McCormack, J. "Eden: an evolutionary sonic ecosystem", in J. Kelemen and P. Sosik (eds), Advances in Artificial Life, Proceedings of the 6th European Conference (ECAL), Prague, Czech Republic, 10 - 14 September 2001, LNCS Vol. 2159, Springer-Verlag, Berlin, Germany, ISSN: 0302-9743, 2001, pp 133-142.

Eden: an evolutionary sonic ecosystem

Jon McCormack

School of Computer Science and Software Engineering
Monash University, Clayton Campus Victoria 3800, Australia
jonmc@csse.monash.edu.au

Abstract. This paper describes an Artificial Life system for music composition. An evolving ecology of sonic entities populate a virtual world and compete for limited resources. Part of their genetic representation permits the creatures to make and listen to sounds. Complex musical and sonic relationships can develop as the creatures use sound to aid in their survival and mating prospects.

1 Introduction

«Man is nature creatively looking back on itself»
Friedrich Von Schlegel [1]

Music, like all complex creative endeavors, has drawn from a vast range of human experiences in search of expression. A great source of this expression is often nature itself. Over all other art forms, music seems open to 'the purest expression of order and proportion, unencumbered as it is by material media' [2], so it therefor seems natural for composers to look to artificial life and artificial nature as a source of creative inspiration.

In many ways, artificial life adopts a process-based methodology, shifting the emphasis from material to mechanisms. Formal process mechanisms have existed in music for some time and had a profound impact in music of the twentieth century [3].

This paper describes a novel music composition system that draws from artificial life techniques in its methodology.

1.1 Simulation and Composition

In discussing computer music in this paper, it is important to differentiate between music *simulation* and music *composition*. With simulation, the primary goal is for the computer to simulate an existing composer, genre, or playing style. For example, Johnson-Laird developed a system to improvise performance similar to that of jazz musicians [4]. Ebcioglu devised an expert system to generate chorales in the style of J.S. Bach [5]. As a rather crude generalization, music simulation is a decomposition problem; original compositional techniques, unique to the computer, fair better when they draw from generative or evolutionary methodologies (like Artificial Life).

Of course, it is possible to use evolutionary techniques to evolve works that mimic a particular style, and this has been the focus of much research into evolutionary composition. In their survey paper of evolutionary music composition systems [6], Todd and Werner make the important observation that in evolutionary systems it is necessary to determine which individuals are more 'fit' than others (in the case of a compositional system, which individuals are better 'composers') and hence should survive and have offspring. Todd and Werner see the fitness evaluation as an integral part of the evolutionary system – to be performed by a critic of some sort, be it human, rule-based, learning-based or even co-evolved. It could be argued that this interpretation while legitimate from an evolutionary computing perspective, limits the potential of computer-based composition. This is because it assumes that (i) critics know at every intermediate step that one approach will be musically 'better' than another, and (ii) that the creative process can be adequately expressed using some formalized system [7]. That is, it assumes that musical criticism is understanding rather than interpretation. It is suggested that the role of the non-human 'critic', requires considerable domain-specific knowledge, and that such knowledge can be difficult to quantify and encode [8], [9].

1.2 Evolutionary Systems

In natural selection, organisms adapt to their environment. It is possible to view a creature's morphology as the result of adaptation to its particular environment (physical, temporal and social). For the system described in this paper, the 'critic' in the terminology of Todd and Werner is the environment itself. The only domain-specific knowledge is a simplified physical system, which (superficially) mimics the dynamics of sound in the real world. The human 'composer' acts in the mode of meta-creator, designing *environments* and observing the results (visually and sonically) as creatures within the world adapt to that environment.

In the natural world, evolution is not only about survival, however. In sexual species, sexual selection plays an important role [10]. In particular, mating calls represent one of the earliest forms of communication [11] and have roots in the origins of music itself. The artificial evolution of mating calls has been studied by Werner and Dyer [12]. Many famous composers have drawn from birdsong and mating calls as compositional material. Janequin, Beethoven, Messiaen, and Rautavaara for example, have all made extensive use of mating calls in their compositions.

1.3 Related Work

Much of the system described in this paper draws its inspiration and methodology from John Holland's *Echo* [13], particularly in the use of rule-based methods for the internal decision-making system of creatures. Many others have used evolutionary systems as a basis for composition, but in the main for compositional simulation [9], [6], rather than as a new form of creative tool for the composer.

A more closely related system to the one described here would be that of Dahlstedt and Nordahl [14]. Their *Living Melodies* system uses a genetic programming framework to evolve an ecosystem of musical creatures that communicate using sound. *Living Melodies* differs from *Eden* in that it assumes all creatures have an innate 'listening pleasure' that encourages them to make noise to increase their survival prospects. *Eden* contains no such inducement, beyond the fact that some sonic communication strategies that creatures discover should offer a survival or mating advantage. Hence, only some instances of evolution in *Eden* result in the use of sonic communication, whereas in *Living Melodies*, *every* instance evolves sonic communication.

2 The *Eden* System

Like many Alife worlds [15], [16], [17], [14], the *Eden* world operates over a two-dimensional rectangular lattice of *cells* that develop globally at discrete time steps. Each cell can contain multiple elements, the principle types being *rock*, *bio-mass* (food) and *evolving creatures* (who are appetizing and carnivorous). Rocks and bio-mass do not undergo evolution, however the bio-mass operates under a feedback control model similar to that of Lovelock's *Daisyworld* [18]. The intensity of radiant energy falling on the bio-mass is seasonally adjusted, giving rise to cycles of growth and decay.

Rocks are placed in the world according to a simple diffusing model. If a rock is placed in a cell, then no further growth is possible in that cell. Rocks do not grow or change, but they do provide places of refuge allowing creatures to hide from predators. Each of the three major entities in the world has a distinctive 'color', allowing those entities with appropriate sensors to use color information to distinguish between the various types of matter in the world.

2.1 Performance System

The evolving creatures use a rule-based performance system, similar to that described by Holland [13]. The overall structure of this system is shown in Figure 1. Creatures have a series of environmental sensors that detect the physical qualities of the surrounding environment. This sensory information is passed as *messages* to a rule-based system that performs internal processing on a list of current active messages at each time step. As a result of this processing, the performance system may decide to output a message indicating that an action should be taken.

A creature's sensors and actions are drawn from a finite set of possibilities and do not evolve. Evolution of the performance system (how the sensory information is turned into actions) does undergo an evolutionary process.

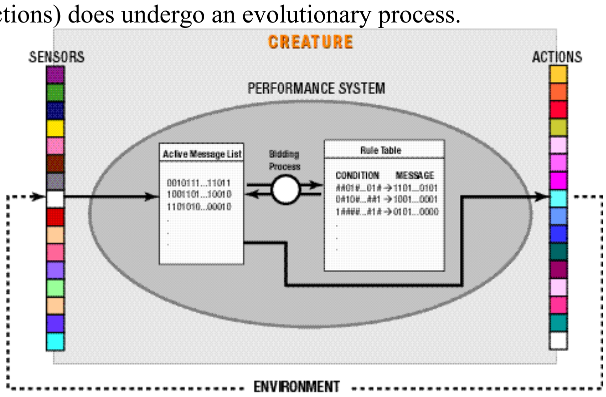

Fig. 1. The principle components of a creature: a series of sensors that detect environmental and local stimuli; a rule-based performance system that evolves; and a set of actions that the creature can perform in the world

2.2 Sensors

A creature may 'sense' the following information from the environment:

- *Color* of the contents of the cell the creature is currently occupying, and the color of the cells in the front, left and right directions of the creature;
- *Nutritional value* of other entities in the current cell (creatures are carnivorous and may kill and eat each other). Rocks have no nutritional value;
- *Sound* information – details of the sound arriving at the current cell;
- *Pain* level – an introspection as to damage the creature is suffering. Creatures will 'feel' pain if for example they are being hit by another creature or are very hungry;
- *Energy* level – an overall measure of how healthy the creature is.

The primary goal of the system is for the creatures to develop interesting sonic behaviour, hence a great deal of bandwidth in the sensors is devoted to sensing sound. Sound sensors detect sound pressure levels across a range of frequency bands, giving the creature the potential ability to distinguish many different types of sound. This ability to 'hear' is complemented by the range of sounds a creature can potentially make, which is detailed in section 2.3.

A physical model determines how the sound is propagated through the environment [19]. Sound arriving at a cell from multiple sources takes into account distance, frequency, and intermediate obstructions. No real attempt is made to simulate the psychoacoustic properties of sound beyond the exponential relationship between pressure levels and perception of intensity. The perceptual mechanism for loudness behaves in an exponential way, as it does for humans,

$$L = 20 \times \log_{10}\left(P/P_0\right) \quad \textbf{(1)}$$

Where L is the sound pressure level in decibels (dB), P_0 a reference pressure corresponding roughly to the threshold of hearing in humans [20, page 1055]. Psychoacoustic properties are perceptual, and additional properties would be difficult to integrate into the system at its current level of physical modeling.

2.3 Actions

A creature can potentially perform any of the following actions:

- *Move* forward one cell;
- *Turn* left or right;
- *Eat* whatever is occupying the current cell;
- *Hit* what ever is occupying the current cell;
- *Mate* with whatever is occupying the current cell;
- *Rest* for a time step (do nothing);
- *Sing* with particular frequency and volume characteristics.

The action of 'singing' means that the creature generates a sound with particular frequency and volume characteristics. The range of frequencies possible mirrors that which the creature can potentially hear.

The intent to perform an action does not necessarily mean it can actually be carried out. Action *messages* are sent to the environment, where they are tested for physical possibility. For example, it is not possible to move into a rock. Other actions may be capable of being performed, but of course may be of little use, or even detrimental to the creature's health (e.g. trying to eat when there is noting to eat on the cell).

The environment of *Eden* enforces a simple physical model on the world and its inhabitants. All actions carry an energy penalty; the amount of this penalty determined by the physical effort needed to perform the action. A creature dies if its energy level reaches zero. A creature's internal representation keeps track of its current physical properties such as mass, velocity, color and energy level. Many of these properties may change over its lifetime. For example, as a creature eats bio-mass, its mass increases and physical actions like moving around cost more.

2.4 Rules

Messages passed from the sensor system are processed by an internal rule-based, message passing system, similar to that described by Holland [13]. Since this system is described in detail by Holland, the description here will be minimal, highlighting the differences between Holland's system and that of *Eden*.

Messages are stored in an *active message list*, in order of their arrival into the system. Messages are 32 bit binary strings, in the case of messages that come from the environment the string represents sensory information. A database of *rules* (called a *rule-table*) is maintained as part of each creature's internal representation.

Each rule consists of two components, a *condition* string and an *output message* string. Condition strings are composed from an alphabet of three possible symbols: *0*, *1*, or *#*. To see if a rule should be applied, the condition string undergoes a bitwise check with a message string from the active message table. For each bit, a *0* or *1* in the condition string matches the same symbol in the message string in the same bit position. A # in the condition string matches either a *0* or a *1* in the same bit position. The successful match of a message from the active message table with a rule's condition string results in the rule's output message being placed in the active message table[1].

2.4.1 Credit Assignment

Rules incorporate a *credit assignment* system, whereby each rule is assigned a credit value, representing the 'usefulness' of the rule to the organism. New rules begin with a default credit value, and must *bid* to be used. When more than one rule matches a given input string, the rule with the highest bid wins. In the case of equal highest bids, the winning rule is chosen at random from amongst the highest bidders. Bidding is proportional to a rule's credit value and it's *specificity* (the less # symbols in the condition string the more specific it is – a condition string comprised of only # symbols will match every input string, but always bid 0).

If a rule is the successful bidder on a message from the environment (a sensor message), it pays its bid to the environment. If it is bidding on a message generated from another rule it pays the rule that generated the message. Only the winning rule must pay out its bid – the losing bidders do not loose any credit. Again, the reader is referred to [13] for details.

[1] Subject to successful bidding, detailed in the next section.

2.4.2 Credit Payoffs

Each creature's health and energy levels are monitored via a *health index*. If the creature is finding food and not being attacked for instance, the health index will increase. If the creature is running around aimlessly or being attacked, its health index will decrease.

Let $H_t(\lambda)$ represent the health index of a particular creature, λ at the current timestep *t,* and $H_{t_i}(\lambda)$ be the health index at some previous time t_i (where $t - t_i > 0$). For each creature, the cumulative differential of the health index is monitored and when it its magnitude exceeds some constant, α_λ, a *credit payoff* is performed on the active rules (i.e. when the inequality $\left|H_{t_i}(\lambda) - H_t(\lambda)\right| \geq \alpha_\lambda$ is satisfied).

For a given credit payoff C_i, all the rules that were successful bidders since the last credit payoff C_{i-1} are kept in a list. In addition, the total credit paid out to the environment since the last credit payoff, E_i, is kept. All active rules are paid out proportionally according to the formula,

$$P_i = \frac{kE_i}{f_{R_i}} C_i \tag{2}$$

Where f_R is the frequency of the particular rule *R* in the list, P_i is the credit value added to the rule's current credit value and *k* is a constant. C_i may be positive (meaning an increase in health) or negative (a decrease in health). Recall that C_i is a differential value, representing the rate of change in health. The number of time steps between successive payoffs will be dependent on how quickly or slowly the creature's health is changing. For example, if a creature is being attacked and losing health quickly, payoffs will be more frequent. The rules involved in letting the creature get hit will also decrease in credit quickly (hopefully soon being outbid by other rules that may prove more successful, if the creature is to survive).

Using this payoff system of equation (2), over time rules that are helpful to the creature's survival gain credit. Using the table-based approach of active rules allows rules that indirectly increase health to receive appropriate credit. For example, while the rule to 'eat when you find food' is a good one, you may need to walk around and look for food to find it first. The rules for walking and turning, although they decrease health in the short term, may result in finding food. This increases health in the longer term. If such rules are helpful in increasing health, their credit will increase.

2.5 Evolution

No artificial life system is complete without evolution[2] and *Eden* is no exception. The genetic algorithm used is based on the *Schemata* approach of Holland [21], [13]. Only the particular implementational aspects of *Eden's* evolutionary components relevant to this discussion are detailed here.

Recall from section 2.3 that mating is a possible action for a creature. The basic units of genetic exchange through mating are the creature's rule-tables. Mating will

---

[2] 'No artificial life systems of this ilk' may be a better generalization.

succeed only if a creature is over a certain age and is healthy enough. Successful mating, while costing energy, does not adversely affect health. For two creatures that mate, the most successful (highest credit) rules are crossed over in the hope that the resultant rules may also be successful. Additionally, rule mutation and creation are possible. The probability of these events occurring can be controlled interactively by the user at run time.

The observant reader will note that the selection of rules based on their strength during crossover represents a *Lamarckian* evolution [22], since learned behavior is passed from parents to offspring. This is done for efficiency reasons, as it results in the discovery of interesting strategies for survival more quickly than for a Darwinian approach. It is quite possible to bypass the Lamarckian components of the evolutionary system, if required.

3 Implementation

Although the primary goal of *Eden* is the evolution of sonic communication for the purposes of music composition, the program has a visual dimension as well. As explained in section 1.2, the composer has the role of 'meta-creator', and a visual feedback system was considered a reasonable way of facilitating this role. It is also an interesting way to observe, anecdotally at least, the behavior of creatures in the world.

Representation of the entities of *Eden* is done using tiling patterns, loosely based on Islamic ornamental patterns [23]. The resultant images formed by a grid of cells (as shown in figure 2), suggest a continuous mass of substance, as opposed to squares containing individual entities.

3.1 Interaction

While *Eden* is running, the user has control over a number of parameters. These include mutation and random rule generation probability during breeding, selection of coefficients that control bio-mass efficiency and distributions, and physical parameters related to the environment. Real-time statistical and environmental information is also provided. A sample screen shot is shown in Figure 2.

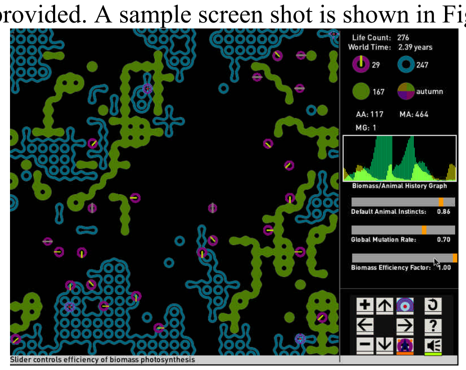

Fig. 2. A screen shot of *Eden* in operation showing the world (left) populated with rocks, bio-mass and evolving creatures. The user interface and statistical controls are on the right side of the image

3.2 Sound

Sound listening and sound generation is split into frequency bands. For simplicity, the current implementation uses three frequency bands. One could imagine the frequency bands as 'low', 'mid', and 'high'. Creatures can generate sound at four possible levels of intensity at each band, giving a total gamut of 64 distinct sounds. This selection is somewhat arbitrary, chosen as a reasonable tradeoff between storage space and sonic diversity. There is no reason why it could not be extended if required.

All 64 sounds are pre-generated, using frequencies within the (human) audio spectrum (20-20,000Hz). The base frequencies chosen to represent 'low', 'mid' and 'high' were 100, 1000 and 10,000Hz. Each sound is of a few milliseconds duration, roughly equal to a single time step. A separate software program pre-generates the sound set based around user-supplied frequencies. Experiments have been performed with different sound sets, and the generation and selection of sound sets could be considered part of the compositional process. It is important to note however, that the creatures do not 'hear' the timberal properties of the sound – they just register as bit patterns on the individual's sensors. The choice of sonic qualities of the sound set is for the benefit of the human listener, as a sonification of the process.

When the program is running, sounds are replayed in real time as the creatures sing them, thus creating the composition. A creature may sing different sounds each time step, thus permitting the generation of more sonically complex sounds that change in tonal characteristics over time. The ability to control levels at different spectra also contributes to the tonal evolution of sounds that the creatures make.

4 Results

As mentioned in section 1.3, there is no 'hardwired' impetus for a creature to actually make or listen to a sound. They will only do so if it increases their chances of survival or mating.

A number of simulations have been run in relatively small worlds (usually a grid size of 50 x 50 cells). Starting with random rules takes a long time to evolve to anything useful, so controls are provided to seed new creatures with some basic 'instinctual' rules to help them survive (e.g. '*if on top of food then eat*'). The user can choose dynamically how much or how little 'instinct' they would like the initial population of creatures to have. If the population falls below a given threshold, the world is automatically populated with a fresh set of creatures.

In many simulation runs sound does not play a major role in survival. While sound can be heard, analyses of the rules show that this is principally due to mutation (for example the 'eat' action mutates into a 'sing' action). Such rules do not survive for long, if they do not provide an increase in the health index. Singing costs energy, so it will only be used if it can provide some survival or mating advantage.

In some simulation runs however, sound does provide interesting survival advantages. For example, offspring of a certain parent used sound to signal an abundance of food, calling its siblings to share in the find. A creature's 'ears' listen in a forward facing direction, over a conical shaped area, so moving forward when hearing a food call is a reasonable strategy. A creature's (fixed) morphology dictates that it can hear at much greater distances than it can see. After many generations of using this method of signaling for food, new behaviors eventually emerged that exploited the food-signaling tendency, calling creatures over and then hitting, killing, and eating them. Why go searching for food when it will come when called…

Creatures also exploit the frequency dependent nature of their 'voices' and 'ears'; often groups (children of a common ancestor) will use only a particular frequency band when communicating. The use of different frequency bands for different tasks (food signals and mating) have also been observed.

More results that are interesting are obtained when some basic 'instinctual' rules for using sound are seeded into new creatures. For example, if the world contains creatures that cry out when being attacked, others soon learn to avoid going near those who are screaming. Behaviors also emerged whereby creatures would use signals to find mates, or to deceive in order to attack and eat another individual.

Creatures also evolve strategies for coping with seasonal change. For example, in the winter months, when food is scarce, some individuals prefer to 'hibernate' resting until food becomes more abundant and easier to find, then bulking up in the summer months in a period of frenetic activity.

Perhaps the most interesting properties that the system exhibits is the 'evolution' of sounds as the generations pass by. Simulations often begin with chaotic cacophony of sound that slowly becomes more 'pure' and sparse as the inhabitants of *Eden* evolve and discover rules that exploit the use of sonic communication. Certainly, the sonic development exhibits characteristics of dynamics and counterpoint, often considered the basis of many good compositions.

5 Conclusion

Clearly there are a number of limitations in the system described. In the physical world, organisms evolved sensory organs because they were useful for survival or mate selection. Whereas in *Eden*, the sensory organs are provided unconditionally and the internal structure of the creature evolves around them.

Secondly, the internal performance system could be improved. While the rule-based system described does permit complex sequences of actions to evolve, the discovery of such actions can be difficult (the longer the dependency of rules on each other, the more difficult they are to discover). It is planned to replace the rule-based system with a non-deterministic finite state automata (NFA) based system, in the hope that this may permit the evolution of a more complex use of communication.

Nonetheless, despite these limitations, as a compositional system, *Eden* is capable of producing compositions that have interesting qualities[3]. The goal of this work is not simulation or mimicry of existing compositional techniques or styles, but to expand the possibilities for composers who want to 'creatively look back on nature' to paraphrase Von Schlegel. Audio samples from some evolved worlds of Eden are available on-line for the listener to judge for themselves at:
http://www.csse.monash.edu.au/~jonmc/projects/eden.html.

5.1 Acknowledgements

Eden was developed during a fellowship provided by the Australia Council.

[3] Of course, 'interesting' is a subjective term and one always tends to find ones own creations interesting.

6 References

1. Schlegel, F.V.: *Lucinde and the Fragments*, University of Minnesota Press (1971)
2. Loy, G.: *Composing with Computers – a Survey of Some Compositional Formalisms and Music Programming Languages*, In *Current Directions in Computer Music Research*. M.V. Matthews and J.R. Pierce, (eds.). MIT Press: Cambridge, MA. (1989) 291-396
3. Xenakis, I.: *Formalized Music: Thought and Mathematics in Composition*, Indiana University Press (1971)
4. Johnson-Laird, P.N.: *Human and Machine Thinking*, Lawrence Eribaum Associates (1993)
5. Ebcioglu, K.: *An Expert System for Schenkerian Synthesis of Chorales in the Style of J.S. Bach*. In *1984 International Computer Music Conference*. San Francisco, ICMA, (1984) 135-142
6. Todd, P.M. and G.M. Werner: *Frankensteinian Methods for Evolutionary Music Composition*, In *Musical Networks: Parallel Distributed Perception and Performance*. N. Griffith and P.M. Todd, (eds.). MIT Press/Bradford Books: Cambridge, MA. (1998)
7. Solomonoff, R.J.: *The Discovery of Algorithmic Probability: A Guide for the Programming of True Creativity.*, In *Computational Learning Theory: Eurocolt '95*. P. Vatanyi, (ed.). Springer-Verlag: Berlin. (1995) 1-22
8. Wolpert, D. and W. Mcready: *No Free Lunch Theorms for Search,* Technical Report SFI-TR-95-02-010, Santa Fe Institute. (1995)
9. Wiggins, G., et al.: *Evolutionary Methods for Musical Composition*. in *Proceedings of the CASYS98 Workshop on Anticipation, Music & Cognition*. (1999)
10. Darwin, C.R.: *The Desent of Man, and Selection in Relation to Sex*. Reprinted 1981 by Princeton University Press ed. London, John Murray (1871)
11. Hauser, M.D.: *The Evolution of Communication*. Cambridge, MA, MIT Press (1997)
12. Werner, G.M. and M.G. Dyer: *Evolution of Communication in Artifical Systems*, In *Artificial Life II*. C.G. Langton, (ed.). Addison-Wesley (1991) 659-682
13. Holland, J.H.: *Hidden Order: How Adaption Builds Complexity*, Helix Books (1995)
14. Dahlstedt, P. and M.G. Nordahl: *Living Melodies: Coevolution of Sonic Communication*. in *First Iteration: a conference on generative systems in the electronic arts*. In: A. Dorin and J. McCormack (eds.): Melbourne, Australia, Centre for Electronic Media Art, (1999) 56-66
15. Ulam, S.: *Random Processes and Transformations*. In *Proceedings of the International Congress on Mathematics* (1952) 264-275
16. Conrad, M. and H.H. Pattee: *Evolution Experiments with an Artificial Ecosystem.* Journal of Theoretical Biology, **28** (1970) 393-401
17. Packard, N.H.: *Intrinsic Adaption in a Simple Model for Evolution*. In C.G. Langton (ed.) *Artificial Life, SFI Studies in the Sciences of Complexity*. Los Alamos, NM, Addison-Wesley, (1988) 141-155
18. Watson, A.J. and J.E. Lovelock: *Biological Homeostasis of the Global Environment: The Parable of Daisyworld.* Tellus, **35B** (1983) 284-289
19. Roederer, J.: *Introduction to the Physics and Psychophysics of Music*. Second ed. New York, Springer-Verlag (1975)
20. Roads, C.: *The Computer Music Tutorial*. Cambridge, Mass., MIT Press (1996)
21. Holland, J.H.: *Adaption in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. Second ed. MIT Press Cambridge, MA (1992)
22. Bowler, P.J.: *Lamarckism*, In *Keywords in Evolutionary Biology*. E.F. Keller and E.A. Lloyd, (eds.). Harvard University Press, Cambridge, MA. (1992) 188-193
23. Grünbaum, B. and G.C. Shephard: *Interlace Patterns in Islamic and Moorish Art*. In *The Visual Mind: Art and Mathematics*. M. Emmer, (ed.). MIT Press: Cambridge, MA. (1993).