

Interaction and Self-organisation in a Society of Musical Agents

Peter Beyls

Interdisciplinary Centre for Computer Music Research (ICCMR)
University of Plymouth, UK
peter.beyls@plymouth.ac.uk

Abstract: This paper outlines a distributed architecture defining a virtual world where musical agents interact according to the expression of mutual affinities. Agents continuously exchange information in their respective neighbourhoods while self-organization takes place. The society functions on a scale between total autonomy and a platform that accommodates compelling man-machine interactions, providing an adaptive musical playground. Agents associate spontaneously into temporary clusters, viewed as emergent structures. These clusters are considered the result of perpetual self-production following the theory of autopoiesis. The fluctuating associations are interpreted as complex polyphonic constructs in real-time.

1 Introduction

The current paper describes aspects of the output section of an interactive music system aimed at supporting compelling man-machine interaction in the domain of open, non-ideomatic improvisation. The project explores both symbolic musical pattern processing [3] as well as artificial life oriented, sub-symbolic computing [12]. The system includes a listening component using genetic algorithms to breed the appropriate sensors to capture significant changes in the input MIDI stream [1]. Musical output adopts the principles of self-organisation [13] in a distributed network of agents engaged in social interaction. In most Alife research, an agent is characterized by two important features: first, the level of its functional autonomy and second, how well it can adapt to large swings in context. Autonomy implies the use of sensors and the creation of relationships between perception and action. An autonomous agent can adapt according to its behavioural history and current perception of the environment. Adaptation is a global emergent quality: surprisingly coherent yet unpredictable behaviour follows spontaneously from the expression of a large mass of inter-agent affinities. Emergence is thus a consequence of self-organization, the creation of structural organization without any central design agency. The survival of any biological system seems to rely on the self-organizing properties of its constituting cells, a process known as autopoiesis [8]. Autopoiesis implies perpetual self-production, the continuous creation of new answers while facing an unpredictable environment. This image is central to our work. We think of musical improvisation as the expression of variable *relationships* amongst a collection of agents, including the human performer. A relationship is thought of as a structural coupling between two interacting entities – either synthetic or organic. Complexity in a living system is viewed as the result of a self-organizing process. External pressure may force changes to the organism but the palette of potential changes is made available by the process of self-organization – a concept known as structural determinism. The interplay between man and machine - thought of as interacting complex dynamical systems [15] - thus follows a series of structural couplings. The

human can be seen as an *environment* with its own structural dynamics in which the machine agents interact. However, human and machine agents are viewed as operationally separate. This implies that the environment cannot simply control the interactions between an organism and that environment. In contrast, the structure of the organism determines what kind of effect an external disturbance will cause. The structural coupling between organism and environment is characterized as a form of influence, rather than instruction. In other words, the focus is less on what kind of information is transmitted, but more on what kind of structures are affected inside the receiver. Therefore, we suggest viewing autopoiesis as an alternative to conventional mapping in interactive composing. Man and machine are seen as sources of mutual influence rather than control. The essence of life, and for that matter deeply engaging musical interaction, is imagined as the existence of continuous structural changes in many dimensions inside the interacting partners – with the understanding that changes happen within the scope of structural limits that defines them guaranteeing survival. When man and machine support rewarding interactions happening over longer time spans, we conclude they are compatible, in other words, they are adapted to each other. This is the process of *ontogenesis*; the individual history of structural changes in the organism while its basic structural organization (its identity) and ability to adapt remains in tact.

In this paper we are foremost interested in the quality of the dynamics of the interaction *inside* the society. Therefore we set up experiments where a *simulated external environment* produces fluctuating activations impacting on individual agents. Other versions of the same system derive activations from a human musician, but this is beyond the scope of present paper. The proposed model even develops interesting structures without any external activation at all. In that case, the society performs as an autonomous system that continuously rearranges the positions of its agent-components. Agents form temporary structural associations that are translated to MIDI streams, the music thus reflects the structural changes in time while global structural integrity (survival if the system) is guaranteed.

2 Definition of an agent

An agent is a both a graphic object in 2D space and a MIDI player object (Figure 1). The following instance variables are significant for the present study:

- Physical position in the action pane and angle of movement. These values are equivalent to the initial conditions of a complex dynamical system.
- Energy level. Agents dissipate less energy when their position remains stationary than when they are moving. An agent becomes stationary when the net sum of all affinities impacting on it sum to zero or when its energy level is lower than a given threshold. The energy dissipation factor when active and the recovery factor when asleep are not equal; this contributes to global non-linear behaviour.
- Critical-distance-1: determines the sensitivity of an agent to communicate with any neighbour.
- Critical-distance-2: all agents within a radius of critical-distance-2 of an agent are considered neighbours of that agent (CritGap fader in the interface).
- Activation: the activation is a signed quantity (-100 to 100) intended to function as a source of qualitative information. In particular, it has great impact on the musical interpretation of clusters in which an agent happens to be associated.

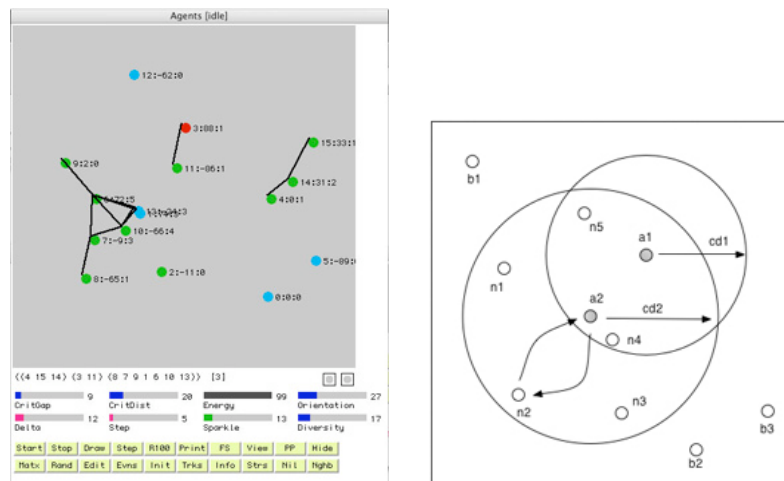
Activation, it is subject to changes forced by pressure from the environment, i.e. a human musician. However, the current paper focuses on autonomous behaviour following experiments with fixed activations or activations that fluctuate randomly. A probabilistic algorithm is used to generate a slowly changing activation level that moves in the range -100 to +100 (please refer to program listing in Section 3).

- Orientation: the orientation bit designates two types of interaction between any two agents. If zero, the agent will address the stress problem; it will attempt to move to an alternative location that will result in less stress impinging on it. If one, it will engage in interaction with all other agents in its immediate proximity.
- Personality dataset: a list of pitch-intervals, durations and velocities from which sub-lists are taken to construct melodies.

3 Implementation

As a simulated physical object, the agent moves in two-dimensional space while dissipating energy proportional to the distance travelled. When energy becomes lower than a given threshold, it enters a sleep state for a number of clock cycles. When waking up, the agent's new full energy level is set according to an external global parameter. Both the energy ceiling and the sleep cycle length introduce non-linearity and unpredictability in the system. Since agents condition each other's movement, they also influence each other's energy consumption. Every agent has two options for local interaction, controlled by the value of its orientation bit. If zero, it will try to lower its perceived stress by changing its physical position. If the orientation bit is one, it will interact only with some of its neighbours. The list of actual neighbours is computed by considering all agents within a critical distance from a referent agent. Both alternatives are detailed next. Agents are viewed as a society in which a social climate exists, a climate that is in constant flux because of a specific affinity that every agent expresses towards every other fellow agent. Affinity is articulated as a scalar value, which forces the agent to be at a certain preferable distance from all other agents. Since these affinities might be conflicting, the result is emergent push-pull behaviour. In addition, the net impact on any given individual agent is made up of contributions of the affinities from all agents that are not asleep. An agent thus moves spontaneously because it aims to minimize its own individual tension towards the rest of the society. The algorithm works as follows; every agent considers a position at a distance (given by the delta parameter, typically a value between 5 and 20), relative to its current spatial position in the 2D field of action. Eight neighbouring positions are considered (equivalent to the Von Neumann neighbourhood in a cellular automaton) and the new tension is computed and catalogued. Finally, the agent will select the position that guarantees the least stress. Since all agents follow the same uniform rule, the network of mutual tensions is pulled towards local minima. In other words, the agents all contribute at minimizing the stress in their society, an example of emergent functionality. Affinities for 16 agents are represented by a matrix of 16 by 16 elements holding values between 0 and 400. This matrix can be inspected and edited in its private interface. High values will produce the effect of agents spreading away from each other. Small values will tend to cluster the agents into dense configurations. By intuition, we may understand that a particular mix of random values will implicitly impose complex, animated behaviour. The society becomes a

complex dynamical system that exhibits all kinds of cyclic and chaotic attractors. By providing certain matrix contents, the society as a whole produces a wide range of interesting oscillations. Experiments 4 and 5 (described in Section 5) reveal different evolutions of structural changes over 40 process cycles.



Figures 1 and 2: (1) The main interface showing the action pane with tree temporary clusters. (2) Example of neighbourhood.

When its orientation bit is set, an agent will consider its proximity and possibly engage in a bonding process (Hofstadter, 1995). All agents within a given critical-distance-1 are addressed – the agent acts as a catalyser, i.e. only its neighbours will execute local rules, not the agent itself. For every agent within that range, its respective neighbours are collected in turn. The *collect-neighbours* function has its own sensitivity parameter: the critical-distance-2. Referring to the situation depicted in Figure 2, as a neighbour of *a1*, *a2* will interact with one of its neighbours, in this case, agent *n2*. Every agent in the subset (*a2 n5 n4*) will exchange information with one of its own neighbours; they swap the current value of their private angles of movement. Strange behaviour emerges because this procedure acts as a specialized delay line. In addition the difference between the two critical-distance values introduces considerable complexity. One may tune both sensitivities independently and so influence - though not exercise explicit control over - global behaviour. In addition, by changing the relative proportions of the agents' orientations, structural development in the society will either (1) follow from the expression of affinities or (2) any two interacting agents will interfere with a subset of their immediate neighbours and exchange information on how they move in space. The *sparkle* parameter acts as a probabilistic source introducing noise in the system. It is intended to keep the system away from point attractor behaviour and follows the knowledge that organic networked systems, like human brains, produce random firing patterns even without any apparent external stimulation (Harth, 1995). When sparkle fires, the orientation bit of the current agent is temporarily inverted. Two performance modes exist: the agents' animation process and the MIDI player process can either be synchronised or run independently. By definition, the animation process (running the

simulation and real-time visualisation) updates at a chosen rate, typically 500 msec to 5 seconds. The player process computes a new melody when the previous one just finished playing. When running independent, many simulation cycles may pass while a melody is played, thus when a next one is computed, its contents will echo how much the world has changed since it was last sampled. Otherwise, when in sync, the simulation waits until the player has finished playing the current melody. In that case, changes in the society will immediately be reflected in the music generated.

4 Musical articulation

Musical output should reflect the history of the relationships between individual agents i.e. the structural changes of the system in time. The resulting melody is computed as follows: every agent first creates a list of its neighbours, i.e. all agents within a physical distance set by its *critical-distance-1* instance variable. In order to evaluate many sensitivity schemes, the *critical-distance-1* value is tuneable: a bias value plus a diversity value (deviation from bias, in percent). Next, the *neighbour-clusters* function isolates individual groups of agents according to their neighbours. The result is a variable series of temporary structures known as ‘clusters’. Clusters are visualized by drawing segments between agents falling in each others zone of influence; i.e. their distance below a given critical distance threshold (Figure 1). Such clusters are viewed as emergent structures. They reflect how individual agents create temporary alliances, very similar to the ‘flickering clusters’ described in [7]. Next, the energy of all clusters is computed by averaging the energy of all its constituent agents. The cluster with the maximum energy is further selected. Note that short clusters (few agents) may provide more energy than long clusters -- size and energy are thus interacting in subtle ways. A reference melody is created next, it will serve as a musical backbone. The ID of the first agent in the strongest cluster (the reference agent) provides the MIDI channel (1 ~ 8) for all events to be added to the reference melody. In addition, the private, stylistic data of that agent is borrowed to compute events. The data includes a list of pitch-intervals, durations and velocities. The energy of the reference agent will influence the number of events generated as well as the density (articulation by way of rests) of the events. Next, we consider the other agents in the cluster. We shall examine their *activation* rather than energy. Only when the status of a remaining agent is *active* it will throw in additional events to the backbone. Two options exist: expansion or contraction. Expansion signifies that supplementary events are generated from a single source event. Contraction, in contrast, creates a single new events from the data supplied by a group of existing backbone events. The results coalesce as a form of emergent musical orchestration.

Event generation by contraction: Contraction is computed using a grouping-algorithm; it typically produces additional events characterized by features of specific groups of backbone events. Its parameters include: *groupings-list*, *source-channel*, *intervals*, *groups-flags-list*, *transposition* and *destination-channel*. All argument lists are of different and arbitrary size; items are addressed cyclically using an incremental, global pointer, taken modulo the length of the list in question. (This implies that arguments interact in irregular ways). A minority of argument lists are created on the spot from a pool of candidate values; these values as a whole constitute a small database of stylistic information. At this point we hit the atomic level of our system,

one of the very few instances where factual knowledge, expressed as numerical data - is specified by an external human designer. Otherwise, stylistic data is borrowed from the private data of every next agent in the cluster: pitch-intervals, durations, velocities, activation and energy level. The algorithm first scans the backbone and creates groupings of 2, 3 or 4 source events - only considering events with MIDI channel equalling the source-channel - summing the duration of the groups for use by the shadow events to be generated. In addition, pitches are at an offset (from the first pitch in the group) guided by the intervals parameter while transposition argument provides a global pitch offset. The groups-flags-list is a Boolean list conditioning the new event to be added to the backbone or not, using the specified destination-channel. The length and the number of true items in the groups-flags-list are proportional to the absolute value of activation of the agent.




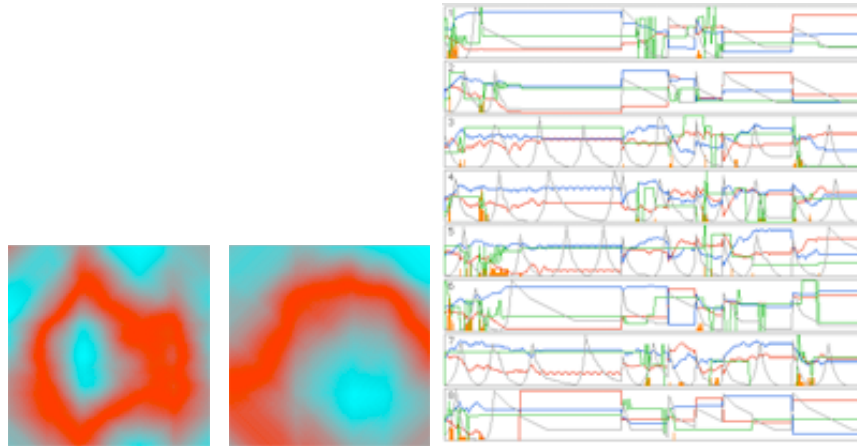
Figure 3: Example of emergent cluster orchestration.

Event generation by expansion: The expansion algorithm performs as functionally opposite to the contraction algorithm, it spawns many events from one source event. It applies the eight following parameters: source-channel, intervals, transposition, minimum-duration, duration-dividers, group-flags-list, new-channel and delay-list. All events of duration equal or higher than the minimum-duration are collected – conditioned by the group-flags-list. The particular mix of Booleans may thus partition the source material in asymmetrical ways. The duration-dividers (typically 2, 3 or 4) split up the collected durations into that many new values. Now the start-time of the nascent event is postponed by a delay, the product of the new duration and the value taken from the delay-list argument. This method yields automatic synchronization: start-times and/or end-times of source and destination events will align in most cases. The application of intervals and transposition arguments is similar as in the contraction algorithm above. Finally, the sign of the activation of the agent acts as a switch for either the contraction (positive) or expansion (negative) procedure to be selected. Figure 3 shows a prototypical example of emergent cluster orchestration. Colours refer to eight individual MIDI channels. The accumulative effect of the clusters is clearly observed.

5 Experiments

Figures 4 and 5 exhibit two different stress views, both as seen from the perspective of a single reference agent. The affinities matrix is filled with random values between 0 and 400; the upper limit is derived from the size of the action pane. The two figures show the effect of two different spatial configurations of the agents while the affinity matrix remains unchanged. The total stress of the reference agent is computed by summing the differences of ideal position (guaranteeing minimum stress) as compared to current position, for all agents in the society. This provides the

magnitude of stress of a single agent towards all other agents given their current physical positions.



Figures 4, 5 and 6: (4, 5) Two instances of temporary tension landscapes. (6) 400 generations of 16 interacting agents.

In order to visualize this global view we create a stress-view array of 80 by 80 elements, it will map to the 400 by 400 pixels action pane. The reference agent visits regular positions in the pane in coordinate steps of 5 pixels, thus implying a lower resolution for this simulation. A global picture of the tension experienced by a single agent thus accumulates in the stress-view array. It is next normalized to values between 0 and 100. The integers in the resulting array serve as pointers in a colour lookup table that interpolates between light blue to deep red. The resulting image shows a two-dimensional tension landscape. It reveals the non-linearity captured in the affinity matrix, more precisely, it shows the relative target tension for all possible positions (80 by 80, spread out over 400 by 400 pixels, implying a resolution of 5 pixels) of a single reference agent. Three experiments with random activations (between -100 and +100) defined at instantiation time, activations remain unchanged during the course of all 3 experiments. Figure 6 displays 400 generations of 16 interacting agents, though only the first 8 agents are shown. Five values are traced: x and y position in 2D space (red, blue), angle of movement (green), energy (grey) and the bonding history (orange). Bonding amplitude refers to the number of agents interacting at every moment in time. All system parameters remain unchanged throughout the simulation. Both critical distances equal 10. The agent's orientations are also fixed: (1 1 0 0 0 1 0 1 0 0 1 0 0 1 0 0). Thus only 6 out of 16 agent will exchange information with their respective immediate neighbourhoods. The positions of the agents are occasionally disturbed by external action but the affinity matrix remains unaffected. One observes the effect of the couplings expressed in the matrix. The system accommodates the disturbance and resettles into orbits of variable periodicity. When the system is pushed out of relative stability, it evolves from erratic to more regular oscillations, occasionally hitting a point attractor. Local stable highly

periodic oscillations occur over extended periods of time, witness the push-pull activity between agents 3, 4, 5 and 7.




Figure 7: History track with all orientations set to zero.





Figure 8: Outcome from the third experiment.

Figure 7 shows a similar history track with, after a few generations of local interaction, all orientations set to zero. Thus the only objective of every single agent is to minimize its locally sensed stress by gradually adjusting its position. The angles stay fixed. One clearly observes the gradual introduction of more consistent paths of movement in space. This shows the capacity of the affinities matrix to create spontaneous order out of initial randomness. The x and y movements become more or less synchronized as well. The energy dissipation process creates less irregular peaking at a slower tempo; the global breathing of the system slows down, apparently settling into a point attractor in the long run. Figure 8 documents the third experiment. The density of positive values in the affinity matrix is only 15 percent. Orientations are (0 1 1 1 0 1 0 1 0 1 0 1 1). We observe four behavioural phases: initial chaotic interactions, then only agent nr. 4 engages in a steady bonding process while most of the other agents remain stationary, then a cycle attractor is hit with only agents no. 5 and no. 11 interacting. Finally, when critical-distance-1 is slightly increased, the oscillations become less periodic and the bonding process starts peaking again albeit in a limited number of agents. The variations in the energy profiles reveal evidence that energy dissipation and bonding density interact.

The next two experiments document the systems behaviour under specific parametric conditions and include a visualisation of the internal structural changes and the resulting score. Figure 9 shows a free running simulation over 40 generations, with a snapshot of every temporal structure. The affinity matrix is filled with random numbers in the range 50 to 200 (the upper limit is half the size of the action pane). All agents receive initial random energy in the range 50 to 100. All orientation bits are zero, so only the critical-distance parameter is significant and no bonding can occur. Critical-distance is 27 for every agent. Delta and step parameters equal 5. The frames of the behavioural history are read top left to bottom right. The simulation starts with random positions of all agents. Three clusters of 2 agents each are visible in the first

generation. All agents move one step according to the social forces expressed in the affinities matrix. The average pull is strong and all agents gradually collide in a single cluster. From here on, the local push/pull activity on the individual agents disintegrates the cluster and makes it shift in space. Figure 10 shows the resulting score orchestrated – over eight MIDI channels -- according to the contraction/expansion algorithm described in paragraph 4. The temporal articulation of the density of the events and how they relate vertically reflect the structural couplings between the agents over time.



Figures 9 and 10: (9) Experiment 4, where structural changes with random contents of affinities matrix. (10) Experiment 5, where structural changes with uniform contents of affinities matrix.




Figure 11: Interpretation of behaviour shown in Figure 9, respective to bars 1 to 9 corresponding to generations 1 to 10.

Figure 10 documents an experiment where the affinity matrix is filled uniformly with the value 150, thus every agent expresses the same social preferential distance towards every other agent. The orientations vector is $(0\ 0\ 1\ 0\ 1\ 0\ 1\ 1\ 0\ 0\ 1\ 1\ 0\ 0\ 0\ 0)$. Thus 6 out of 16 agents may engage in a bonding process with a neighbour agent within the critical-distance-2 (value 22 in this simulation) from that agent. The critical-distance-1 is 39 for every agent. Frame 1 shows the initial random positions of the agents. All agents express strong attraction at the start of the simulation resulting in the creation of a single complex cluster. This cluster gradually disintegrates and transforms itself into a diamond-like shape clearly visible as from generation cycle 17

and remains relatively stable during the next few generations. The critical-distance-1 is increased to 67 at generation 33 and a single cluster emerges.

Since the affinities matrix remains uniform, the cluster does not change dramatically over the next generations. The musical interpretation in Figure 11 reveals more articulated coherence between neighbouring MIDI channels and a much more regular overall picture than the score in Figure 12. However, it is difficult to detect where respectively the contraction and expansion algorithms were applied because their musical effect is very similar.

6 Discussion

We may think of the agents as virtual musicians with private MIDI channels, floating in a 2D space. They create highly complex temporal superstructures of great plasticity. This interesting non-linear behaviour is a consequence of the inter-agent affinities, the flux of external activation and the dissipation of energy. Thus a simple biology inspired control structure exists supporting complex musical interaction. Also, we may exert influence over the quality of the internal interactions by using only a very limited number of control parameters: in particular the critical distances and the orientations vector. In terms of musical control theory we speak of a strong control structure implying coherent, minimal control over maximal complexity.

The advantages of the current model can thus be summarized as follows: (1) it supports a strong musical control structure yet expressing mutual influence rather than explicit control, (2) it provides musical behaviour that seamlessly integrates subtle external activation in an otherwise completely autonomous system, and finally (3) it supports self-organising behaviour which we hypothesise to be instrumental to the synthesis of both interesting internal inter-agent interaction and rewarding man-machine interaction.



Figure 12: Interpretation of behaviour shown in Figure 10, respective to bars 1 to 11 corresponding to generations 1 to 15.

7 Related work in context

Many theories addressing complexity analysis make use of distributed models. Consider the work of Minsky [11] aiming to explain cognitive processes as globally emergent properties as a consequence of local interaction amongst simple components. The fluid dynamics metaphor developed in [7] also aims to explain cognitive processes. In particular, the ‘flickering cluster’ image is highly suggestive. The stable-seeming fluid-like properties of thought emerge as a statistical consequence of a myriad tiny, invisible, independent, subcognitive acts taking place in parallel. However, our implementation views clusters as unpredictable and variable partial configurations within a fixed pool of agents objects.

The animated microworld of Ventrella [16] explores attraction and repulsion in a simulated society. The party planner model developed earlier by Goldstein [5] also investigates similar dynamic principles to model social configurations by way of the specification of idealized physical distances between people. Both approaches are significant sources of inspiration.

In addition, the discipline of artificial life has spawned many incarnations of distributed computational models. According to the field of application we speak, for instance, of particles, molecules, agents or artificial creatures. Particle systems have a long history in the computer simulation of complex natural phenomena. They are instructive, early examples of studying complexity by considering relationships amongst small buildings blocks – in contrast to using differential equations. More recent work by Reynolds [14] describes the flocking of birds (boids) as an emergent process. The Swarm simulation environment developed at the Santa Fe Institute [10] is also a noteworthy example. The swarm and boids ideas were adapted for musical purposes by Blackwell and Young [2]. Swarm Music is an interactive music improviser. It maps the positions of particles/boids to positions in MIDI space. An external human improviser may act as a temporary target for the swarm. Style-scripts provide additional parametric control over the nature of the interaction thus further conditioning the musical output. Dahlstedt and McBurney [4] is a recent example that addresses musical autonomy in a society of music composing agents. The Eden interactive installation project by McCormack [9] is a sophisticated example of distributed audiovisual intelligence. This system features agents that listen, act and otherwise evolve in a virtual world connected to the physical world via infrared sensors. Finally a paper by this author [1] details a molecular collision model of musical interaction.

8. Conclusion

This paper introduced a self-organizing real-time musical system that views improvisation as a process of perpetual renewal. It takes significant inspiration from the theory of autopoiesis and the fluid dynamics theory developed by Hofstadter’s team. From the musical point of view, the evolving associations between individual and group behaviour are particularly interesting. Even without any external activation, the society can sustain complex interactions for hundreds of generations.

One can also view the system as a strong musical control structure because of the extreme economy in parametric specification. The most significant parameters are the critical-distances, energy, activation and the affinity matrix. Because all parameters are continuous rather than discrete, the implied behavioural space is virtually infinite.

Yet one can easily tune the system to function within specific behavioural boundaries. Subsequently, powerful emergent musical personalities do surface from the interactions either involving human interactors or not.

References

1. Beys, P. 2005 *A molecular collision model of musical interaction*, Proceedings of the 8th Generative Art Conference, Milan, Italy
2. Blackwell, T and Young, M, 2004 *Self-organised music*, Organised Sound 9(2), Cambridge University Press, UK
3. Cope, D. 1991 *Computers and musical style*, Vol. 6, The Computer and digital audio series, A-R Editions, Inc. Madison, WI
4. Dahlstedt & McBurney, 2006, *Musical Agents: Toward Computer-Aided Music Composition Using Autonomous Software Agents* Leonardo, Vol. 39, No. 5
5. Dewdney, A. K. 1987 *Computer Recreations - Diverse personalities search for social equilibrium at a computer party*. Scientific American, September 1987
6. Harth, E. 1995 *The Creative Loop. How the Brain makes a Mind*. Penguin Books, UK
7. Hofstadter, D. 1995, *Fluid Concepts and creative analogies, Computer Models of the Fundamental Mechanisms of Thought*, Basic Books, Harper Collins
8. Maturana and Varela, F. 1984 *The tree of knowledge, the biological roots of human understanding*. Shambala, Boston
9. McCormack, J. 2003 *Evolving Sonic Ecosystems*, in A Adamatzky (ed), The International Journal of Systems & Cybernetics - Kybernetes, Vol 32, Issue 1/2, Emerald, Northampton, UK
10. Minar, Burkhart, Langton & Askenazi, 1996 *The swarm simulation system: A toolkit for building multi-agent simulations*. Technical report, The Santa Fe Institute
11. Minsky, M. 1985 *The Society of Mind*, Simon and Schuster, New York
12. Miranda, E. R. 2002 *Sounds of Artificial Life*. Proceedings of the ACM SIGCHI International Conference - Creativity and Cognition 4
13. Prigogine, I. 1980 *From Being to Becoming, Time and Complexity in The Physical Sciences*, WH Freeman & Co. NY
14. Reynolds, G. 1987 *Flocks, herds and schools: a distributed behavioural model*, Computer Graphics, vol. 21, nr. 4, Siggraph Proceedings, ACM Publishing
15. Thomson and Stuart, 1986 *Nonlinear dynamics and chaos*, John Wiley & Sons, Ltd. UK
16. Ventrella, J. 1993 *Attractions and Repulsions, A Gravity-based Animated Microworld for Exploration of Dynamical Systems*, The Media Lab, MIT