Cyborg dancing: generative systems for man-machine musical improvisation.

Alice Eldridge.

Creative systems lab., Evolutionary and Adaptive Systems group, University of Sussex, UK.

alicee@sussex.ac.uk

Abstract

One of the major motivating forces in generative art is the desire to explore uncharted spaces, to create artefacts that escape the designer's control: to attain emergence. This paper focuses on the design of digital systems that would be suitable partners for man-machine collaborative exploration of these spaces. Limitations of existing approaches that impede the artist/user's creative explorations are reviewed. The central problem is framed as a constraint on the variety of outcomes that are possible. Taking inspiration from existing musical practices, an alternative approach is proposed and illustrated with a working example of a generative man-machine system for improvised musical performance. The principle difference between this and many other digital generative artistic tools is that whereas the material from which the final artefact is made is usually defined, here it is provided in real time by the performer. This appears to increase the creative freedom of the user, whilst preserving the independence of the digital generative process and offers a practical alternative to the slippery concept of creative emergence in increasing the variability of possible artistic artefacts.

Key words: Generative art, Alife art, man-machine improvisation, music performance tools, creative emergence.

1. Introduction

From process experimentalists such as Cage, through conceptual artists such as Hanns Hacke to practitioners today such as Jon McCormack and Ken Rinaldo, the use of generative processes reflects a desire to at once relinquish control, and to explore the open-ended dynamics of the universe. These endeavours are inevitably and fruitfully informed through a close and mutual exchange with Artificial life (ALife) and Adaptive systems research, both in terms of conceptual agendas and practical mechanisms. It has been suggested that the application of Alife techniques holds promise for digital generative systems to overcome "the hackneyed replicable paths of interactivity that have thus far been presented by the arts community" [1], and suggests ways of creating a "cybernetic ballet of experience with machine and human involved in a grand dance of each sensing and responding to the other" [ibid]. Such possibilities spark the imagination and fuel the desire to explore and express the "computational sublime" [2], and create "sculptural and virtual algorithmic manifestations which will surpass our wildest dreams" [1].

This paper presents a consideration of the characteristics that we might desire in a generative system suitable for man-machine collaborative exploration of these realms and illustrates these choices with a working example of a generative system for man-machine musical improvisation. Whilst there exist many examples of truly engaging interactive exhibits [see 3 for a review], examples of generative tools, particularly for live performance, are less evident[1]. Some of the limitations of existing design approaches to generative tools are considered in section 1. These can be framed as constraints on the variety of possible artistic outcomes (due to both design decisions, and the inherent restraints of digital systems). Taking cues from autonomous robotics, it has been suggested that one way of overcoming these constraints is the creation of systems which exhibit 'creative emergence' [6,7]. Apart from the fact that this concept is difficult to define and has yet to be demonstrated in digital systems, it is suggested that the associated characteristics of such systems may not be amenable to the creation of artefacts that tickle the human aesthetic. In section 2, McCormack's [2] scheme of the generative process is presented and applied to standard musical practices. Taking inspiration from these, a practical alternative to creative emergence in autonomous systems is proposed which preserves the generative aesthetic and liberates the user's creative freedom, broadening the arena for true cyborg dancing. Section 3 outlines a working implementation of this framework which is currently being used for generative man-machine improvised musical performance.

1.1. The creative limitations of existing generative tools

In [8] Dorin presents an pictorial version of Borges' [9] *Library of Babel* and puts it to the reader that we might question Picasso's artistic genius if he had merely picked the *weeping woman*, along with all his other works, from its vaults. This is presented as an analogy for aesthetic selection which he suggests is, for the non-programmer user, a creatively empty tool: "the 'art' of creating an image using aesthetic selection is

[1] A notable exception is Al Biles' GenJam [4] which has perhaps made more public performances than any other generative system. It is pertinent however, that over the years, the principle system has moved away from evolutionary computation to a knowledge based system containing a database of 1001 best jazz licks [5].

indeed mindless" [p.6 ibid]. Whilst there are a myriad of other generative processes and methods for user interaction, the creative potential of most generative digital systems to date are severely constrained by the designer's programming decisions.

1.1.1. The constraints of digital tools

Design constraints are a problem for software tools in general. Almost any tool or medium leaves its characteristic mark on the artwork with which it is created. And often it is these very characteristics that inspire their use. Although these characteristics could be seen as constraints, they can be distinguished from the *creative potential* of a tool or instrument, which can be considered in terms of both the possibilities it affords, and its usability. As Golan Levin puts it: "a feature of a successful instrument is that its results are inexhaustible and extremely variable" [10, p.54]. Variable results are of course little use, if they are difficult or impossible to control. Another of Levin's desiderata then is that an instrument be "instantly knowable, and indefinitely masterable". [10, p.56] i.e, that the rules of operation should be simple, but the possibilities endless. This is achieved in physical tools like the humble pencil, which has the extraordinary property "that an individual may eventually, through its use, reveal a unique and personal voice in that medium". This is rare in any digital tool. As Dorin notes, in systems such as Latham's *Mutator* [10]: "none of the pixels voice the thoughts of the wanderer" [8, p.10]. Whilst the interface may be instantly knowable it offers no scope for excellence : "there is no means for distinguishing a master from a relatively inexperienced user" [8, p.6].

Dorin goes on to suggest that the designer's control may be relinquished by using aesthetic selection to steer the non-linear interactions of self-organising primitives in order to generate complex higher level emergent phenomena. The programmer would still specify the basic elements and how they interact, but the user could then enter an open-ended conceptual space, sculpting the system into a unique complex emergent structure un-envisaged by the author. This seems to open the space of artistic possibilities, as offered by other tools i.e. to readdress the balance between the artistic skill of the tools' creator (e.g. Stradivarius) and it's user (e.g. Menhuin). Within a generative art framework the thought of such control whisks us away for a brief cyborg pas de deux around the grounds of the computational sublime. But the problem is, as has been noted elsewhere [6], that whilst certain types of emergent behaviour can be demonstrated in silico, there exists no un-contended digital system that exhibits truly open-ended dynamics [12]. The emergence of multi-level phenomenon is a deep open problem in biology [13], leaving intuition as the principle guide in the initial selection of suitable primitives. Finally as Bird and Webster [6] suggest, the mapping of these (yet-to-be-digitally-attained) dynamics into a perceptual medium for artistic ends is non-trivial.

1.1.2 The problem with creative emergence

An alternative proposal for escaping the designer's control and broadening the scope of possible outcomes, is the creation and artistic application of 'creatively emergent' systems [6,7]. The concept is linked to Cariani's taxonomy of adaptive robotic systems and refers to a system's capability for constructing new primitives [14]. In robotics terms, these would be the sensors, effectors and control mechanisms that map sensory input to action. Bird et al describe two physical systems that are capable of such feats: Gordon Pask's Electrochemical Ear [15] and Paul Layzell's evolveable hardware [16]. The creative emergence exhibited by these systems enables them to escape the designer's control, and it has been suggested that if employed in generative art there would be "no clear sense of a 'creator' and an 'author', as the artist and the device both play participatory roles in the creative process" p.44 [6].

Such systems open the space of possibilities. Pask's Ear in particular, also exemplifies an interesting collaborative process. The structural adaptivity required for creative emergence means that the system cannot be 'mastered' in the normal sense of the word. Pask suggests that learning to 'steer' the evolution of the system is a process more akin to animal training [15], where the user must become familiar with the behavioural characteristics and recognise emerging trends. Such approaches to control offer interesting possibilities for man-machine collaboration, but one of the problems is that the incumbent epistemic autonomy in creatively emergent systems implies an 'aesthetic autonomy', i.e. it creates its own aesthetic norms. If we are concerned with creating artefacts for human consumption this may not be an attractive property[2].

Our digital partner for generative man-machine creative endeavours must at once preserve its operational autonomy (it must also be 'interesting'!), escaping the constraints of its designer, and facilitate creatively meaningful interaction (you must be able to 'use' it) toward the production of aesthetically valuable artefacts. Three desirable characteristics then are:

- **independence of generative process** (which is a defining feature of generative art),
- **inexhaustible and extremely variable outcomes** (which gives the user a space for creative exploration)
- **intuitive yet masterable means of interaction** (which is desirable for any artistic tool).

In the next section ways of achieving these characteristics are considered by examining the nature of the interactions and influences in generative arts practice and comparing them with other arts practices.

2. Interactions and influences in arts practice

2.1 Interactions and influences in the generative process

The generative process can be conceptualised using the biological notions of genotype and phenotype as adopted in ALife research [17] (Fig.1). The designer constructs a generative process (the genotype), and typically stands back to observe the phenotype unfolding in the hands of an automated procedure (the enaction of the specification). The genotype invariably acts to structure a pre-specified medium, whether it be pixels [11], MIDI notes [18], old washing machine parts [19], mould on photographic film [20], or the behavioural characteristics of robots [21], creating the phenotypic realisation or artefact with which the

2 In this paper we are concerned with the creation of artefacts which speak to an existing human aesthetic. The types of generative process we are interested in are those whose outcomes occupy the space of 'art-as-it-could-be-which-might-tickle-the-human-fancy'. There is much to be said for considering the remaining vast swathes of art-as-it-could-be-which-might-tickle-the-simulated-agent-fancy (as in artificial creativity research, or perhaps Hard Alife Art) but this is not the focus of this paper.

audience engage. If the genotype specification includes mechanisms which are responsive to environmental feedback, the audience can also 'interact' with the phenotype and potentially influence future outcomes of the system. However, in most systems the material from which the phenotype is formed, as well as the process which structures it are specified by the designer. This restricts the possibilities, limiting the user's creative freedom.

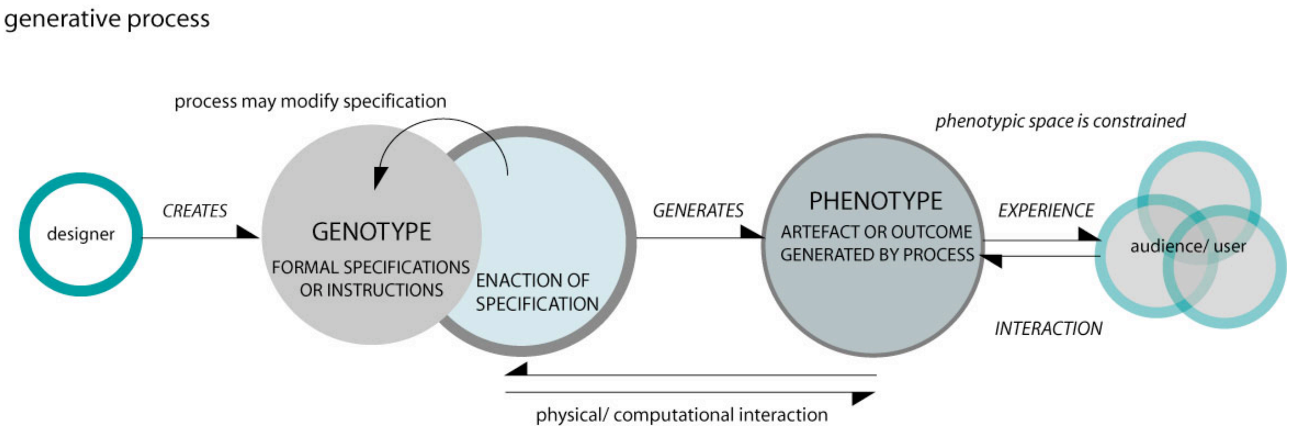

Fig. 1 Overview of the interactions and influences in the generative process. (with kind permission from McCormack [3]). The user's influence on the final outcome is constrained by the designer's decisions over the genotype, enaction process and material from which the phenotype is formed.

Theoretically, creatively emergent systems escape this constraint as environmental feedback can induce structural change and the creation of new primitives. Experience-based modification of genotype and enaction mechanism is the key to the creation of inexhaustible and variable phenotypic possibilities. This is evident in many real-world systems, but remains an open problem in silico [12]. But do we necessarily need to take on this hard and slippery research agenda in order to develop generative arts practice [22]? If we are interested in practical creative possibilities rather than epistemological concerns of ALife, perhaps we can find alternative non-digital solutions to this problem in existing arts practice which can be integrated with generative practices in order to inspire the design of creative cyborg instruments.

The successful modification of the genotype in the biological organism requires a feedback mechanism from the environment to the organism, and some pressure for the modification to be adaptive e.g. evolution and natural selection. A similar process occurs in all sorts of social, cultural and economic systems. As well as in many creative practices. As we have seen above, this process is restricted in standard generative arts practice due to a fixed genotype and enaction mechanism. In the next section McCormack's scheme for generative art is applied to standard musical practice to illustrate the existence of 'open-ended' creativity.

2.2 Interactions and influences in a scored musical performance.

Adapting McCormack's scheme to scored musical performance, we can draw an analogy between the designer's ultimate constraint over outcome in digital generative art and that imposed by the composer in a scored musical performance. Each note value and placement is predefined. Even though the enaction mechanism is now a truly complex adaptive system (a human musician), which is sensitive to environmental feedback (relevant environmental cues being the global sonic outcome and audience's reaction), only intricacies of expression (minutiae of tuning, phrasing, volume etc) can be modified. The phenotypic *structure* is tightly constrained.

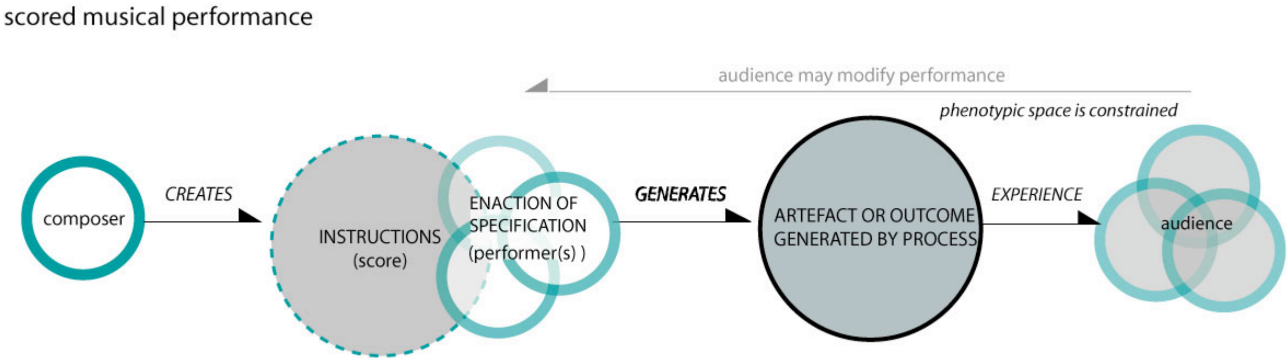

Fig. 2 Influences and interactions in a scored performance. The performers' influence on the final outcome is limited by the composer's specification.

2.3 Interactions and influences in an improvised musical performance.

Compare this to an improvised situation. If we consider the same interactions and influences that take place in an improvised musical performance, we might get something like Fig. 3. Here the outcome is a direct realisation of the performers' musical ideas. The genotype is the collective musical intuitions of members of the group and the enaction mechanism is their control of their instruments. Both their ideas and their actions are inspired by immediate and dynamic feedback from both the instantaneous group sound and audience reaction. Modification is the very basis of the practice, and its musical adaptivity is (generally) assured by the musical sensibilities and skills of the players (they won't get booked again if its not!).

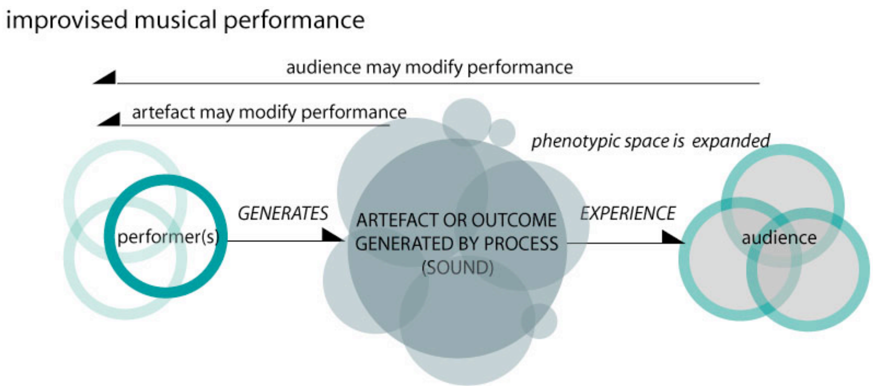

Fig. 3 Interactions and influences in an improvised performance. The performer determines the final outcome and adapts according to its current state.

In this improvisation model, there are no compositional constraints: the final outcome is an artefact of the dynamic process of its production. This in itself can be seen as a generative process. But part of the draw of generative art is the relinquishing of control, the exploration of the 'computational sublime'. How can this model be integrated within the generative model such that the generative system maintains independence from, yet affords interaction with the user to achieve a creatively powerful generative cyborg system? One possible conceptualisation of this integration is given in Fig. 4.

2.4 A model for generative man-machine improvisation

A genotypic specification and enaction mechanism are still specified (and can be digitally implemented) by the designer. The crucial difference is that this enaction mechanism demands human collaboration. The resultant phenotype then is the co-creation of a hybrid man-machine mechanism. The space of phenotypic possibilities is less tightly constrained by the designer, as we now have a truly autonomous open-ended system (the human) partaking in its creation. The digital generative process does not lose its independence in that there still exists a pre-specified genotype, the outcomes of which cannot be entirely predicted by the performer. The automated generative process is still privy to feedback as in the original model. The human presence in the enaction process also creates the feedback seen in the improvised model: it can now be adaptively modified to creative ends. The performer has freedom of creative expression: the constraints of the programmer are relaxed.

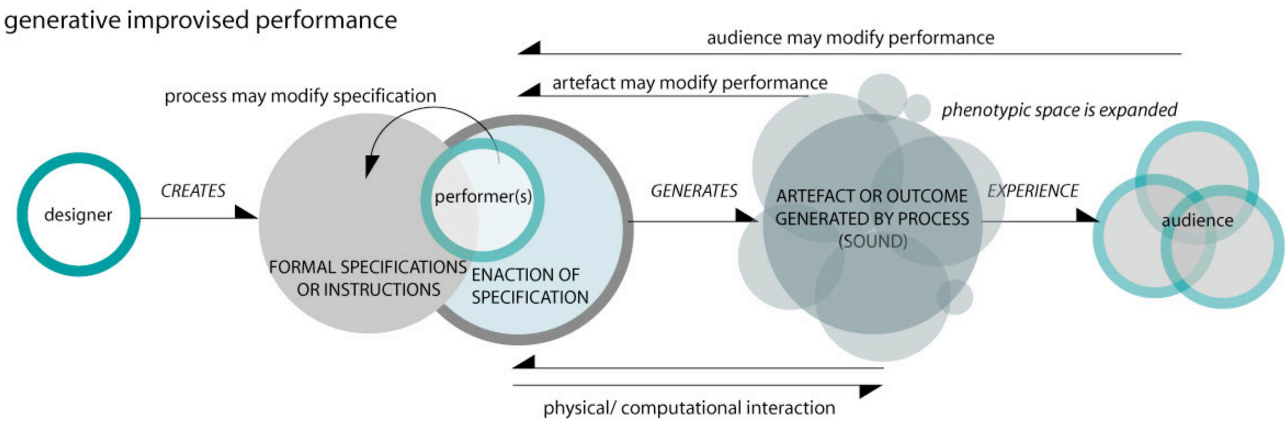

Fig. 4 Interactions and influences in a generative man-machine improvised performance. The enaction process is a collaborative effort between the designer-specified enaction process and the performer. The performer now has greater influence on the final outcome.

This scheme has been realised in a digital system which is currently used for generative improvised musical performance. As noted above, the genotype typically unfolds automatically to structure a pre-defined medium. This pre-definition of the medium is one of the reasons that it would be tricky to make a *biomporph* [23] with Latham's *Mutator* [11]: the system designer determines both the dynamical/ structural properties of the generative system, and the material which is structured. In the current approach the digital generative algorithm is specified by the designer. But the substance from which the phenotype is formed is determined in real-time by the user: in this case by providing sonic material sampled live during the performance. The enaction mechanism itself is a cyborg effort.

3. Fond Punctions: a realisation of the generative man-machine improvisation model

Adopting this approach, a system has been built which is currently used for live improvised performance. The digital genotype of the generative system is based on two conceptually distinct but algorithmically coupled systems: a homeostatic network and a physics simulator. These act to re-compose musical fragments of the performer's improvisation by controlling a granular synthesis engine which processes live samples. The homeostat operates at a rhythmic/phrasal level, the physics simulator influences longer term structure at the level of musical form. Details of the interactions and influences between components of the system are given in Fig. 5. For full details of the system see [24].

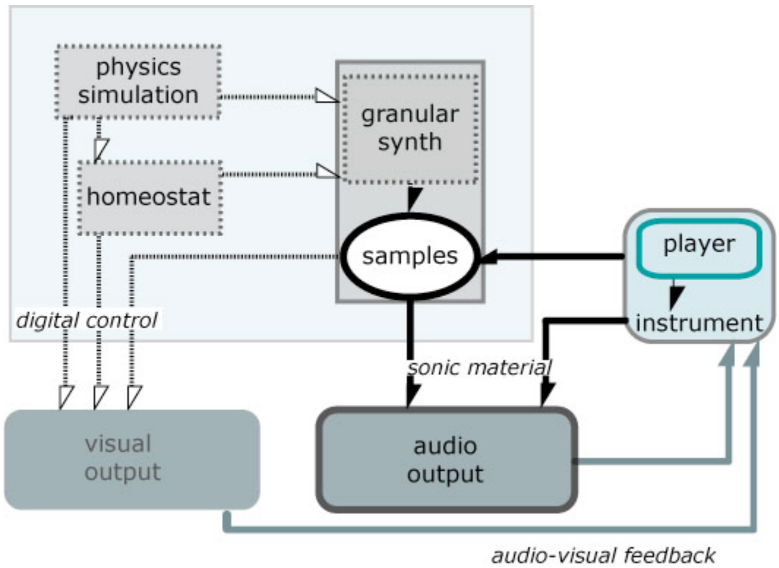

Fig. 5 Diagram of influence for the whole performance system.

The performer provides the sonic material on which the system operates. These raw samples are processed by a granular synthesis engine which is controlled by the outputs of the homeostat. This outputs directly to the external acoustic environment, to which the player responds (audio-visual feedback). The physics simulation parameterises both the homeostat and the granular synth on a longer time scale, creating higher level structure (digital control). The visual output is derived from the physics simulation, homeostat and samples and also influences the player's performance.

3.1 Outline of system components

The homeostatic network is based on the system described by cybernetician Ross Ashby [25]. This acts as a responsive pattern generator, controlling and triggering the sampling engine to create digital re-compositions of the musician's acoustic improvisations. Multiple audio samples are taken during the performance, and the output values of individual units in the homeostatic network are used to control the granular synthesis engine, determining *when* sound grains are triggered and from *where* in the sample they are taken. Different grain sizes and densities vary the acoustic/electronic or melodic/rhythmic feel, creating digital re-interpretations or timbral reflections of the performer's improvisations.

The motion-collision equations in the physics simulation describe the movements of various objects in a virtual space which is depicted in an accompanying video projection (Fig. 5). These are used to parameterise both the granular synthesis engine and homeostat network, and act to generate longer term structure. One set of equations describes the movements of two separate cell-like aggregations. The motion dynamics of each aggregation control parameters such as windowing length, volume and density in the synthesis engine and the 'viscosity' of the homeostat. Another set of motion equations describes the fixed trajectories of the three white bubbles shown in Fig. 6. As each collides with the left and right boundaries the initial sample is triggered (forward and reversed) at a speed determined by the length of the trajectory, creating polyphonic drones. Collisions between the bubbles and the cells perturb the homeostat, forcing it into new trajectories, resulting in different state dynamics, which create new re-compositions of the original acoustic material.

The visual output is based on the outputs of both these systems (see Fig.6 caption), providing an abstract representation of the current state dynamics of the digital generative system. This acts as both a cue for the performer (for example allowing anticipation of imminent changes in state-dynamics caused by collisions) and an invitation to the audience to consider the origins of the sounds they hear.

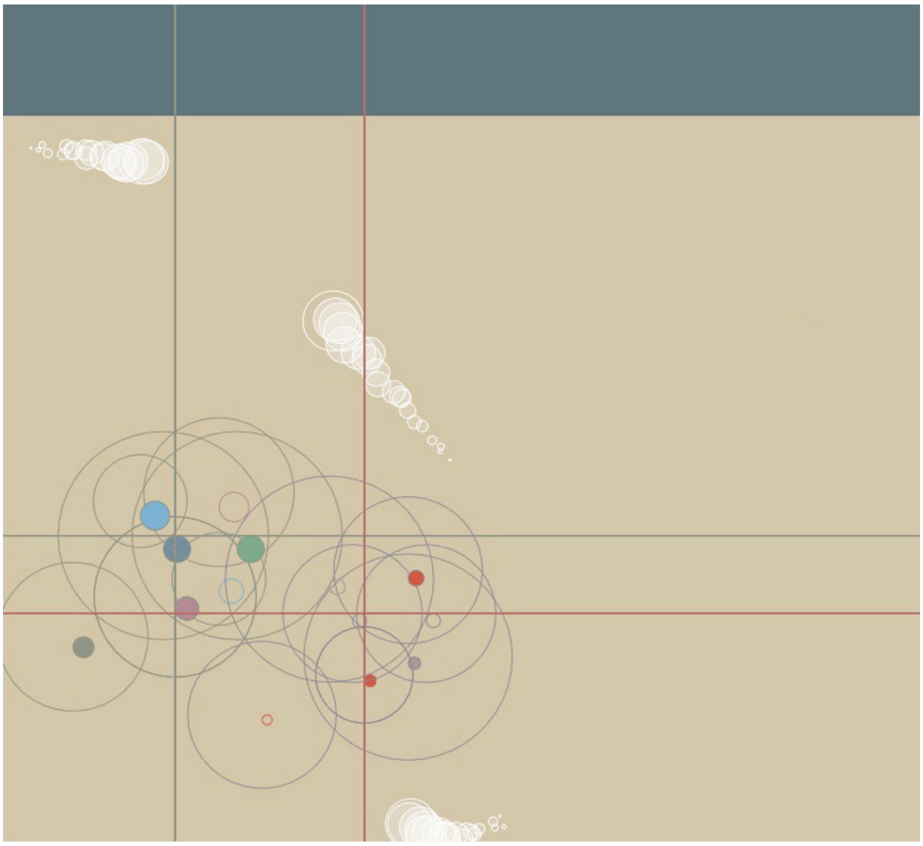

The cross-hatches mark the spatial position of the centres of two cell-like aggregations which move according to the motion equations. The large rings represent stored samples. The coloured centres are positioned in this sample space according to the current outputs of the homeostat, representing the sounds that are voiced. The white bubbles trace fixed trajectories, and trigger the playback of the first retained sample. Collisions between the bubbles and cell-like aggregations perturb the homeostatic system, forcing it into new trajectories which create new musical constructs.

Fig. 6 Screen shot of Fond Punctions video projection.

3.2 Design rationale

Although obviously not created to replicate any particular musical style, the system was designed to capture broad musical characteristics. A sample based system was chosen for the reasons given above. A sampling interface also makes the system intuitive, yet leaves room for more sophisticated musical expression. Using granular synthesis techniques specifically means that the system has the potential to explore a full range of electro-acoustic timbres, which seems aesthetically apt for an exploration of man-machine improvisation. By varying the windowing function size and shape, playback speed, and grain density, a wide spectrum of sounds are achievable, from the original timbre of the sampled instrument, through recognisable distortions of it, to the grainy textures typically associated with granular synthesis.

Two different models were used in the generative engine in an attempt to build a simple hierarchical structure. As the melodic and timbral material comes primarily from the performer, a low level patterning algorithm capable of producing interesting and varied rhythmical behaviours was sought. Previous experimentation suggested that the dynamics and properties of the homeostatic network would be suitable for this task. The system is capable of producing qualitatively different rhythmical patterns: state dynamics vary from periodic to wild oscillations. These states are controllable as the system is parameterised by a single 'viscosity' variable (in Ashby's original mechanical system, this was literally the viscosity of the liquid in which extensions of each electromagnetic component trailed). Low values, create a 'stiff' system, which settles rapidly to cyclic states, producing repetitive musical patterns, high values promote wilder oscillations, creating more exploratory musical behaviours. In addition, the network responds to environmental input. Perturbations cause a departure from the current attractor, and depending upon the severity of perturbations the system will either return to its previous path or reach a new stable state. In this implementation, these perturbations are caused by collisions between objects in the physics simulation. As these are shown in the visual display, the audience gains a simple but powerful sense of the adaptive nature of the system as visual collisions trigger a change in musical behaviour.

A simple physics engine was used as a 'container' for the other components, as motion equations offer an accessible metaphor which is realised visually and aurally to create simple compositional structure. Taking a common metaphor of an increase in height to represent an increase in energy, as the cells rise up, the viscosity of the homeostat increases, creating more exploratory musical behaviours, and the volume and pitch of the granular synthesis engine increase. The physics engine can be easily configured to create different compositional structures for specific performance/ installation needs. In this implementation, a hand-coded switch is added such that after the second ring hits the surface, (when musical 'energy' is at its maximum), the gravitational force is increased, causing both rings to descend (the sounds slowly 'dying out' to a natural end). If no new samples are added during this descent the system is left to 'explore' possible counterpoints of previously-played material. This configuration was designed to create a loose compositional structure which although simple, is effective in presenting a temporal exploration of the state dynamics of the generative system, bringing some musical variation and arguably rudimentary development of musical ideas.

4. Discussion

Playing with this system, provides an exciting mix of influence and inspiration: you are at once twirled around un-foreseen musical spaces, yet able to shape the dance in musically interesting ways which seem to be appreciated by the audience. The system arguably takes a first step toward relaxing the constraints of the programmer and opening up the space of possibilities for generative performance. To what extent then, does the system fulfil the desirable characteristics suggested in 1.3?

Independence of the generative process. The way in which the algorithmic generative systems is employed here is comparable to many existing generative works in that the dynamics of an automated process act essentially as pattern generators whose outputs structure a given medium [18, 26]. The attraction of this approach within the musical domain, as in all generative art, is the prospect of creating novel musical material, of relinquishing control to a headstrong system whose outputs lead us into unimagined spaces.

There have been fairly extensive explorations of abstract, or 'extra musical' generative systems in the musical domain, but it is of note that the results typically wander around spaces which do not capture the musical imagination of the wider audience: outputs generally make it into the studio or onto the stage only once fragments have been re-composed under more traditional composition methods [27]. It has been suggested that one reason for this is that extra-musical generative systems 'lack the cultural references that we normally rely on when appreciating music' [18]. The approach described here represents one way in which we can bring the "sublime computational poetics" [2] of generative systems onto existing stages.

Intuitive yet masterable means of interaction: Operating the system is fairly intuitive: the simplest click or whistle into a microphone is sufficient to create rudimentary compositions. An installation version of the system was recently shown at this years big blip (a festival of creative arts, science and technology), and succeeded in grabbing and holding the attention of children as young as five years. The sample–based interface is instantly accessible to the musically-naïve. With a bit of musical foresight, however samples can be taken which will facilitate the creation of fairly complex counterpoints. Musicians visiting the exhibition, developed significantly more sophisticated creations. With additional practice, it is also possible to create longer-term harmonic structure (something that traditional loop-sampler performances lack). Modulations can be achieved by feeding in harmonically ambiguous samples at crucial times and only submitting clarifying notes once you have changed key. The musical sophistication and compositional complexity can also be enhanced by increasing the sample length, allowing manipulation of larger musical ideas. The usefulness of the system as a creative tool is currently being investigated by distributing it amongst musicians of various backgrounds.

Inexhaustible and extremely variable outcomes. In the current version, the dynamics of the homeostatic network and physics simulation are affected by each other, but are not influenced by the performers actions[3]. The user cannot directly steer the course of the system. However because some long windowing functions are used in the granular synthesis engine, the rhythmic and melodic structure of the player's input is preserved. This means that although the dynamics of the structuring process itself cannot be modified, the musical structure of the material it organises is preserved. The possible musical outputs then are open-ended as they reflect both the musical structures created by the performer and the dynamics of the digital.

---

[3] Developments toward this end are currently underway.

As an improvisation, the player's actions (both what to play and what to sample) are strongly influenced by the digital sonic output of the system. The human contribution to genotype and enaction mechanism is moderated by the environment. This system does *not* answer the ALife search for *automated* robust modification of genotype and enaction mechanism, nor does the digital system itself exhibit creative emergence. However, in making the enaction mechanism a man-machine collaborative effort, the advantage that creatively emergent systems offer in expanding the possible phenotypic space is fulfilled by the musician who themselves introduce new musical primitives. This approach does not answer an epistemological research agenda, but offers a practical solution to an artistic need.

5. Conclusion

By enjoining man-machine collaboration at the heart of the generative process, the approach taken in Fond Punctions offers one way of opening up the artistic possibilities and liberating the user's freedom of expression in creative generative systems. The inclusion of a human influence on the realisation of the generative process undermines the research agendas of many related areas (e.g. automated composition and artificial creativity focus on ways of *removing* the human from the creative loop), and sidesteps key ALife issues surrounding computational emergence. But if our concerns are primarily creative, do we need to take on these hard epistemological issues? The approach taken here seems to facilitate creative exploration and expression of the computational sublime in a man-machine collaboration without having to wait for an analytic understanding of complex adaptive systems. Is seems possible that by playing with these systems in musical and other artistic ways, we may gain insights into their behavioural dynamics which evade us when we sit staring at the computer screen. If these insights led to the generation of testable hypotheses, we could begin to take seriously increasingly common propositions of 'performance as research'[4]. There are things you can only learn about someone by dancing with them.

References

1. Rinaldo, K. E. 1998. Technology Recapitulates Phylogeny: Artificial Life Art. *Leonardo* v.31 no.5.
2. McCormack, J. 2004. Impossible Nature: The Art of Jon McCormack, Victoria, ACMI.
3. Whitelaw, M. 2004. Metacreation: Art and Artificial Life Cambridge, Mass.: MIT Press.
4. Biles, J. A. 1994. GenJam: A Genetic Algorithm for Generating Jazz Solos Proceedings of ICMC
5. Biles, J. A. 2001. Autonomous GenJam: Eliminating the Fitness Bottleneck by Eliminating Fitness. In *Proceedings of the 2001 Workshop*, San Francisco.
6. Bird, J. and Webster, A. 2001 The Blurring of Art and ALife. In Dorin, A. (Editor) *Proceedings of Second Iteration, CEMA, Melbourne*. pp38 - 46
7. J. Bird, J. Layzell, P. Webster, A. and Husbands P. 2002. Towards epistemically autonomous robots: Exploiting the artistic potential of physical systems *Leonardo* Vol. 36, No. 2

[4] Raised most recently at one of the Interactive mind symposiums held in Liverpool. Mike Wheeler personal communication

8. Dorin, A. 2001. Aesthetic Fitness and Artificial Evolution for the Selection of Imagery from The Mythical Infinite Library", *Advances in Artificial Life*, Proc. 6th European Conference on Artificial Life, Kelemen & Sosik (eds), Prague, Springer Verlag pp. 659-668

9. Borges, J.L. 1970. *The Library of Babel in Labyrinths* Harmondsworth: Penguin, pp 78-86

10. Levin, G. 2000. Painterly Interfaces for Audiovisual Performance. *M.S. Thesis,* MIT Media Laboratory.

11. Todd, S. & Latham, W., 1992. *Evolutionary Art and Computers*. Academic Press.

12. R. M. Smith and M. A. Bedau, 2000. "Is *Echo* a Complex Adaptive System?" *Evolutionary Computation* 8 (2000): 419-442.

13. Bedau, M. McCaskill, J.S. Packard, N.H., Rasmussen, S., Adami, C., Green, D.G., Ikegami, T., Kaneko, K. and Ray, T.S. 2000. Open problems in artificial life, *Artificial Life* 6 (2000): 363-376:

14. Cariani, P. 1992. Some epistemological implications of devices which construct their own sensors and effectors. In F. Varela and P. Bourgine (ed.), *Towards a Practice of Autonomous Systems*. MIT Press, Cambridge, Massachusetts

15. Pask, G. 1959. Physical Analogues to the Growth of a Concept in *Mechanisation of Thought Processes: Proceedings of a Symposium held at the National Physcial Laboratory on 24-27 Nov, 1958* (London: H.M.S.O)

16. Layzell, P. 2001. Hardware evolution: On the Nature of artificially evolved electronic circuits (Ph.D Thesis, School of Cognitive and Computing Sciences, university of Sussex)

17. Dorin, A. and McCormack, J. 2001. First Iteration, Generative Systems Guest editors introduction. *Leonardo*, 34(3), p. 335

18. Miranda, E.R. 2003. On the Music of Emergent Behaviour: What Can Evolutionary Computation bring to the Musician*? Leonardo* 36 (1) pp.55-59

19. Berry, R. Percy's Organ. 1986 http://www.abc.net.au/arts/adlib/stories/s885079.htm15 (visited 30.08.05)

20. Montag, D. 2000. *Bioglyphs: Generating images in collaboration with nature's events* - PhD Thesis University of Hertfordshire

21. Rinaldo, K. 2000. The Flock http://accad.osu.edu/~rinaldo/works/flock/flock.html (visited 30.08.05)

22. McCormack, J. 2005 Open Problems in Evolutionary Music and Art, Lecture Notes in Computer Science, Volume 3449. pp 428 – 43618

23. Dawkins, R. 1986. *The Blind Watchmaker*. Longman, Essex.

24. Eldridge, A.C. 2005. Fond Punctions: Generative Processes in Live Improvised Performance. *Proceedings of Generative Arts Practice*. Sydney. (submitted)

25. Ashby, W.R. 1960. Design for a Brain. Chapman and Hall, London

26. Dahlstedt P. and Nordahl M.G. 2001. Living Melodies: Coevolution of Sonic Communication. *Leonardo*. 34(3) pp. 243-248(6)

27. Miranda, E.R. 2001. Entrée o Absurdo e o Misterio porto Alegre: Goldberg Edicioes Musicais.