of the ‘music’ produced by the system can be found at:

<http://www.cogs.susx.ac.uk/users/alicee>, and a transcription of an example extract is given in figure 6.

3.1 Public response

The AdSyM was very well received at two public appearances made at different stages of development. An early version, using just the simple four-unit homeostat to produce microtonal chords was used to open and close an electro-acoustic live set for the last night of a festival in Brighton (www.occulture.tv). Several audience members commented on its strong atmospheric effect.

The version as presented here, was played for several hours in a relaxed setting to members of *Blip* - a Brighton based Science/Art discussion forum (www.blip.alturl.com). 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 (www.paul-brown.com), the evenings’ main speaker, showed similar enthusiasm.

3.2 Listener Survey

In a more rigorous setting, 20 volunteers, all unaware of the nature of the compositional process, listened to an example of the audio produced. Each then filled in a survey consisting of scaled questions and open questions (details can be found in [14]). There was significant agreement

The image displays two systems of musical notation. Each system consists of five staves. The first staff of each system is labeled 'I' and contains sustained four-part harmonies. Above these staves are fractional accidentals: the first system has $\frac{13}{16}$, $\frac{13}{16}$, $\frac{3}{4}$, $\frac{1}{4}$, $\frac{5}{16}$, $\frac{7}{16}$, $\frac{9}{16}$, $\frac{1}{8}$, $\frac{5}{8}$, $\frac{3}{16}$, and $\frac{5}{16}$. The second system has $\frac{9}{16}$, $\frac{1}{4}$, $\frac{15}{16}$, $\frac{1}{2}$, $\frac{5}{8}$, $\frac{3}{4}$, $\frac{1}{2}$, and $\frac{3}{8}$. The second staff of each system is labeled 'II' and shows a melodic line with notes and rests. The bottom three staves of each system are labeled 'III' and show a rhythmic part with notes and rests. The notation includes various accidentals (sharps, flats, naturals) and rhythmic values (quavers, crotchets, minims).

Figure 6: Score of output from start of example run, with pitch bend factor set to 2. Top stave (marked I) shows the sustained four part harmonies generated by the homeostat, second stave (II) shows the ‘melodic’ part: the harmonies are played at times determined by the randomly selected numbers 3,14,2,15 with the quaver as unit time. Bottom four staves (III) show the rhythmic part derived from the CA (harmonies are inverted and played an octave down). Fractions at the top of each chord qualify the traditional accidentals shown: for example the first chord should be read (in ascending order) $D \frac{3}{4} \flat$, $E \frac{13}{16} \flat$, $G \frac{13}{16} \sharp$, $C \frac{13}{16} \sharp$. These fractional accidentals apply across all parts.

that the audio produced was “interesting”, “musical” and “would be described as music”.

19/20 listeners 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 standard musical elements: ‘*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.*’.

4 Evaluation

The exploitation of the process of self-regulating algorithms for the generation of novel music seems promising. In contrast to the ‘top-down’ imposition of rules of conventional harmonic systems, the multiple feedback mechanisms in the homeostat algorithm provides a bottom-up mechanism for producing harmonic relations in a micro-tonal system, which are interesting if unconventional. Similarly, the local rules of the 1D CA employed produce a higher level patterned structure that bears a strong sense of rhythm despite lack of consistent metre. Simultaneous presentation of several consecutive time steps voiced at different pitches produces some interesting accents and creates an almost melodic line. Even the selection of random numbers used to generate a ‘melodic’ line is capable of producing some interesting phrases - although the success of this algorithm is much less consistent due to its stochastic basis.

Perhaps the most promising achievement of the current system, is the capability for the creation of higher-level structure derived from the dynamics of the self-regulating network. Examination of the outputs of the entire network over 2000 iterations (approximately 1 hour of listening time), shown in figure 7 suggests that the system passes through several qualitatively different states. These apparent phase changes are evident throughout the whole system and can be summarised:

iteration	0-250	250-700	700-950	950-1250	1250-1650	1650-2000
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. The stable-runaway-stable pattern is characteristic of the behaviour of a simple system at the level of individual response to perturbation. Here however, a similar pattern is also manifest at a higher level. This self-similarity across scale begins to increase the richness of the internal structure of the system.

5 Discussion

“...weird and surprising yet strangely familiar..” listener 17

The main algorithms employed controlled musical events at the lowest level - simply defining positions in pitch-time space. However, the audio produced was conclusively agreed to be ‘musical’. It is not clear from this initial work *which* factors promote the perception of sounds in time as music. In this instance, it could be simply the familiarity of the timbre of the MIDI instruments with which the lines are voiced, or the somewhat arbitrary presence of dynamic and tempi changes. However, listener response suggests that any musical success may be attributable




Figure 7: Graph of outputs of whole network over 2250 iterations

to the internal *structures* of the music which reflect the dynamics of the algorithmic processes.

The basic stable-runaway-stable (S-R-S) pattern characteristic of the simple homeostat when perturbed produces a basic balance of repetition and novelty. (A similar balance is present in the CA rhythm and stochastic melody, as the timings are repeated, but played at changing pitches). This balance arguably encourages cognitive engagement in the listener. Interestingly, this pattern of stability-perturbation-stability is evident in many musical and narrative forms: Larivaille proposed a 5 stage elementary scheme for narrative based on equilibrium-disturbance-reaction-consequence-equilibrium [15]. The same scheme is often applied to film narrative [16], and a very similar pattern is common in many traditional musics such as the exposition-development-recapitulation of Sonata form. In the homeostatic network used here, the S-R-S pattern of the individual unit is also manifest at a higher level in the overall cycle through stability, oscillation and runaway behaviour (see figure 7). This development of a pattern on multiple levels is seen in many complex musical forms. It seems possible that the higher-level dynamics of the network provide an internal structure that promotes ‘musicalness’, perhaps by engaging analytic processes in the listener, despite the absence of traditional musical conventions of form. This possibility must be methodically investigated. There are many other algorithms of this class within the adaptive and dynamical systems fields that have similar potential, opening up a new approach to algorithmic music. The reactive nature of the homeostat also offers potential for interactive systems as an alternative to genetic algorithm approaches to provide a degree of user control.

5.1 Potential for modeling

Consideration of the role of the internal dynamics of the system in producing musical material may also provide a basis for modeling musical processes. It has previously been suggested that CAs may share similar organisational principles to some elements of music [17]. The potential for modeling musical processes using CAs and other dynamic systems could provide an in-road into the development of an abstract theory of music, complementing current musicology of specific traditions.

5.2 Sonification in the arts and sciences

The system presented here is essentially a sonification of abstract simulations of natural adaptive processes. Scientific data has long been used as a basis for composition, and the use of auditory display in scientific tools is not new (eg the Geiger counter). However, increases in computational power provide new possibilities for artists and demanding consideration of new visualisation techniques across scientific disciplines.

Computer music pioneer Charles Dodge worked with natural functions such as fluctuations in the earth's magnetic field [18] and such approaches are becoming more popular with advances in enabling technologies [2]. This interest in objective structure has roots in the serialist techniques of the Second Viennese school (Boulez [19], Webern etc.). Such music still has a somewhat elite, but increasing, audience: this school of composition is perhaps more aptly considered as experimentation or research in contrast to the pure aesthetic aims of mainstream subjectivist composers. Within this framework, the exploration of simulated systems (rather than data from the natural world) offers great potential, both in terms of complexity and controllability.

Increasing computing power also creates ever-increasing amounts of data throughout the scientific community. The need to comprehend this abundance of data, together with advances in media technologies have promoted consideration of the use of auditory techniques in data visualisation. The rapidly developing field of computer modeling in particular - where data is often high-dimensional - presents particular difficulties for current graphical visualisation. Audio's natural integrative properties are increasingly being proven suitable for presenting high-dimensional data without creating information overload for users [20].

Further exploration of the sonification of simulated systems has the potential to contribute to scientific as well as artistic practices. The approach epitomises a current trend in sci-art collaboration, and demands interdisciplinary research effort with musicians and aesthetic experts assisting scientists in the development of intuitive mappings for sonification toward epistemological ends, and computer scientists contributing to artistic practitioners' understanding of the dynamics of aesthetic processes as well as the creation of new art forms.

6 Conclusions

Initial work employing self-regulating algorithms for the generation of new computer music is promising: the dynamics of such systems seem capable of producing structures of sufficient complexity to engage the human mind in the absence of traditional rules of form and harmony. The reactive nature of this class of algorithm also offers potential for developing more interactive systems. In addition to extending artistic practice, the sonification of simulations has potential

scientific application as a data visualisation tool and would encourage a growing interest in science-art collaboration.

Acknowledgments I thank Phil Husbands, Adrian Thompson and Inman Harvey for support and critical comments on the manuscript. The project was made possible through EPSRC MSc and PhD funding.

References

- [1] E.R. Miranda. Evolving cellular automata music: From sound synthesis to composition. downloaded from www.csl.sony.fr, 2002.
- [2] J. Dunn. Life music: The sonification of proteins. *Leonardo*, Dec 1997.
- [3] D. Cope. *Computers and Musical Style*. OUP, 1991.
- [4] J.A. Biles. Genjam: A genetic algorithm for generating jazz solos. *Proceedings of the 1994 International Computer Music Conference*, 1994.
- [5] E.R. Miranda. *Granular Synthesis of Sounds by Means of Cellular Automata*, volume 28. Leonardo, 1995.
- [6] D. Keane. *On the Wires of Our Nerves*, chapter The Quest for "Musically Interesting" Structures, pages 97–109. Associated University Presses, 1989.
- [7] C. Dodge and T.A. Jerse. A cellular universe of debris, droplets, defects and demons. *Scientific American*, May, 1985.
- [8] I. Xenakis. *Formalised Music*. Indiana University Press, 1971.
- [9] D. Worrall. Studies in metamusical methods for sound image and composition. *Organised Sound*, 1(3), 1996.
- [10] H. Simon. *Formal Representation of human judgement*, chapter Pattern in Music, pages 219–250. Wiley, 1968.
- [11] M. Markus and I. Kusch. Cellular automata for modelling the shell pigmentation of molluscs. *Journal of Biological Systems*, 3:999–1011, 1995.
- [12] M. Gardner. Mathematical games the fantastic combinations of john conway's new solitaire game "life". *Scientific American*, 223:120–123, October 1970.
- [13] W.R Ashby. *Design for a Brain*. Chapman and Hall, 2nd edition, 1956.
- [14] A.C. Eldridge. Adaptive systems music: Algorithmic process as a compositional tool. Evolutionary and Adaptive Systems MSc Dissertation 2002, School of Cognitive and Computing Sciences, University of Sussex.
- [15] P. Larivaille. L'analyse (morpho)logique du recit. *Poetique*, 19:368–388, 1974.
- [16] T. Todorov. *Genres in Discourse*. Cambridge University Press, 1990.
- [17] E.R. Miranda. *Composing with Computers*. Focal Press, 2001.

- [18] C. Dodge. *Earth's Magentic Field*. Nonesuch, 1962.
- [19] P. Boulez. *Penser al Musique aujourd'hui*. Gonthier, 1963.
- [20] G. Kramer, B. Walker, T. Bonebright, P. Cook, J. Flowers, N. Miner, and J. Neuhoff. Status of the field and research agenda. In *Sonification Report*. National Science Foundation, 1997.