Fond Punctions:
Generative processes in live improvised performance.

Alice Eldridge

Fond Punctions is an improvised audio-visual performance for 'cello and real-time generative processes. Adaptive mechanisms are used to control the granular synthesis of acoustic material sampled during the performance, producing a musical confluence of live 'cello and cybernetic extemporisations. This paper describes the system and contextualises the approach, suggesting that the use of generative processes in live performance offers exciting practical and creative possibilities for both performance and generative art practices.

Fond punctions was inspired by and amalgamates two separate personal desires: as an artificial life researcher/practitioner, to share the excitement and unnerving sensation experienced when a simulation displays life-like behaviour; as a 'cellist to be able to perform solo electro-acoustic improvisations without having to touch my laptop. This paper presents the what? how? and why? of this exploration of generative processes in live improvised performance. A description of the performance is given in section one, followed by an algorithmic outline of the system in section two. Section three contextualises the approach and explains the practical and creative motivations behind combining generative art and improvised performance practices. Technical details of the system are given in the appendix.

1. What ?

Fond Punctions is a short improvised 'cello and laptop duet based on the performer's interactions with a generative live-sampling system. The piece is based around a metaphor of artificial pond life: the digital system feeds from and disgorges a re-organisation of the acoustic environment. The performance develops from nothing. The opening 'cello gestures are sampled and spawn audio-visual 'bubbles' that rebound along trajectories in a virtual underwater space (Figure 1). Subsequent acoustic motifs are captured and emerge as cell-like structures in floating conglomerations, sonically creating rhythmic reflections of the 'cello lines. As the aggregations accumulate they rise up, overpowering their acoustic origins. When the performer stops playing, the objects begin to sink slowly to the bottom, algorithmic explorations of the stored audio material creating counterpoints on the original themes. As they sink out of audio-visual view the piece ends.

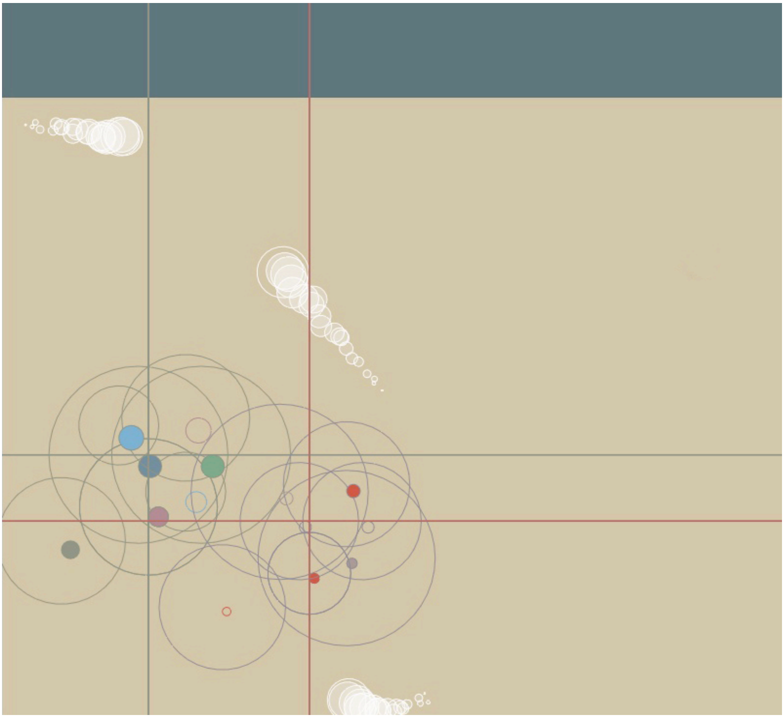

Figure 1. Screen shot of Fond Punctions video projection. The cross-hatches mark the spatial position of the centres of two cellular aggregations which move according to motion equations, confined by the limits of the fawn-coloured area. The large outlined rings represent individual stored samples. The coloured centres are positioned in this sample space according to the current outputs of the homeostat. 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 (see below for further details).

2. How ?

2.1 System overview

Algorithmically the system is based on two distinct but interacting systems: a homeostatic network and a physics simulator. These both act to parameterise a granular synthesis engine, operating on live-sampled material. Structurally the systems operate at different levels: the homeostat operates at a rhythmic and phrasal level, the physics simulator determines longer term structure at the level of musical form.

The homeostatic network is based on the system described by cybernetician Ross Ashby (Ashby 1965). This acts as a responsive pattern generator, creating poly-rhythmic re-compositions of the musician's acoustic improvisation. Multiple audio samples are taken during the performance, and the output values of individual units in the homeostatic network are used to control *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 the impression of digital re-interpretations or timbral reflections of the performer's improvisations.

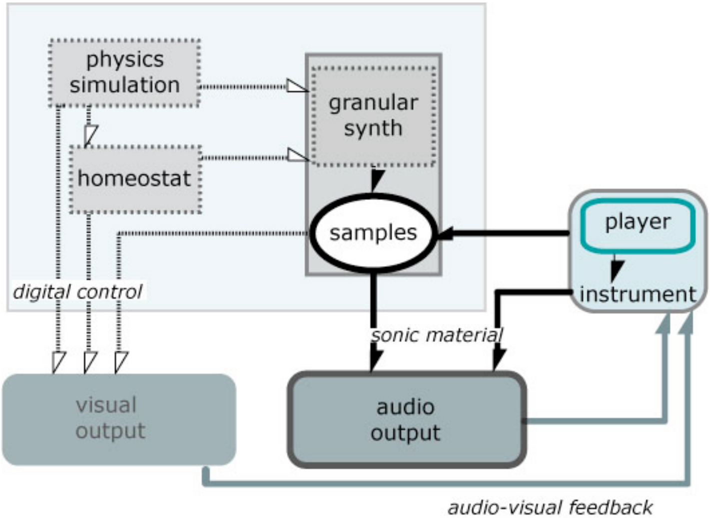

Figure 2. Diagram of influence for the whole performance system. The performer determines the acoustic material on which the system operates. These raw samples are processed by a granular synthesis engine which is controlled by the outputs of the homeostat. This outputs directly to the external acoustic environment, to which the player responds. The physics simulation parameterises both the homeostat and the granular synth on a longer time scale, creating higher level structure. The visual output is derived from the physics simulation, homeostat and samples and also influences the player's performance.

The motion-collision equations in the physics simulation describe the movements of various objects moving in the space depicted in the visualisation. One set of equations describes the fixed trajectories of the three white bubbles shown in Figure 1. As each of these collide with the left and right boundaries the initial sample is triggered (forward and reversed) at a speed determined by the length of the trajectory. This creates a polyphonic drone which shifts throughout the performance as the path lengths are incommensurate. Another set of motion equations describes the movements of the two cellular aggregations and parameterises the granular synthesis engine and homeostatic network, creating a higher-level compositional structure. Details of these mappings are given in the appendix. Collisions between the bubbles and the cells perturb the homeostat, forcing it into new trajectories, resulting in different state dynamics, which create re-compositions of the original acoustic material.

The visual output is based on an abstract representation of the dynamics of these systems and is also affected by the samples taken. The dynamics of the physics simulation control the global movement of bubbles and buoyant cell-like aggregations that twitch to the pulse of the homeostatic oscillations. Individual sample lengths determine their size (see Figure 1 caption). Details of the each component can be founding the appendix.

The performer controls only *when* to take the samples – and of course what to play, which as an improvisation is directly influenced by the sonic output of the system. In system terms, this closes the feedback loop on a macro scale; in performance terms this creates new musical ideas that trigger fresh improvisations. The system could perhaps be conceived as a 'generative extended instrument', creating counterpoints on player-defined fragments, effectively enabling the musician to perform a time-extended improvisation.

2.2 Implementation

The homeostatic algorithm is implemented in C++ as a Max/Msp[1] external. All audio processing is done in Max/Msp, the granular synth being built using on an object by Nathan Wolek[2]. Sample triggers are taken using a MIDI foot switch. The motion dynamics are implemented in Java in the Processing[3] environment in which the graphics are also produced. Communication is achieved between processes via OpenSoundControl[4] using the CNMAT updwrite object[5] in Max and Andreas Schegel's OscP5[6] library in processing. The application runs on a G4 power book running OS X tiger.

3. Why ? Contextualisation and motivation.

Generative art is often defined in terms of the creation of a process, which is left to run with some degree of autonomy, using a computer, machine or some other 'procedural invention' (Galanter 2003). The relation between artist, artwork and audience can be elucidated using the biological concepts of genotype and phenotype (Dorin and McCormack 2001) as shown in Figure 3. The artist constructs a generative process (the genotype) and typically stands back as the process unfolds in the hands of the 'procedural invention' to produce the resultant artefact (the phenotype).

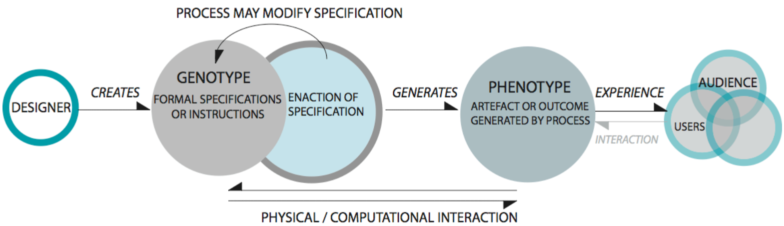

Figure 3. Interactions and influences between artist, artwork and audience in generative art. (Reproduced from McCormack (2004) with kind permission).

In some early process-based works by experimentalists such as John Cage, this desire for autonomy in the realisation of the genotype was arguably primarily conceptually motivated. Process-experiment for Cage centred on bringing about acts "the outcome of which are unknown" (Cage, 1961). For others today, such as John McCormack, the draw of generative art seems to answer a more aesthetically motivated desire: to share the experiences of 'sublime computational poetics' (McCormack, 2004).

Whilst the creation of a generative process remains central, many generative arts practitioners today explore possibilities for interaction by creating generative processes that are sensitive to environmental feedback. Pieces such as Richard Brown's Memetic Starfish (Brown 2000) demonstrate just how effective simple interactive mechanisms can be in engaging audiences: in this case by engendering attributions of intentionality, and even personality to what is really nothing more than digitally-controlled projected light.

By encouraging interaction with adaptive generative processes we can enrich the creative possibilities of generative processes in art, and in particular music. On a practical level, by bringing generative processes into interactive performance practice, we create new possibilities for man-machine performance. In addition, adding an element of interactivity to generative music within a traditional improvisation framework, offers a means of 'humanising' the machine aesthetic, which arguably constrains the accessibility of some generative music.

3.1. New possibilities for man-machine improvisation.

As a performing instrumentalists the possibilities offered up by advances in digital technologies are tempting. But in reality, when sitting or standing behind a 'cello, bass, or any instrument with both hands fully deployed, it is physically awkward and invariably musically disruptive to turn to the track pad and keyboard of a laptop. The practical motivation behind the current system then, was to develop a 'hands-free' performance system or extended instrument that was both flexible enough for improvisation, but reliable enough for live performance.

A range of exciting new approaches to electro-acoustic, or man-machine improvisation are constantly being developed, however these invariably require a human to sit at the controls of the electronic system. It is very hard to pre-programme digital systems that both avoid repetitious tedium and can be 'trusted' to behave appropriately in a live musical setting. The beauty of generative systems is that they allow a designer to compose a space of possibilities in which the machine is free to roam. Some regions of the space may be richer than others, but the use of simple adaptive generative mechanisms seems to provide a workable balance of reliability and unforeseen inspirational novelty.

In the system used in Fond Punctions, part of the control comes from simply constraining the limits of tempo and pitch range. More importantly, the performer also controls which musical fragments are sampled during the performance. This offers greater creative control over the final outcome of the system than is afforded by many generative systems. Typically the generative process (Figure 3) acts to structure a pre-defined medium – whether pixels, MIDI notes, or robotic behavioural repertoire. In the current system the material is determined by the performer *whilst* the genotype unfolds, so the selection of phenotypic 'substance' becomes a dynamic, artistically defined part of the generative process.

Novelty is achieved, as although the performer determines the samples, the system does not play them back verbatim as a traditional sampler would. Instead, as described in section 2.1, the dynamics of the homeostatic system make selections from this sample base, re-composing fragments of the original phrases. Thus the system has enough freedom to produce fresh musical material, which can potentially inspire the musician's improvisations. The performer can never know exactly what state the homeostatic network will enter next, but because the projection reflects the dynamics of the physics simulation, it provides a visual cue. This allows the performer to anticipate the collisions which trigger the homeostat, causing a change in the sonic output of the material.

3.2. Humanising generative music

If on the one hand generative processes allow human performers to invite the computer onto the stage, then the application of generative processes within a traditional improvisation framework may represent one way of making the world of machine music more inviting to human ears.

Most approaches to generative music use 'closed' generative systems. In line with the scheme depicted in Figure 3, the designer/composer attempts to create an algorithmic specification of the musical output that is realised automatically by the machine. Such systems tend to only make it onto the stage or into the studio once generated fragments have been 're-worked' using more traditional performance methods (eg Miranda, 2001). It has been suggested that the success of such systems (in terms of creating engaging music for human listeners) is hampered by the fact that they tend to 'lack the cultural references that we normally rely on when appreciating music.' (Miranda 2003) p.1.)

The alternative of course is to attempt to embed musical knowledge directly (eg Ebicioglu, 1984), or implicitly, through a learning process (eg Cope 1991). Whilst both these approaches may produce something that is 'more recognisable' as music, they are inherently tied to existing musical styles, and leave little room for exploring the new potentials offered by generative art.

The approach taken here then, is an attempt to preserve the novel musical possibilities of generative systems but to sculpt the algorithmic outcomes into a humanly accessible aesthetic. This is achieved as described in section 2.1, by letting a human musician interact with the generative process at the enaction stage, so that the final artefact is a co-product of the human-machine process. This approach seems to be one way in which we, as generative arts practitioners, can share new experiences of 'sublime computational poetics' within existing cultural practices.

References

Ashby, W.R. (1965) Design for a Brain: The Origins of Adaptive Behaviour. Chapman and Hall Ltd.

Brown, R. 2000. http://www.memetics.com last accessed 27/09/05

Cage, J. (1961) Silence, Middletown, Connecticut, Wesleyan University Press.

Cope, D. (1991) Computers and Musical Style, Winsconsin, Madison.

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

Ebicioglu, K. (1984) An expert system for Schenkerian synthesis of chorales in the style of J.S. Bach. In Proceedings of the 1984 International Computer Music Conference San-Fransisco: ICMC, pp.135-142.

Galanter, P. What is Generative Art? Complexity theory as a context for Art Theory. Proceedings of Generative Art 2003, Milan.

McCormack, J. (2004) Impossible Nature: The Art of Jon McCormack, Victoria, The Australian Centre for the Moving Image.

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

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

Links for resources

1. Max/MSP: http://www.cycling74.com/index.html
2. Nathan Wolek: http://www.nathanwolek.com/
3. Processing: http://www.processing.org/
4. Open Sound Control: http://www.cnmat.berkeley.edu/OpenSoundControl/
5. OSC Max Objects: http://www.cnmat.berkeley.edu/OpenSoundControl/Max/
6. OscP5: http://www.sojamo.de/

Appendix – System details

I.I Homeostat

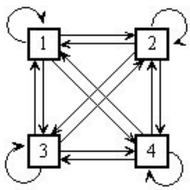

$$Output_{i(t+1)} = \sum_{j \in c} Output_{j(t)} \times Weight_{ij(t)}$$

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

Figure i. Schematic of 4 unit homeostat and update equations (right)

System description

The homeostat is modelled on cybernetician Ross Ashby's system of the same name, which was built to illustrate how machines could be at once fully deterministic and yet adaptive (Ashby (1943)). Like all homeostatic systems, internal stability is maintained in the face of environmental perturbation through internal reorganisation. Algorithmically, the output of each unit in the system is determined by the weighted outputs of every other unit including a recurrent connection (Fig.i), creating the necessary feedback loops. The system is considered to have been critically perturbed whenever the output of any one unit exceeds a pre-specified critical value. Reorganisation is achieved by randomising the weighted connections between each unit. So by a process of trial and error, the system invariably achieves a stable state. When stable, the output values of each unit typically settle to a N-point limit cycle.

The whole system is parameterised by a 'viscosity' variable which controls the overall stability of the system (in Ashby's original electro-mechanical device this was literally the viscosity of the liquid in which extensions of the electro-magnetic components trailed). Low viscosity values therefore produce a 'stiff' machine which more commonly and more rapidly settles to a steady state, higher values produce wilder oscillations and the system typically takes longer to achieve stability. In this digital implementation, the viscosity variable also affects the output values themselves, as for a certain mid-range, outputs tend to be constrained to the viscosity value itself. This is utilised to musical ends as described below.

I.II Mappings and the Granular Synthesis Engine

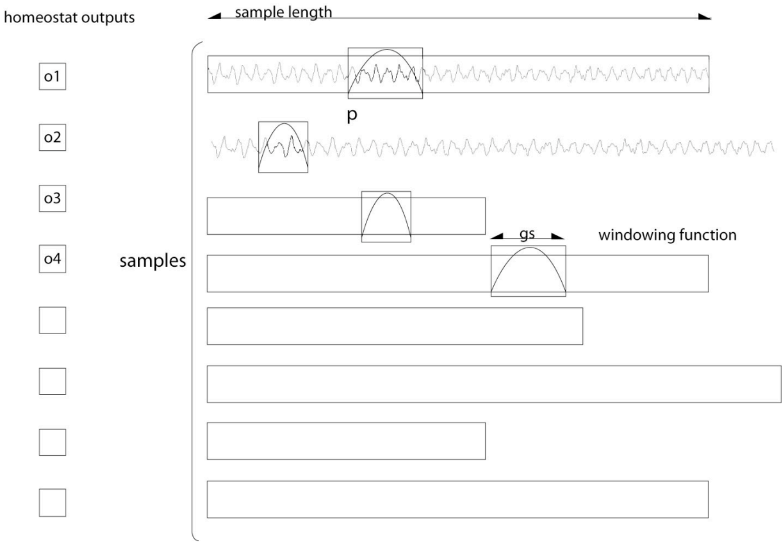

Figure. ii The outputs (O_n) of the homeostat are used to determine to point P in the sample that grain is taken from.

A granular synthesis engine (GS) was implemented in MAX/Msp using Nathan Wolek's gran object. During the performance, the length of each sample is stored and the current output range of the homeostat is mapped dynamically to the individual samples. Up to 8 different samples are held at any one time, (typically these are 5 – 20 secs long although nothing prevents times outside this range) and can be overwritten throughout the performance. The outputs of each of the 8 units in homeostat are used to trigger sound grains taken from the stored samples.

The stored samples 1-4 are read by 8 gran objects with grain sizes in the range 400:2000 ms, at original pitch. This preserves the pitch and timbral characteristics of the original sample and for higher values even melodic/rhythmic fragments can be recognised. The position of the grain in each file is determined by the output value of individual homeostat units. Grains are triggered whenever the output of the hom is NOT equal to the max change.

Stored samples 5-8 are read by 8 gran objects operating at smaller grain lengths in the range 90:300ms and variable (4x – 32x) pitch. This produces the pops and clicks characteristic of sparse granular streams. These are triggered whenever outputs of the same homeostat are equal to max change. This typically occurs for mid-range viscosity values which create wild oscillations.

Rhythmic Enhancement

This second set of samples produce a more rhythmic texture which is enhanced with some simple stochastic elaboration. In order to avoid very repetitive rhythms, a probabilistic filter is used: a random no between 1 and N is selected, and only the nth trigger will cause an output. N is reset each time a collision between objects occurs. An opposing process creates a higher density of and variation in rhythmic output: delay lines are set on half the triggers, so that when a trigger does arrive, it is duplicated at varying fractions of the regular beat.

In this implementation, the network consists of 8 fully connected units and the update time of the algorithm is slowed to around 200ms, so that the system acts a responsive pattern generator. This creates a rhythmic regurgitation of the original sounds produced by the performer. The effect is in some ways akin to a remixing-loop sampler. The 'remix' being determined by the state dynamics of the homeostat system.

I.III Physics Simulation

The physics engine describes the motion of various objects. Two aggregations of rings move around a finite space, rebounding off the perimeters, and colliding with another set of 3 bubbles which traverse fixed paths described by simple functions (eg sine, quadratic). For performance purposes, this system is shaped to provide an overarching compositional form: initially a cluster of cells (C1) accumulate as samples are taken, rising with each new sample. These represent the 'melodic' sound files 1-4. The second set of cells (C2), representing rhythmic files 5-8, appear only once the first set have hit the surface. When both aggregations hit the surface, the 'buoyancy' of the simulation switches so they start to sink.

The visual display reflects the dynamics of the physics simulation, but also the homeostatic network, and aspects of the performer's actions. Variables defined by the physics simulation also control parameters of the granular synthesis (GS) engines and events act as environmental input to the homeostatic network as detailed below.

Simulated height of C1 centre	controls	viscosity of homeostat Grain length of GS1 Grain amp of GS1
Simulated height of C2 centre	controls	Grain length of GS2 Grain amp of GS2
Simulated collision	controls	Colour change in visuals Homeostat perturbation Stochastic rhythm change
Outputs from homeostat	controls	Position of grains in sample Position of cell centres in visuals
Sample length	controls	Size of cell boundaries in visuals