Autonomy, Influence and Emergence in an Audio-visual Ecosystem

P. Beyls, PhD
LUCA School of Arts, Brussels, Belgium, &
School of Arts, University College Ghent, Ghent, Belgium
www.luca-arts.be, www.kask.be, www.beyls.org
e-mail: peter.beyls@telenet.be

Abstract

This paper provides an informal introduction to Petri, an interactive audio-visual installation built as a society of communicating virtual agents developing autonomous behaviour but equally sensitive to outside human-originated influence. From a systems point of view, this work consists as a dynamic interface connecting life in an artificial society with life in a given tangible physical environment. Design principles include the synthesis of sustained forms of emergent audio-visual/behavioural complexity over time, short-term complexity from simple local interactions and long-term complexity made available through genetic optimisation. Petri suggests that partial understanding of systems behaviour blended with the accommodation of surprise makes for a rewarding aesthetic experience.

1. Introduction

An early morning walk in a forest is a rewarding experience to most of us; one participates in a completely autonomous audio-visual universe acting as a self-sustaining biotope. The quality of the experience is modulated by a delicate blend of recognition and surprise; one becomes an unobtrusive kind of “discreet participant” in an autonomously unfolding (natural) world. The theory of anticipation (Huron 2006) applies – the intensity of the individual experience is highly conditioned by a dynamic process, i.e. the continuous cognitive evaluation of the relationship between expectation and surprise. In fact, the forest metaphor promotes constructive principles for designing artificial ecosystems: (1) a continuous scale from microscopic to macroscopic emergent complexity seems basic to natural environments, and (2) natural ecosystems are responsive to perturbations and gracefully accommodate disturbance by external forces such as human beings, natural artefacts remain flexible in terms of behaviour and metabolism.

Petri characterises human-machine interaction as a process of mutual influence. The outside world as considered an “environment”, a physical world where, typically, multiple people exercise spontaneous body language as well as engage in natural social interaction. *Petri* explores the aesthetics of interaction by way of direct implementation of a virtual world developing autonomous activity, however also open to external influence. *Petri* is conceived as a large-scale audio-visual installation; a projection made available as a dynamic interface to be appreciated by humans. As depicted in figure 1, we think of two co-existing environments being interfaced; (1) a synthetic universe formalised in computer algorithms and (2) a physical (public) real-world space. Many parallel trails of activity in both worlds coalesce as macroscopic human-centred aesthetic experience. This approach contrasts with call-and-response systems where user actions trigger a selection from a fixed palette of responsive options.

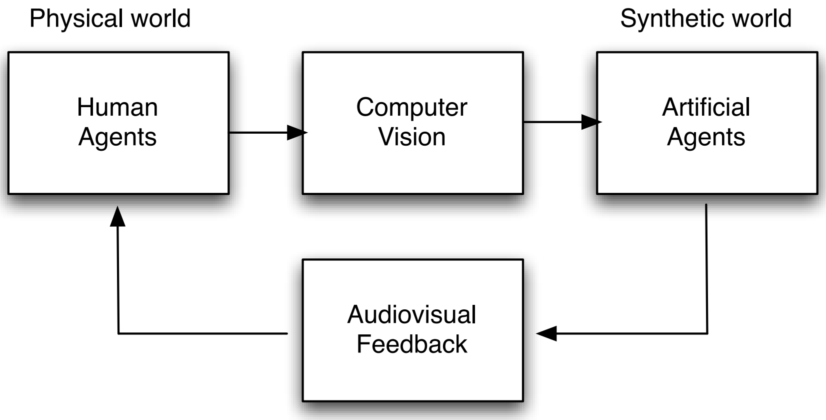

Fig.1. Global systems organisation.

Meaning is created explicitly by a clear sense of correlation between action and response. *Petri* views interaction as cognitive interactivity: the psychological, emotional and intellectual *participation* between a person/a crowd and a system. *Petri* implies implicit interaction in the sense that the system is (1) considered a micro-world of interacting components in itself and (2) a global system behaving in relation to human influence. By definition, such systems abandon the notions of authority and control in favour of a distributed systems architecture supporting life in shared biotope.

Also, the work reported here explores the concept of a reciprocal system where user behaviour exerts influence over an otherwise completely autonomous system. The net result is complex emergent behaviour in the agency merging seamlessly with spontaneous human bodily behaviour. This study suggests that rewarding human-machine interaction is sustained by a delicate interplay of prediction and surprise while appreciating complex emergent systems behaviour. Petri is a typical example of emergence because simple, local interactions between system components (called “particles”) give rise to complex emergent global systems behaviour.

Another design principle is the quest for “diversity”. In a television interview on Dutch TV, while referring to the massive variations in size and shape of shells of land snails Harvard palaeontologist Stephen Jay Gould defined beauty as an enjoyment of evolution-based variations and change; “For this is beauty for evolutionary biologists because we love diversity”. The notion of emergence teaming up with a simple genetic algorithm provides for the maximization of diversity in Petri. As noted by Dorin (2005), diversity in biological organisms is one trait that recommends the ecosystem artistic composition; ecosystems guarantee the autonomous production of novelty. Evolutionary pressure instructs organisms to develop novel structures and behaviours to improve reproductive ability.

Rewarding human-machine interaction is viewed as an interface between two universes, both holding agents engaged in social interaction. Humans interact according to salient social rules, while particles interact in a synthetic universe managed by invented rules. Both universes exhibit self-organising complexity and their confrontation reveals a particularly rich platform in support of spontaneous aesthetic human-machine interaction.

2. Related work

Some early work in agent-based interactive environments involved animated avatars claiming the suspension of disbelief as a key ingredient of rewarding human-machine interaction; the *ALIVE* project is highly exemplary (Maes 1997). Other work necessitates an immersive relationship with the human participant fully engaging the senses; the *Emergence Engine* (Mendelowitz 2000) supports interaction through voice, hearing and touch. Obviously, the suspension of disbelief requires the implementation of reasonably realistic real-world physics. In addition, the works present themselves as a visualisation of a kind of artificial landscape inhabited by virtual creatures. The project introduced in this paper avoids real-world connotations in favour of total synthetic authenticity.

Various interactive music systems are take inspiration of Reynolds' flocking model (1987), including *Swarm* (Blackwell 2004). A mapping function specifies a link between the spatial position of agents and features of the environment – it thus implements the general statement of stigmergy. For instance, an external sound may function as a temporary attractor while the flocking rules still force self-organised behaviour.

Two excellent examples of audiovisual ecosystems are *Eden* (McCormack 2001) and *Diseases Squared* (Dorin 2005). Eden is a reactive cellular virtual world, driven by a genetic algorithm using an implicit fitness function – systems behaviour acknowledged by participants moving in space receive higher fitness ranking. *Eden* evolves according to audiovisual cues from an audience and its implementation takes heavy inspiration of John Holland's *Echo* system.

Diseases Squared is conceived as a virtual machine for artistic pattern creation supporting multi-scaled temporal complexity, autonomous production of novelty, susceptibility to constraints and maintenance of coherence and unity. Bio-diversity is seen as way to enhance novelty without however destroying coherence. This implies that the system must function within a range of possibilities, holding a balance between sufficient novelty and sufficient coherence. The extremes of novelty and coherence are related to respectively high mutation rates and convergence in the context of evolutionary algorithms. In addition, in the context of artistic creation, the extremes are viewed as randomness and uniformity. Dorin makes a connection by suggesting that the implicit fitness function, which aims to maximise reproductive success in a biological ecosystem, may work as a method to navigate between novelty and coherence.

This project takes additional inspiration from distributed thinking in the fields of cognitive science (Minsky 1985), the flocking model devised by Reynolds (1987) and the discipline of artificial chemistry. In addition, *Petri* takes its name from the shallow *Petri* cell culture dish commonly used in the field of microbiology and builds on previous research viewing human-machine interaction as interfering with colliding molecules (Beyls 2005).

3. System description

Petri's universe consists of a collection of animated particles moving about in a virtual world while also interfaced to the real world by way of computer vision. The implementation is fully object-oriented, let us first address the nature of the Particle Class.

A particle moves in 2D space (with bouncing walls) and is sensitive to neighbouring particles as well as human interactors. In addition, particles consume energy and occasionally enter a sleep cycle for some time and wake up with full energy later on. Velocity and angle of movement are important additional instance variables. The gender variable plays a significant role in the reproduction process and holds four different values; a lookup table specifies specific options for particles to breed.
All particles hold a data structure to control a FM software synthesiser defined and instantiated in SuperCollider (Wilson et al. 2011).

Particle display is conditioned by the current values of various instance variables such as status (awake or asleep), sensitivity and the number of neighbours.
Sensitivity to human activity is an adaptive feature; when within the zone of influence of an external interactor (or possibly a group of people) sensitivity is scaled up else it is scaled down. At specific moments in time, visualisation is further conditioned – for instance, when a nascent particle is about to be bred in Petri's world.

When outside the zone of influence, a particle engages in a process of competition with one of its neighbours if that potential neighbour itself is outside the zone of influence and it has at least one neighbour. A particle will delete all its neighbours that are of a different gender and are currently in sleep state.

A particle continuously interacts when awake (status = 1); when the distance to any other particle is within range of the sensitivity parameter, it increments its angle of movement and slightly lowers it velocity.

A given particle will influence the behaviour of any neighbouring particle when (1) it is currently awake, and (2) the distance to a potential neighbour is lower than its neighbour-sensitivity instance variable. A particles slows down when interaction and when its velocity is lower than a given threshold, it enters a sleep state (status = 0) for a given number of process cycles. When waking up later on, maximum velocity is restored and signalled by a local blink in the graphics display.

The World-class object in Petri holds a variable collection of particles (between 50 and 250) and many critical system variables. The world cycles through a number of top-level procedures handling (1) inter-agent interactions, (2) computer vision, (3) accommodation of user interaction, (4) the particle breeding and elimination process and (5) sound handling via OSC to Supercollider.

Computer Vision (CV) procedures intend to connect the internal, virtual world of Petri with the dynamics of a specific 'external' world, typically a public space setting. As a sensing technology, CV offers the advantages of being non-obtrusive; it provides detailed spatial information and may be implement with great efficiency. CV is a five-stage algorithm in this project. First, a frame is captured (50 times per second) and a low-resolution (180 by 120 pixels) version is computed. Second, the difference between the contents of the last frame and the previous frame is calculated in terms of the absolute value of the changes in brightness of every pixel. Next, an image buffer is updated according to the data made available at stage two; it acquires changes in a dynamic memory structure. For every image/memory location, the data is boosted when significant changes occur while it is slightly scaled down otherwise. The net effect is that of a very much responsive yet slowly fading "image memory". Fourth, the "direction of change" in the image is computed from the comparison of the locations of the brightest pixels in the current frame and the previous frame. Fifth, at any point in time, the locations of the maximally changing pixels are considered "centroids of change", they are captured in a dynamic list and visualised as a kind of trail signalling changes in the external, physical environment.

The next software module acknowledges the information made available by CV; it will condition the particle behaviour in specific ways. Particles are attracted by changes in the physical world and tend to move towards the location where these changes actually take place. In addition, a centroidLevel variable is scaled up or down in proportion to the total amount of change between consecutive CV frames; in absence of changes the level gradually moves towards zero.

Changes in the physical environment exercise global (the virtual world as a whole) and local (at specific locations) impact; for instance, particles may reposition themselves if the centroidLevel is peaking at its maximum (i.e. a normalized value of 100%). Also, the higher the centroidLevel, the stronger the particle attraction of moving toward a target defined by the centroidPosition; the current XY-location of maximum change in the CV image buffer. Typically, the CV module senses a dynamic environment, external activity leaves a variable trail in terms of target dynamics and target location – therefore, local particle behaviour is influenced in especially subtle ways in addition to the fact that the history of changes introduces complex non-linear global systems behaviour.

The next software module handles the reproduction process. New particles are instantiated in proportion to the number of particles within the zone of human influence. A new particle will be born out of the genetic make-up of two randomly selected particles within the zone of influence. The newly bred particle appears physically close to one of its parent particles, its gender is decided from the consultation of the current reproductionTable. The contents of this table is subject to continuous evaluation: it contents is filled with random numbers when either the total number of particles drops too low or there is too much convergence towards a world holding particles of a uniform gender. This higher-level adaptive mechanism guarantees global survival of the world yet offering persistent genetic diversity – in this basic sense, the world is viewed as self-contained autonomous system merging two qualities; the sustained synthesis of novelty and elastic responsiveness to external influence.

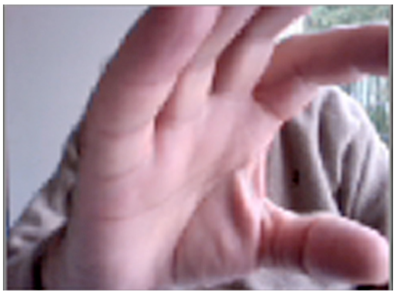

Fig. 2. Current image.

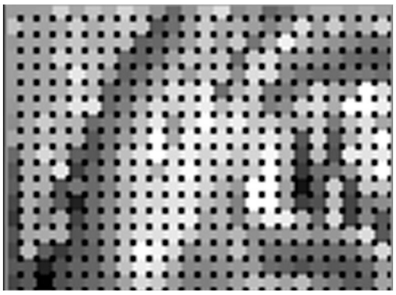

Fig. 3. Inter-frame changes.

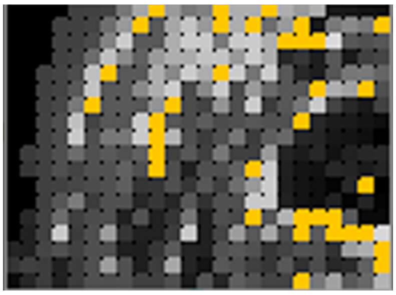

Fig. 4. Image memory.

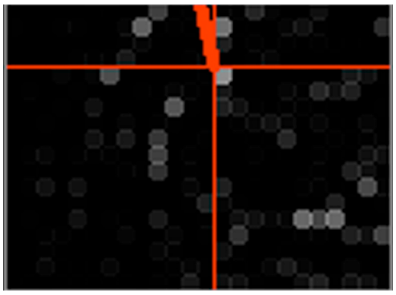

Fig5. Angle of current change.

Notice the implicit nature of our fitness function; any particle within the zone is considered a potential parent though the first two that manage to get within that zone earliest on (in a given process cycle) are effectively designated as parent objects. The reproduction process addresses the reproductionTable – a data structure holding all legitimate breeding options within a four-gender format. Offspring objects merge their parents' parametric data by way a crossover operator and mutation is added in order to avoid convergence.

A particle holds five items of parametric data considered control variables of a frequency modulation software synthesizer running on the Supercollider sound server engine; (1) carrier frequency, (2) carrier partial, (3) modulation partial, (4) modulation index and (5) amplitude. As time goes by, the evolutionary process echoes in the sound produced by the installation as a whole. In addition, some parameters may receive further articulation in the sound scheduler process, as described next.

Particles consult their neighbours in order to produce sound. Control parameters are computed from the consideration of a specific particle and its given neighbouring particles; sound is generated when (1) the total number of neighbours exceeds the minimum number of neighbouring particles threshold, (2) the neighbour has actually moved in space over the last process cycle, (3) it has not output any sound in the last 10 seconds. A mapping algorithm imposes additional articulation on the sound controlling parameters (as identified by the breeding process) the xy-position to respectively carrier-frequency and modulation-frequency – the other three sound parameters remain untouched. In other words, a genetic algorithm conditions the timbre qualities of sound (within a given epoch) while particle position exerts impact only on frequency parameters.

Many parallel dynamic visualisation processes contribute to a life-like inclusive display. Particles are visualised according to their gender, sensitivity, current state and location. The trajectory of the centroids of the circular zones of human influence is traced in 2D space. Neighbouring particles are visualised as variable clusters using line segments. Occasional super-structures emerge; when a particle has at least three neighbours, a curve object is visualised inclosing all neighbours in question.
One may think of the visualisation process as another instance of emergence; a critical mass of systems modules contribute visualised data and, while taken as a whole, all contributions coalesce to offer the perception of a complex life-like synthetic universe.

Figure 2 shows the image currently captured by the camera. Figure 3 displays the amplitude of the local differences between any two consecutive camera frames. It shows the absolute values of the changes as mapped to a grey scale image. Figure 4 reflects the accumulation of successive changes accommodated in a dynamic memory structure. Figure 5 displays the way the centroids of external activity change in the physical world; the locations of the trail are mode visible proportional to their strength and the angle (the direction of change) if equally depicted by a red vector.
As mentioned above, visualisation in Petri reflects (1) the internal values of instance variables inside individual particles and (2) follows from the local interactions between spatially neighbouring particles.

The current status of particles’ instance variables is reflected in their visualisation. For instance, when particle status equals 1 (i.e. awake), visualisation is simply a coloured dot. When within the zone of human influence and given maximum sensitivity a complementary white circle is displayed. The green circle marks the centre of the history trail of human activity.

The interaction between individual particles is also status (awake or asleep) dependent. For any particle, when awake and the distance to any other particles is smaller than its neighbouringSensitivity parameter, its velocity will slow down and, when below a given threshold, its status will switch to 0 (asleep) and this momentary event will signalled by a flashing circle. When the particles happens to be asleep and the sleep-time is over, its status will switch to 1 (awake) and this event similarly triggers a visual spark.

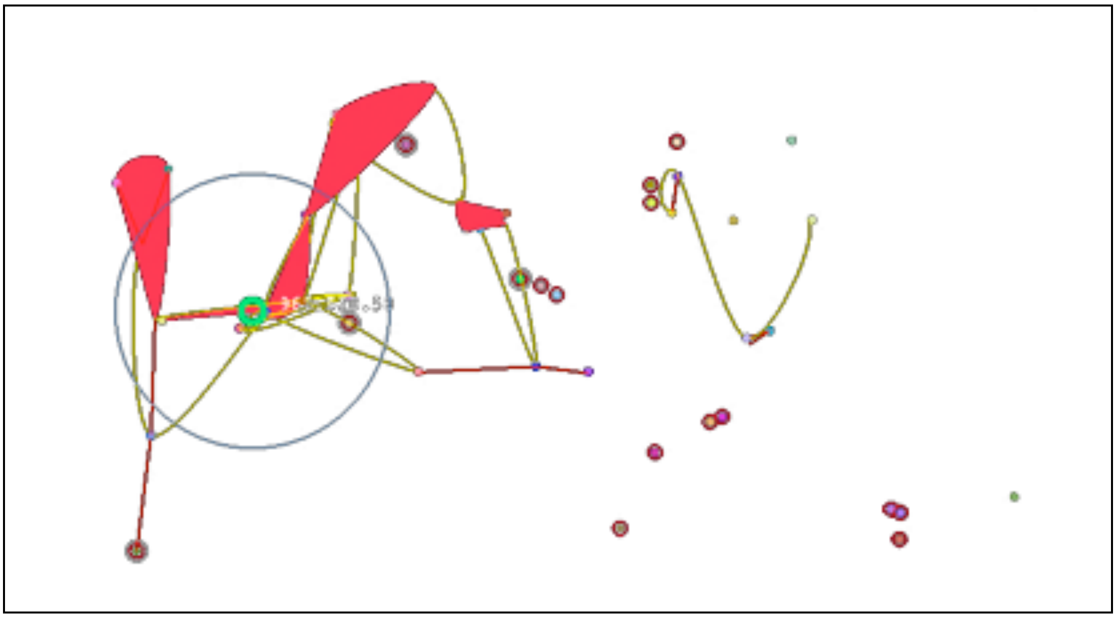

Fig. 6. Prototypical clustering of interacting particles within the zone of influence (the blue circle) (colour inverted).

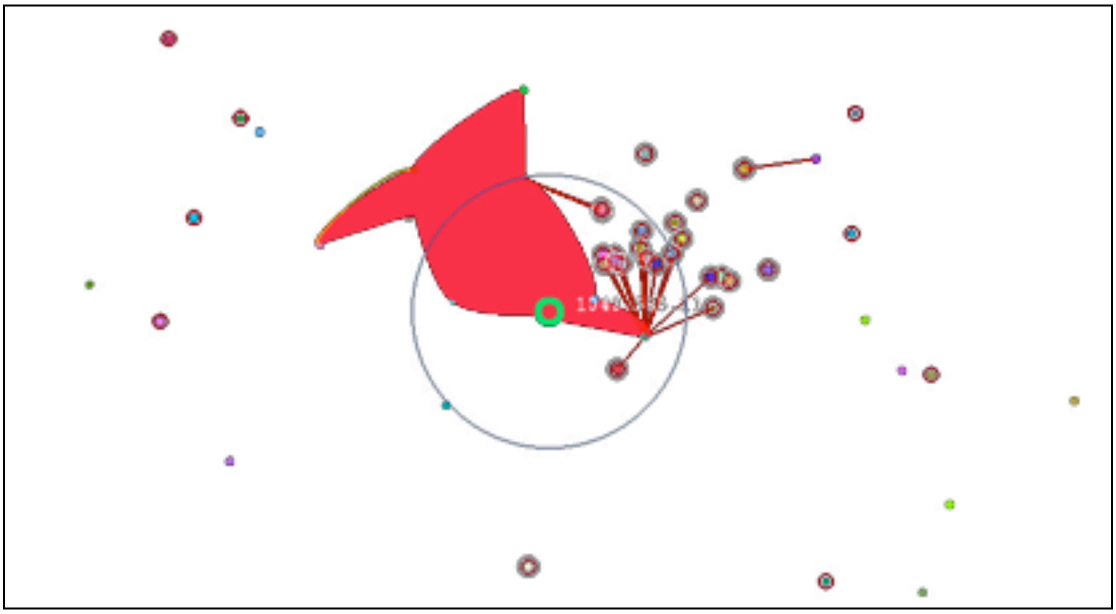

Fig. 7. Visualisation emerges from the consideration of particle status and current neighbourhood (colour inverted).

Further particle visualisation echoes its number of neighbours. When a particle features three particle neighbours, a red blob will enclose them, given more than 3 neighbours, a yellow curve will connect them. In case of just one neighbour, a line segment is drawn between the particle and its neighbour. As shown in figures 6 and 7, complex dynamic visualisations emerge as a consequence of the fluctuations in status, position and sensitivity of the current particle population in the world.

4. Discussion and conclusion

Petri can be seen as a crossing point between nature and culture in the sense that a synthetic universe is interfaced with the physicality of the natural world.

This project was first shown in Spring 2010 at the Update II international exhibition of New Media Art in Ghent, Belgium and was nominated for the Liedts-Meesen New Media Award. People express fascination by the continuous generation of audiovisual novelty while – at the same time – the system keeps a structural integrity. Remarkably, many "visitors" of the installation feel urged to capture the attention of the system by engaging in exaggerated bodily behaviour such as vigorously waving their hands. Their idea of "participation" lends towards exercising control over a self-organising process in continuous flux. Incidentally, many people expect dealing with a responsive (rather than an interactive) system much like interacting with a computer game in which clear mappings exist between user gestures and machine responses – a computer game is said to be transparent to the user. Petri does not aim for transparency but rather suggests the interaction paradigm of mutual influence.

Few systematic studies exist aiming to connect social processes in contemporary society with scripted processes in artificial societies. Human social processes are rooted in cultural histories, unconscious and mainly hidden from deliberate observation. However, spontaneous human social group behaviour might externalise dynamic patterns reflecting deeply hidden cultural logic. The work reported here is a first step in challenging that logic with invented/artificial logic though more systematic experiments are obviously compulsory.

A short note on implementation: Petri exists as two concurrent software modules. Visualisation, computer vision and interaction are implemented in Processing (Fry & Reas 2007). Real-time sound synthesis is handled by James McCarthy's outstanding Supercollider programming platform (Wilson et al. 2011). Data transfer between Processing and the sound server running in Supercollider is handled by OSC (Open Sound Control).

In conclusion; Petri exists as the suggestion of a temporal interface between a virtual, autonomous world populated by simple creatures and a specific physical world inhabited by humans. It exists as a playground and expresses faith in the idea that rewarding human-machine interaction may emerge from the articulation of human originated influence over an otherwise autonomous process. This approach is in eminent conflict with most commonly observed interaction protocols that imagine accurate control over a given process. In contrast, the project documented in this paper views spontaneous bodily behaviour of a human participant as complementary to the internal behaviour of an artificial world.

This parallel, synthetic universe is thought of as a distributed system consisting of a population of basic entities called particles. Particles interact locally using very simple rules. However, when considering the population as undivided, simple local interactions give rise to interesting, complex global behaviour that could not be anticipated by the systems designer – emergence is said to happen implicitly without the need for global explicit human-engineered guidelines.

Human “participants” (rather than “interactors”) may develop a degree of sensitivity and fractional understanding of what is actually happening inside the population of interacting particles. However, they may never really develop a complete understanding of the installation in its entirety, even given repeated visits. This particular mix of meaning and mystery acts as a source of rewarding human-machine interaction; the idea of interaction itself is extended into a profound, machine mediated aesthetic experience.

5. References

[1] Beyls, P 2005, A molecular collision model of musical interaction, Proceedings of the Generative Arts Conference, Milan Polytechnic University, Milan, pp. 375-386.

[2] Dorin, A 2005, ‘Artificial Life, Death and Epidemics in Evolutionary, Generative Electronic Art, Proceedings of 3rd European Workshop on Evolutionary Art, Applications of Evolutionary Computing, Lausanne, pp. 448-457

[3] Huron, D 2006, Sweet Anticipation, Music and the Psychology of Expectation, The MIT Press

[4] McCormack, J 2001, ‘Eden: an Evolutionary Sonic Ecosystem’, in Advances in Artificial Life, Kelemen, J and Sosik, P (Eds), Springer, pp. 133-142

[5] Maes, P Darell, T Blumberg, B Pentland, A 1997, ‘The ALIVE system: wireless, full-body interaction with autonomous agents’, Multimedia Systems, Special issue on multimedia and multisensory virtual worlds, Vol. 5 Nr. 2, Springer

[6] Mendelowitz 2000, ‘The Emergence Engine: A Behavior Based Agent Development Environment for Artists’, AAAI/IAA 2000, pp. 973-978

[7] Minsky, M 1985, The Society of Mind, Simon and Schuster, New York.

[8] Reas, C and Fry, B 2007, Processing: A Programming Handbook for Visual Designers and Artists, The MIT Press

[9] Reynolds, G 1987, ‘Flocks, herds and schools: a distributed behavioural model’ Proceedings of the 14th annual conference on Computer graphics and interactive techniques, ACM, New York, pp. 25-34

[10] Wilson, S, Cottle, D & Collins, N 2011, The SuperCollider Book, The MIT Press