CHAOS AND CREATIVITY

The dynamic systems approach to musical composition.

Peter Beyls
Expanded Media Studios
Brussels, Belgium

Extended Abstract

Traditional channels for introducing intelligence in computer music systems are firmly rooted in the knowledge-based approach; methods and computational strategies borrowed from the field of artificial intelligence. Expert systems for composition and pattern-directed inference systems for real-time man-machine improvisation are exemplary. In general, the aim is to introduce independent creative decision making through computer simulation of human creativity. Impressive statements have been produced along these lines, in music as well as rule-based computer graphics. Two observations have led to the consideration of a totally different methods. First, expert systems become problematic if situations occur that were not anticipated by the programmer and sooner or later, the programmer is faced with a complexity barrier. Second, appreciation of the pattern making potential of nature led to the study of concepts like self-organization. Complex dynamical systems are an alternative to the constructivist approach in composition, i.e. the critical assembly of architectures of time according to some explicit scenario. Complex dynamical systems, on the other hand, consist of many elements interacting according to very simple laws but giving rise to surprisingly complex overall behaviour. Composition becomes experimenting with attractors -- instead of creating a rule-base -- as well as designing tools that allow the topology of the composer to interact with the system's internal activity. The idea is to critically push the system out of equilibrium using tactile motor control as to explore the various degrees of freedom of a given system. The implicit behaviour is then mapped to the musical problem domain. Improvisation becomes navigation in a hypothetical world of which the composer is both inventor and explorer. Strange and intricate imagery, in both space and time, is found in physics, biochemistry, fluid dynamics, ecology and nonlinear mathematics. We have implemented and evaluated various models for spontaneous pattern formation, including one-dimensional cellular automata, direct computer simulation of chemical instabilities as witnessed in the BZ-reaction and a spatial model exploring equilibrium behaviour in a society of interacting agents moving in 2D space. The present paper outlines a connectionist-like model, a regular structure of agents engaged in local interaction, using forces of activation and inhibition between neighboring agents. Randomness/determinism and chance/necessity seem at the heart of creativity and happen to be central to the music of our time. We propose to view emergent properties from initial random configurations as a subtle alternative for both constraint-based, reductionist handling of randomness as well as rule-based composition by way of some generative grammar; complex dynamics as a creative, generative principle and a channel toward higher levels of man-machine interaction. This paper was prepared for the Second International Symposium on Electronic Art, Groningen, Holland, November 1990.

1. Introduction.

Ever since the early days of computer music, composers have aimed to introduce musical intelligence in a machine by trying to imitate certain aspects of human musical intelligence in a computer program. Pioneers (Hiller & Isaacson, 59 and others) built programs that simulated a given musical style by establishing musical rules borrowed from existing musical paradigms. Constraints were used to filter the output of a random number generator. This method, generate-and-test, proved to be highly inefficient. More control was needed and by the early seventies, crossfertilisation between the emerging discipline of AI resulted in the adoption of the rule-based method for musical composition. In addition, more sophisticated programming methods, i.e. object oriented programming, were introduced. Style imitation *remained at the heart of much work; consider the Flavors Band (Fry, 84) and, more recent, EMI (Cope, 89). However, tools for explaining musical decision making were now available by tracing the history of the computational process. More important, composers could learn about their own musical objectives from the circular process of rule specification and the appreciation of the rules' consequences. Expert systems have led to powerful statements in the arts in general, for instance, consider the knowledge-based drawings of (Cohen, 79) and intricate compositions guided by very high level musical abstractions (Barlowe, 81). In contrast to music, many visual artists continue to see computer media as tools, with minor interest toward process oriented production methods. Cognitive approaches remain even sporadic. In the field of music, expert system technology has proven to be useful for the creation of composers' assistants and intelligent sound editors, introducing very high levels of abstraction within the man-machine dialogue. In interactive composing, so called pattern directed inference systems allow for real-time composition following a scenario of rules while at the same time keeping channels open for perception and interpretation of outside influence. These programs are capable of expressing an individual musical character while simultaneously accommodating requests for attention of an external, human musician (Beyls, 88). Responsiveness and ease of adaptation to large swings in context determine flexibility. However, if we insist on real-time performance, a clear definition and concise description of the problem area is needed. In addition, only simple representation methods and efficient search techniques lead to successful applications. Rule- based systems represent aspects of the world in symbolic form. The idea is to reconstruct problem solving behaviour (as seen in human experts) in a program which reasons and searches through this symbolic space. Logical inference is at the heart of these programs.*

2. Complex dynamics.

In recent years, two observations have led to the consideration of alternative programming methods in AI. First, conventional expert systems remain helpless when faced with situations where knowledge is missing or incomplete. Moreover, expert systems are constructed "by hand", the expected performance of the system has to be formulated in explicit statements. However, at a certain level of complexity, it becomes very difficult to keep track of this performance as well as to debug such systems. In summary, traditional expert systems do not show graceful degradation when situations occur that were not anticipated by the system's designer and the programmer faces a complexity barrier. Second, observation of the pattern making potential of nature has led to the investigation of the constructive forces behind forms of natural organization. For instance, organized patterns are created spontaneously in biological workspaces. Consider the self-organizing behaviour in societies of termites; large artefacts are constructed without any plan saying which actions should be taken and when. Snow crystals are a combination of order and disorder; they grow according to the delicate interplay of microscopic and macroscopic forces. The geometry of nature was put in perspective (Mandelbrot, 77) by drawing attention to the fractal dimension of naturally grown shapes. These are said to be scale-invariant, which means that detail is everywhere and more detail is seen as we approach the shape. I suggest to view this as a metaphor for zooming in on complexity as such because our appreciation of artefacts seems strongly influenced by dynamic navigation between detection of detail and simultaneous perception of overall appearance. In addition, the dynamics of zooming in and out within the process of creation is characteristic of an exploratory attitude; detail may lead to better understanding of the problem at hand. Incidentally, the most popular images of chaos are mathematical formulae expressing fractals as a static generator of visual design, in sharp contrast to seeing dynamical systems as a metaphor for exploration and discovery. There are many other fields where natural dynamics are observed such as biochemistry, ecology, biology, fluid dynamics, neuroscience and nonlinear mathematics. Any system consisting of many properties evolving in parallel over time may be considered complex if it exhibits emergent properties. Emergent properties are sudden, spontaneous structural changes in a system out of equilibrium and in constant interaction with its environment. Such patterns, in time and/or space, are a product of self-organization, (Prigogine, 84) refers to them as dissipative structures.

3. Behaviour.

So, the idea is to avoid the problems mentioned above and to borrow generative principles *from examples found in nature. Besides, as an artist I am more interested in models of evolution and change than in theories of structural design. What can we learn from the creative forces shaping natural phenomena within the realm of our objective: the introduction of aspect of human creativity in artificial computer music systems? Current approaches to this problem are strongly polarized, some express faith in symbolic computing, adhering to knowledge based strategies of problem solving (Laske, 89). Others claim that only a behavioural approach using methods of subsymbolic computing may lead to successful results when modelling aspects of human musical cognition (Leman, 89). We will not continue the debate but provide evidence of the strength and weakness of both approaches from pragmatic experimentation in the problem domain of real-time improvisation and interactive composing. The differences are briefly summarized in simple terms as follows. Symbolic computing is based on the exploitation of knowledge stated explicitly as facts and rules acting upon symbols under the guidance of a supervising mechanism. Subsymbolic methods use analogical representations, which keep what they represent implicit in their representation. Examples are regular arrays as seen in cellular automata and matrixes expressing weights in connectionist networks. Such distributed representations are attractive because of their direct, visual appeal: what you see is what you get. The activity in these systems is no longer guided by a supervisor but issues from the local interaction of many participating agents. The distinction between knowledge based and behaviour based intelligence is important. For instance, when designing a building, an architect definitely draws on knowledge of materials, construction practice, financial considerations etc. He knows how to tackle very specific problems by reasoning and making choices. Does he use similar knowledge to express himself in words; when he speaks about his problem? Definitely not. Speech, perception, locomotion are all examples where behaviour is at the heart of the activity. Many great jazz musicians are exemplary here; the musical intensity of a virtuoso keyboard improvisation has more to do with spontaneous motor control than, say, declarative knowledge of musical scales. Some cultures emancipate the behavioural idea to its fullest. According to Indian musical practice you have to study raga's for 20 years, yet when going on stage, you forget everything you know and just let it happen. In the Indian language it's called "uppaj", which means imagination or "flying like a bird". You let go of all knowledge you have accumulated and take off. Cognitive activity during real-time interactive composition includes perception, imagination and reaction. Behavioural strategies seem appropriate here since they can establish a direct relationship between perception and action. It takes too much time to interpret auditory stimuli using search and mapping over symbolic representations. In addition, the flux of relationships between performer and computer program may be totally unpredictable.*

4. Creativity.

So far, we have traced the AI-context of our basic problem: the construction of intelligent systems for musical composition while insisting that our system should exhibit aspects of creativity. A creative statement should be new and useful. Something new raises questions in the perceiver, fundamental questions, therefore, are products of advanced creativity. If we expect new ideas, would a program functioning as the logical consequence of a set of rules be considered creative? For instance, Chomsky's theoretical work on generative grammars was inspirational for computer based musical composition (Roads, 84) and computer-aided visualization (Smith, 84). Grammars are devices for advanced productivity, but human creativity is needed to design them. One may add meta-level reasoning about the rules in a wish to change and adapt them according to the circumstances they generate. Again, methods of circular thinking are characteristic in human creative behaviour. Ultimate creativity seems connected to the discovery of new paradigms, new ideas that go beyond the potential embedded in rules. Creativity may be seen as searching through a very large problem space (Steels, 86). However, focus in a creative process is mobile. Perhaps a solution is found for a problem that was not anticipated by the programmer -- and the original problem is forgotten altogether. In other words, creativity works as an unpredictable, non-linear process. Incidentally, non-linearity also happens to be a native characteristic of complex dynamical systems. In a wider context, true creativity has more to do with self-revision than with self-confirmation, i.e. the application of procedures, ideas etc. that have proven to be useful in the past. A way to introduce unpredictability is the use of random numbers. Randomness is often used to simulate musical intuition. However, total randomness, like absolute repetition or total predictability, does not carry much meaning. Incidentally, the paradox is that while some artists express faith *that chance procedures may help to imitate human creative playfulness with originating ideas, others view it in opposite terms. Indeed, surealist artists were confident that random techniques would assure exclusion of personal involvement and intuition in the creative process. In interactive composing, the illusion of musical intelligence is a byproduct of minimal decision making, randomness injects energy, it tries to activate all available levels of activity in a given software defined process. Interactive composing offers schemes where motor activity of a human performer adjusts levels and degrees of freedom in such a process. It is important to note that interaction happens in real-time; a composition emerges from intimate man-machine interaction. The performer/composer provides feedback to a generative process of his own design and the emergenge of musical shapes in this abstract, conversational process may be taken literally in the light of the emergent properties seen in complex dynamical systems. Random methods involve selection by imposing constraints on the output of random generators. Grammars, as mentioned above, are examples of constraints imposed at the generative level. In interactive composing we should be able to provide feedback to such programs. Predetermined rules and constraints would be optimized according to continuous evaluation at both sides of the screen, i.e. the program would learn which suggestions are more successful than others (1) and the performer would learn about the current direction of his musical objectives (2). Such natural awareness of unspoken relationships is still unique to human ensemble improvisation. Here, appreciation of human musical intelligence includes awareness of a collective physicallity where musical intensity radiates from unspoken, deep rules/constraints imposed by culture as well as the topology of the human body. By the way, consideration of physical parameters is extremely well developed in ultra low-tech environments like those found in ethnic music, overdeveloped in many a popular musical idiom, while very often underdeveloped in the avant-garde.*

5. Heuristics.

The task of simulating human musical creativity has been recognized as a very difficult -- if not impossible -- one. In the context of advanced knowledge-based programming methods, composition may be seen as a problem solving process. The idea is to find a solution for a given musical problem by exploring a very large search space. Since the search space is by far too large to be explored by exhaustive search, specialized short cuts, known as heuristics, are needed. Heuristics are formalized in rules saying what to do given certain circumstances, we may think of rules as surface knowledge. However, the deeper knowledge are the constraints which are based on the physical properties of musical material or on a particular esthetic theory (Steels, 86). Composing means scanning a search space looking for a musical structure which satisfies the constraints. This involves the creation of a schedule (Ames, 1983); priotizations of all available options from most to least desirable. The program then evaluates this schedule to find an option that satisfies all constraints. However, in case no single option proves acceptable, the program resorts to backtracking, trying to revise previous decisions in the decision tree and, then, trying to obtain a valid solution from there. Expert composers exhibit a highly developed, natural sense for applying appropriate heuristics facing given constraints. It is exactly this expert knowledge which is extremely difficult to capture because it is active on the subconscious level. This leads many to believe that artificial approaches to creativity are doomed to fail since expert knowledge, referred to as musical consciousness or intuition are missing while these are prerequisites for true imagination. However, intensive work in the field of machine learning reveals the potential of having programs learn their own heuristics, leading to second generation expert systems (Steels and Van de Velde, 85). The idea is to build up a heuristic knowledge base automatically through introspection of the program's own problem solving behaviour. This seems absolutely necessary since heuristics are not consciously available. In addition, we know that even if they were, they would be too complex for handcoding, one of the very reasons that -- apart from the introduction of learning -- led to the consideration of a completely different programming paradigm suggested by interesting behaviour in complex dynamical systems. Indeed, in contrast to the approach expressed above, explorers of non- linear systems believe that the laws of unpredictability, chaos and irreversible time are a key to natural creativity -- the ultimate example might be the emergence of life through natural selection, a process where the non-equilibrium dynamics of the environment leads to self-organized structure. Prigogine even suggests the concept of an entropy barrier, meaning that any system of sufficient complexity will become unstable losing its initial conditions and behave unpredictably. With such systems, time is irreversible, it is *impossible to backtrack previous states, there are too many. This is easily and convincingly demonstrated, even with simple cellular automata rules like those found in Conway's Game Of Life. Life is not backward deterministic, a given pattern often has many patterns that may have preceded it; a pattern has only one future but many possible pasts. The recognition of probabilistic evolution and chaotic propagation as a driving force in creative processes is in sharp contrast to the reductionistic view of creativity.*

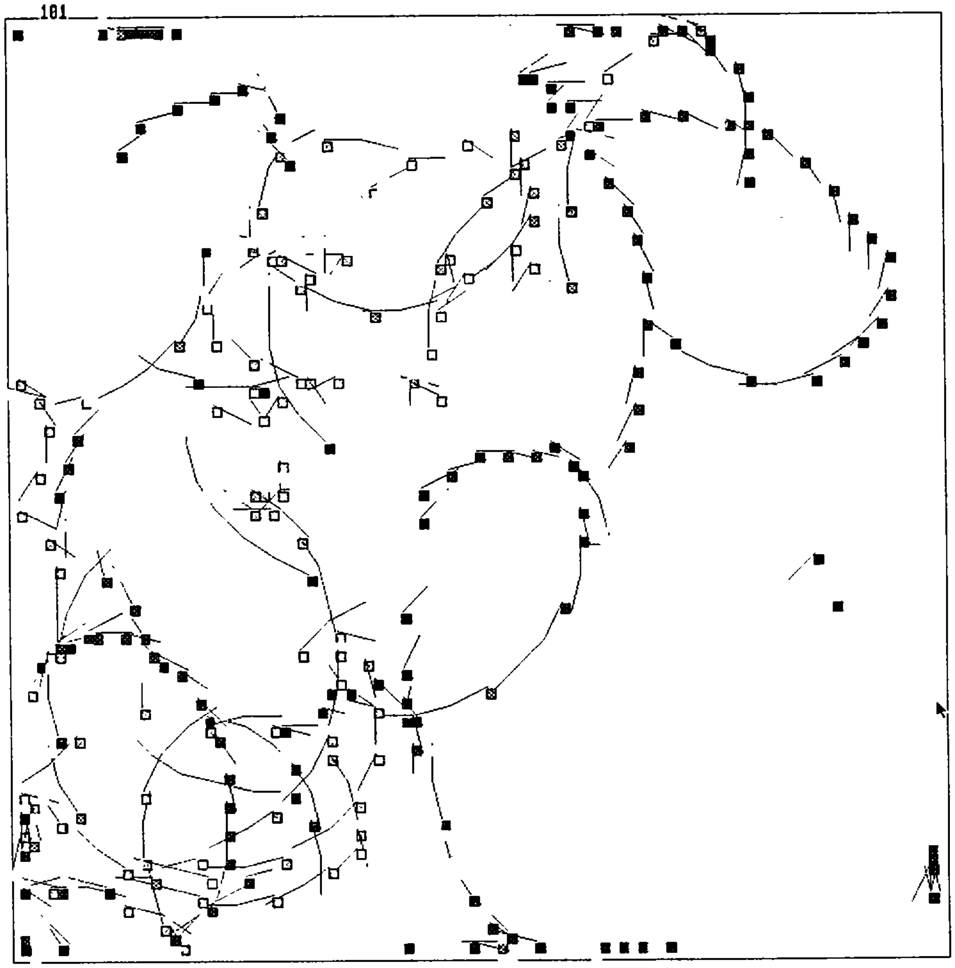

Figure 1: tracing activity in a spatial model.

This figure shows the state of affairs after about 100 generations of 8 agents moving in 2D space. All agents express individual affinities for integration and expression towards all other agents in this virtual society. In addition, agents interact (execute simple rules) when navigating in eachother's neighborhood. As a result, complex macroscopic behaviour results from simple microscopic activity.

6. Randomness & chance.

In a complex dynamical system randomness may be used to explore all degrees of freedom potentially available. However, principles other than constraint-based critical filtering are active for stimulating the birth of automatic musical structures. These principles include: self-organization, activation-inhibition, spreading activation and relaxation. Strong arguments for using dynamical systems is their inherent flexibility and speed of reaction (when compared to rule based systems), they are adaptive by nature, they use distributed ways of knowledge representation (more robust than rule-based systems), in addition to allowing the conception and control over musical structures in terms of underlying images of these structures. These images, known as analogical representations, let us grasp the effect of generative principles in a single picture and, indeed, one picture is worth a thousand words. In addition, gestural techniques, common practice in computer graphics (e.g. pointing, dragging), may directly act upon these images. In fact, any sensory mechanism may be mapped to such representations. For example, in a recent piece for violin with computer extensions (Beyls, 89), sequential melodic material from the instrument is accumulated in a transition network. The content of this network activates a 2 dimensional cellular automaton. Simple local rules push global overall structures to surface. These are then interpreted by a rule-based algorithm that arranges for automatic orchestration of the original monophonic lines into full 8-voice polyphony. In terms of nonlinear jargon, we say the systems exhibits attractors because points of relative stability fluctuate through time and space. The output of such systems ranges from quasi-periodic oscillations to the building up of strange attractors because of internal feedback.

We have studied many methods for acquiring spontaneous activity in environments for interactive composing including one-dimensional cellular automata (Beyls, 89), direct computer simulation of nonlinearity as observed in biochemical processes (Beyls, 90a) as well as a spatial model (Beyls, 90b). The latter case is an example of a combination of constraint-satisfaction and local operations: abstract musical entities, uniquely defined "actors", express individual opinions about each other and engage in interaction when meeting in 2-D space. This micro-world, as a whole, accommodates external gestures from the composer. As a result of both internal spontaneous activity and external goal-directed activity, strange patterns of variable coherence emerge. Figures 1 and 2 provide a snap shot of 8 interacting agents taken after about 100 generations. As a final example we will briefly describe yet another model, inspired by research in fluid dynamics. It is an example of activation/inhibition.

7. A connectionist model.

The initial idea for the current model came from the appreciation of strange phenomena in fluid dynamics such as Benard convection (Babloyantz, 86) and the dripping faucet experiment (Wolfram, 84). Such systems are inspiring because they allow gradual navigation in degrees of complexity by controlling a single external parameter. The suggested model is a connectionist-like structure. You may see it as a specialized cellular automaton, a micro-world represented as a regular array of cells, each cell representing a 'virtual musician'. Cells are active units called agents according to the terminology suggested in (Minsky, 86). These are very simple in themselves but, overall, collective complexity results from mutual interaction combined with external influences. We proceed as follows:

1. *design a simple agent, i.e. definition of a native character*
2. *create initial random affinities between agents*
3. *specify global constraints acting as global parameters*
4. *provide gestural input to particular (groups of) agents*

The idea is to map the resulting behaviour to our problem at hand; the creation of musical structures from emergent properties in a complex dynamical system. The system is seen as a collective of agents exhibiting evolving connections amongst each other. Any two agents connect of they exhibit sufficient affinity towards each other. The agent's other attributes include individual levels of activation and inhibition, energy and position in space. The principle responsible for pattern formation in such organized networks is known as "spreading activation". Agents are thus linked in a network, the links being of variable strength because every agent features an activation level and an inhibition level. The agent's flexibility to move is a function of its level of activation, which in turn is a non-linear function of external constraints. In addition, the system as a whole dissipates energy, and the value of the dissipation factor will introduce variable inertia. Figure 2 illustrates the dynamics of 36 agents organized in a regular 6 by 6 array, time runs left to right, top to bottom. For twenty-six generations, activation, inhibition and local gradient (tension) are shown. Notice peculiar, oscillatory behaviour and how the system moves to a different limit cycle from the accommodation of external disturbances at generation 9 and 19. The full potential of this connectionist model is documented in a separate paper (Beyls, 90c).

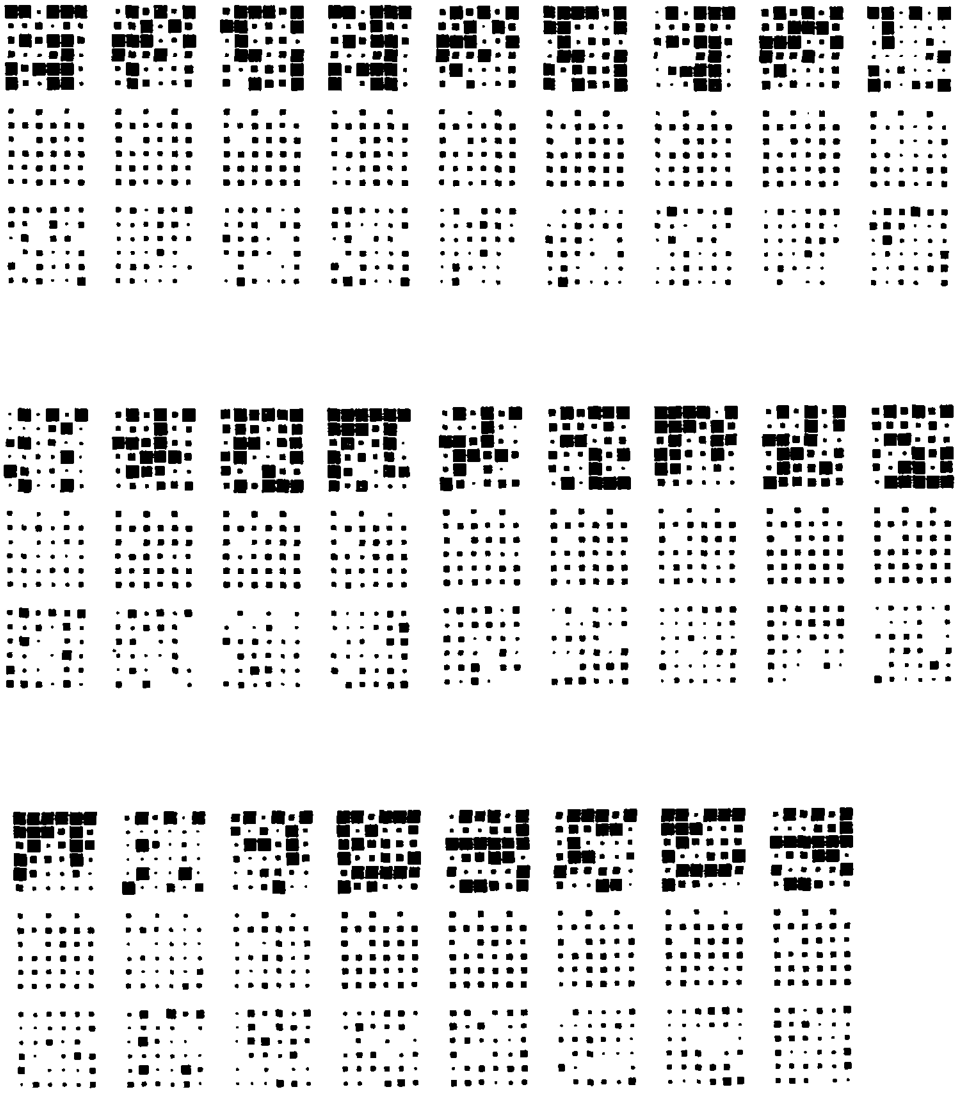

Figure 2: activation/inhibition in regular structure.

Per generation, levels of activation, inhibition and gradient field are visualized.

8. Conclusion.

Does it make sense to view the disorderly behaviour of complex systems as manifestations of true creativity? For the mathematician, the concise formulation of some complex, physical phenomenon in a neat, compact, simple formula may be experienced as an aesthetic artefact. The overwhelming visual complexity of fractal pictures is intriguing. However, our focus *should not be with visual appeal but with what these pictures represent. In addition, the unexpected behaviour of chaotic systems -- the delicate interplay of chance and determinism -- is observed after implementing a computer simulation; that is, possibilities are discovered in retrospect. Self-organization may certainly be seen as a powerful alternative to the constructivist approach in musical composition. We no longer specify recipes for the critical assembly of musical atoms, e.g. the hierarchical structuring of notes, phrases, etc. Here, composition is seen as the architecture of time. In contrast, we aim for spontaneous pattern formation from experimentation with attractor systems. This experimental attitude is strongly related to the invention of hypothetical worlds of which the composer is both inventor and explorer. Composition becomes navigation in attractor fields, the interactive, conversational exploration of levels of stability and sensitivity. In essence, we observe a continuous confrontation of two dynamical systems; one embedded in a computer program, the other as present in the composer's opinion and reaction to its current output. Since the interaction relies heavily on real-time visualization, it is a particular example of the integration of computer music and computer graphics.*

9. Acknowledgement.

I am indepted to the complex dynamics group at the AI lab of Brussels University for instructive feedback. Work described in this paper is supported by Yamaha and Atari Corporation.

10. References.

Ames, C. 1983, Stylistic automata in Gradient, Computer Music Journal, Vol. 7 nr. 4, MIT Press

Babloyantz, A. 1986 Molecules, Dynamics and Life. An introduction to the self-organization of matter. John Wiley and Sons, NY

Barlow, C. 1988 Journey to Parametron, Feedback Editions, Cologne, W. Germany

Beyls, P. 1988 Introducing OSCAR, Proceedings of the International Computer Music Conference, Feedback Editions, Cologne

Beyls, P. 1990 Musical morphologies from self-organizing systems, Interface, Vol. 19, Swets and Zeitlinger BV, Holland

Beyls, P. 1989 The musical universe of cellular automata, Proceedings of the ICMC, Columbus, OH.

Beyls, P. 1990 Subsymbolic approaches to musical composition. A behavioural model. Proceedings of the ICMC, Glasgow

Chadabe, J. 1977 Some reflections on the landscape in which computer music systems are designed. Computer Music Journal Vol. 1-3

Cohen, H. 1979 What is an image? Proceedings of the IJCAI Conference, Tokyo, Japan

Cope, D. 1989 Experiments in musical intelligence, non-linear, linguistic-based composition. Interface Vol. 18, nr. 1-2

Fry, C. 1984 The Flavors Band: A language for specifying musical style Computer Music Journal Vol. 8 nr. 4, MIT Press

Hiller, L. & Isaacson 1959 Experimental Music, McGraw-Hill, NY

Laske, O. 1989 Articulating musical intuitions by way of rules. Proceedings of the Inernational Workshop on AI and music, IJCAI 89, Detroit

Leman, M. 1989 Symbolic and subsymbolic information processing in models of musical communication and cognition. Interface, Vol. 18, nr. 1-2

Mandelbrot, B. 1977 The fractal geometry of nature, Freeman, San Francisco

Minsky, M. 1986 The Society of Mind, Simon and Schuster, NY

Prigogine, I. & Stengers, I. 1984 Order out of chaos, Bantam Books

Roads, C. 1985 Grammars as representations for music. Foundations of computer music, MIT Press 1985

Smith, A.R. 1984 Plants, fractals and formal languages, Computer Graphics Proceedings, *Siggraph Conference, vol. 18-3*

Steels, L. 1984 Artificial intelligence and complex dynamics, AI Memo, VUB, Brussels

Steels, L. 1986 Learning the craft of musical composition, Proceedings of the ICMC, The Hague, Holland

Steels L. and Van de Velde, W. 1986, Second generation expert systems. Journal on future generation computer systems. Vol. 1, nr. 4 North Holland Publishing.

Wolfram, S. 1984 Computer software in science and mathematics, Scientific American, March 1984