Aural Potential of Emergence

Pablo Castro
Bachelor's Thesis
Institute of Sonology

May 2014

Abstract

The intention of this thesis is to guide the reader along the path I followed in order to design a generative process for the production of sound structures. Some concepts from complexity theory and emergence will be discussed in order to create a framework to better understand the context in which the ideas leading to the generative process developed. After relevant aspects of the theory have been covered some of the ideas behind the design of a generative process will be discussed in more detail.

Acknowledgements

I would like to express my gratitude to Paul Berg for his patience and encouraging attitude during my time at The Institute of Sonology and to Kees Tazelaar for his dedication to running and maintaining such an inspiring and special environment.

Contents

Introduction

Patterns and processes giving rise to them have captivated my interest since I attended art academy, before enrolling in the Sonology Bachelor. At the time I was impressed by how most patterns that are the result of natural processes make very effective, if not perfect, use of the design principles (balance, gradation, repetition, contrast, harmony, dominance and unity) and design elements (line, shape, direction, size, texture, hue and tone). For example, the intricate and balanced lines in the structure of a snowflake, very organized and simple yet so graceful and expressive. As I became more interested in the topic I began noticing the seeming simplicity of the processes by which most of those patterns are formed. Their dynamic nature and adaptability were very attractive features which I thought had some potential for artistic and creative exploitation.

After starting the Sonology Bachelor my interest in the topic grew but my focus shifted towards the aural domain. During the second year of my studies I started relating some of the concepts and ideas I had previously learned to sound. I embarked on the research for this thesis approaching the end of the third year of the program. By then, some sense about what type of process could be designed in order to materialize those aspects of the theory that I considered to be relevant became more clear. The question was not just about how to design a process that has audible emergent characteristics or a generative system. What was really crucial for me was whether or not emergence is something I can hear and appreciate as much as I do when I look at it in its different manifestations in the world around me.

It is the intention of this thesis to guide the reader through the path I followed in order to design a generative system for the production of sound, based on some concepts and ideas regarding the theory of emergence from different fields. Although a comprehensive summary of these theories is beyond the scope of this thesis, a few relevant aspects will be discussed in chapter 1.

After having some understanding of the mechanisms and rules that drive natural processes which display emergent qualities, a link to the aural domain was relevant but still missing. Some level of abstraction was required in order to apply those concepts when designing a system that would produce sound material which exhibits such qualities. A description of the transition from the theory and visual perception of emergent phenomena to the aural perception of it will be described in chapter 2.

Chapter 3 will discuss more specific, non-technical, aspects of the generative system I have designed. Expectations before the process started to take form will be reviewed as well as how some of the obstacles I came across actually paved the way for new abstractions which in turn produced a richer output of the system.

Chapter 4 will expose my views and insights on how this research and the product of it have shaped my views about using systems, processes and rules for the generation of sound.

Chapter 1

Emergent Properties

There are probably more theories explaining emergence or emergent behavior than there are fields researching its potential. A quick search in Google Scholar for ‘emergent behavior’ since 2013 until yields ‘About 18,000 results’. There is, up to now, no agreement on a single theory that clearly defines emergence (Corning, 2002, p. 6). An all-embracing definition of the concept remains elusive. The absence of an unified theory doesn’t devalue the novel and innovative ways in which complex systems that display emergent behavior can solve problems. How efficient they are at adapting to changes in both their inner and outer environments or how patterns and structures arise out of interaction between the components of the system. It is the process that matters. Its dynamic nature and the seeming coordinated and thoughtful way in which states of order are achieved and sustained over time. Understanding the characteristics that these systems have in common aids in comprehending emergent phenomena and highlights aspects of it which could be useful when designing emergence.

1.1 A Brief Introduction

Emergent behavior is exhibited by complex systems. I will limit the definition and discussion of complex systems to those aspects which I think are relevant in order to understand and apply the notion of emergence within a sonic context. Understanding the properties that different complex systems have in common is crucial if one wants to determine whether or not the higher-level patterns or behavior exhibited by a system are truly emergent.

Mitchell (2009, p. 12) proposes the following properties as being shared by most complex systems:

- Complex collective behavior
- Signaling and information processing

- Adaptation

She then goes on to define a complex system as follows: ‘a system that exhibits nontrivial emergent and self-organizing behavior’ (Mitchell, 2009, p. 13). It is worth noticing the distinction that is made between nontrivial and trivial emergence. He explains how ‘sometimes a differentiation is made between *complex adaptive systems*, in which adaptations plays a big role, and nonadaptive complex systems, such as a hurricane or a turbulent rushing river’. Trivial emergence could be understood by observing those patterns or behaviors exhibited by a system which do not play a significant role in its development or existence. That high-level organization arises as a result of the interaction among the components of the system. It can be seen as manifestations of the constant internal reconfiguration of the system. Think of the chladni patterns or cymatics. Patterns are formed over time but they have no other function.

Such patterns demonstrate the system’s tendency to self-organize whenever the conditions in the environment have changed, and serve as a an example of how effective the components of these systems are at following the rules which dictate their behavior. In other words, ‘Emergent complexity without adaptation is like the intricate crystals formed by a snowflake: it’s a beautiful pattern but it has no function’ (S. Johnson, 2001, p. 20). Non-trivial emergence does play a significant role in the system’s development, stability and existence. It is the ability of each component to adapt to changes in its immediate inner environment, which in turn enables the whole system to adapt to changes in the outer environment. That in which the phenomena is perceived by the “subject of cognition” (Di Scipio, 1994, p. 206).

What is remarkable about emergent behavior is not the mere fact that a system can adapt to changes in the environment where it exists. Is the fact that there is no single special component in the system which is in charge of keeping track of changes or alterations in the environment and passing on that information to another component. All the components are constantly assessing each others state. That leads to all of the components constantly exchanging information relevant to their functioning as a group. As coordinated as it may seem, the overall state of the system depends on decentralized decisions which are based on the execution of very simple rules which are inherited by each component. This implies that every component of the system plays a significant role in the system but is not essential to its existence. Each component is interchangeable. Regardless of how old any of them may be, the overall behavior of the system is maintained and its ability to adapt and come up with solutions suited to its environment tends to grow ‘smarter over time’ (S. Johnson, 2001, p. 20).

The study of emergence has its roots in the field of biology, more specifically in that of biological evolution. ‘The term “emergent" was first used by G.H. Lewes (1875) in his “Problems of Life and Mind", but the concept was adumbrated by the “heteropathic laws of causation" of John Stuart Mill’s “Logic" (1843)’ (Ablowitz, 1939, p. 139). For a history of the concept see Goldstein

(1999), Corning (2002) and (Ablowitz, 1939). Before complexity became a field of study in its own right, different attempts at explaining emergent phenomena had been pursued by experts in different fields.

> People had been thinking about emergent behavior in all its diverse guises for centuries, if not millennia, but all that thinking had consistently been ignored as a unified body of work —because there was nothing unified about its body. There were isolated cells pursuing the mysteries of emergence, but no aggregation (S. Johnson, 2001, p. 17).

A difference tends to be made between weak emergence and strong emergence. See (Bedau, 1997) for an interesting discussion on the topic. Although interesting ideas can be drawn from these perspectives, the system that will be described in chapter 3 is more related to the notion of weak emergence.

> Weak emergence applies in contexts in which there is a system, call it S, composed out of “micro-level" parts; the number and identity of these parts might change over time. S has various “macro-level" states (macrostates) and various “micro-level" states (microstates). S’s microstates are the intrinsic states of its parts, and its macrostates are structural properties constituted wholly out of its microstates. Interesting macrostates typically average over microstates and so compresses microstate information. Further, there is a microdynamic, call it D, which governs the time evolution of S’s microstates. Usually the microstate of a given part of the system at a given time is a result of the microstates of “nearby" parts of the system at preceding times; in this sense, D is “local". Bedau (1997, p. 4)

The way I grasped the concept of emergence was through the perception of it rather than by trying to digest the multiple philosophical discussions on the topic. It is something to be understood in terms of behavior, relationships between the components of the system, the environment and the subject of cognition. It is strongly determined by the amount of time one is exposed to the phenomenon and requires some degree of engagement with it. A good classic example describing emergent phenomena rather than the theory behind it is:

> If I play two notes together on the piano, there is an aspect or quality of this sound which is not the property of either of the notes taken separately. The chord has the characteristic of “chordiness"; the harmonious combination of sounds has a new attribute which no one of its individual components had, but which is due solely to their togetherness (Ablowitz, 1939, p. 2).

Technological advancements in the last century, along with the introduction of the computer as an essential tool for scientific research, have expanded not just our understanding of complex systems and emergence. It has allowed us to simulate them and even introduce some of the concepts to most software applications that have become essential to our daily lives. As is the case with understanding complex systems, identifying the common properties of emergent behavior displayed by different systems is more meaningful than trying to define what emergence is in one universal statement. Goldstein (1999, p. 49) suggests the following common properties:

- Radical novelty: emergents have features that are not previously observed in the complex system under observation.
- Coherence or correlation: emergents appear as integrated wholes that tend to maintain some sense of identity over time. This coherence spans and correlates the separate lower-level components into a higher-level unity.
- Global or macro level: since coherence represents a correlation that spans separate components, the locus of emergent phenomena occurs at a global or macro level, in contrast to the micro-level locus of their components.
- Dynamical: emergent phenomena are not pre-given wholes but arise as a complex system evolves over time.
- Ostensive: emergents are recognized by showing themselves, i.e., they are ostensively recognized.

Goldstein (1999, p. 55) adds that 'emergence requires systems with at least the following characteristics', which have become the 'backbone of complexity theory' :

1. Nonlinearity
2. Self-organization
3. Beyond equilibrium (multi-, non-, or far from equilibrium)
4. Attractors

1.2 Emergence or randomness?

Although it may be difficult to understand and sometimes predict the outcome of complex systems, there is a driving force that these systems have in common. What is perceived at the highest level as sophisticated and organized, either as a pattern or as a particular type of behavior, is produced by a series of underlying processes carried out by simple components who have no overview of the macro-behavior produced by their actions. 'The movement from low-level rules to higher-level sophistication is what we call emergence' (S. Johnson, 2001, p. 18). The idea may be puzzling at first. 'It is one of the most fundamental mysteries of emergence, which is how complicated organisms, with a wide variety of building blocks, can develop out of such simple beginnings' (S. Johnson, 2001, p. 84). Such a statement seems counterintuitive. How can something that by definition is referred to as complex be the result of layered interactions between simple elements following simple rules? We are not used to thinking about systems in terms of processes or collective actions. Our conception of a system, a misconception I would argue, is defined by the existence of an authority figure of some sort and a set of, far from simple, rules that need to be followed in order for the system to "work". Nature provides a majestic example that goes against the centralized view of systems, colonies of social insects such as termites and ants. Mitchell (2009) provides a quite accurate and short description of the way ant colonies work and how we perceive their ways:

> An ant colony, for instance, can consist of hundreds to millions of individual ants, each one a rather simple creature that obeys its genetic imperatives to seek out food, respond in simple ways to chemical signals of other ants in its colony, fight intruders, and so forth. However, as any casual observer of the outdoors can attest, the ants in a colony, each performing its own relatively simple actions, work together to build astoundingly complex structures that are clearly of great importance for the survival of the colony as a whole (Mitchell, 2009, p. 4).

What is interesting is that if we were to witness the behavior of a single ant it would appear as random, even purposeless. It is the interaction among many components that separates their random individual behavior from the self-organizing tendencies that characterizes these systems.

1.3 Perception

Unpredictability is central to the understanding of emergence. According to C. W. Johnson (2006) how we might interpret the resulting patterns or behavior is highly subjective to our knowledge of the inner workings and previous experiences with a given system. When first exposed to a system whose behavior is hard to predict, it may be looked on and perceived as random. If a pattern is perceived within a reasonable time window one may be tempted to label it as an emergent quality of the system. However, after gaining more experience with and knowledge about the system, the real nature of the behavior exhibited tends to become clear. What seemed random at first becomes predictable, sometimes even monotonous. When having enough information about the characteristics of each individual component, it's rules of interaction with other components and with it's environment it is then possible to make a more accurate assessment of the system. If the overall behavior or pattern still can't be predicted or anticipated after such assessment, are we being exposed to emergent phenomena.

It is common to define the unexpected behavior of a particular system as an emergent property of that system. The perception of emergent properties is directly linked to the amount of knowledge and understanding that one may have of a particular system. 'The term emergence is often used to reflect the limitations on our understanding of complex systems' (C. W. Johnson, 2006, p. 5). He suggests that the designer of the system, or those directly involved in its development, tend to be the ones with better knowledge of the possible outputs the system may yield. He describes different ways by which a complex system could be analyzed or investigated in order to determine if the system's properties are indeed emergent instead of random. I will include the ones that were relevant to my research:

- Observing component parts: Provides enough information about each component but doesn't say much about its behavior in an environment interacting with other similar components.
- Resulting behavior or pattern: Gives an impression of what the system can do, but not much can be learned about how the system does it.
- Interactions: If enough is known about the micro dynamics of the system, the interaction occurring at the low-level of the system, we get valuable information about the behavior of the components as a group, and in the context of the environment where they exist.

Another important aspect of how we perceive emergent phenomena is that of scale. What we are exposed to when observing emergent phenomena tends to be one of many possible final states. Before that state is reached, however, the system undergoes a series of internal processes that render whatever pattern it is we perceive at that moment. In some case it doesn't take long to perceive emergent behavior.

For example, when observing the flocking behavior of birds one knows almost immediately that such behavior is characteristic among groups of birds. It remains fascinating, at least for me, but it doesn't require much time before one can relate to the process. If one could get really close to the flock, the perspective and understanding of it would change. That "zooming in" or amplification of the dynamics could provide new insights about an individual's behavior and its effect on neighboring birds.

In the case of ants, a colony can live up to 15 years. S. Johnson (2001, p. 80) describes how scientists trying to understand the group dynamics and collective intelligence of ants need to spend years tracking their behavior. Only after a few years of observation can they begin to draw well grounded conclusions about the colonies behavior, trends and adaptability to their environment. The ability to zoom in and out of the system and amplifying or magnifying the interactions taking place, lends itself for the exploration of elements that are not so obvious at first glance but that are nevertheless important characteristics of the system.

1.4 Design

In order to design a system where a certain type of emergent behavior is a defining feature one is after, understanding how emergent properties come to be is quite relevant. By doing so, a line is drawn dividing random behavior, with its unique aesthetic features, from that which is emergent. Even if the behavior may seem random to those being exposed to it for the first time. Those emergent qualities should be an essential part of the systems design considerations rather than a by-product of them.

Complex systems with emergent properties are common in nature. Despite the multiple manifestations of such systems around us, 'realistically complex systems (e.g. organisms, societies, ecologies...), however, are characterized by a multi-level structure' (Heylighen, 1989, p. 2). The complexity of the interactions between components, which can be 'subsystems which are themselves the product of passed emergences' (Beurier, Simonin, Ferber, et al., 2002, p. 1), increases at each level. Better stated: 'an emergent whole at one level is merely a component of an emergent system at the next higher level' (Heylighen, 1989, p.2). Observing what happens at each level may provide an explanation for the pattern or behavior perceived but says little about the overall functioning of the system itself. It provides an overview of the state of the system at a given point in time under certain conditions.

S. Johnson (2001, p. 78) proposes 'five fundamental principles' that need to be followed if one wants to design 'a system where macro-intelligence and adaptability derive from local knowledge'. Although the system I will describe in chapter 3 does not have an adaptive nature at this point, it is a nonadaptive complex system, following these principles was of great help when thinking about the process I wanted to design.

- More is different: It's only by observing the entire system at work that the global behavior becomes apparent
- Ignorance is useful: Better to build a densely interconnected system with simple elements, and let the more sophisticated behavior trickle up
- Encourage random encounters: Those encounters eventually allow the individuals to gauge and alter the macrostate of the system itself
- Look for patterns in the signs: Pattern detection allows meta-information to circulate
- Pay attention to your neighbors: Local information can lead to global wisdom

The philosophical discussion on this topic is very interesting and sometimes confusing. The application of some of the core ideas of the theory within an artistic context provides original aesthetics experiences. The following sections will try to guide the reader along the path I followed when applying some of the principles described previously in the design of a generative system for the production of sound structures. The flexibility of tools such as SuperCollider allows one to explore these ideas and produce interesting results with very basic knowledge of the programming language.

Chapter 2

Emergence and Sound

The ear is a complex system. Much is known about its anatomy, 'but our sense of hearing is understood only in part' (Pierce, 1983, p. 96). Other factors such as memory, upbringing and training influence the way one may experience and interpret sound. There are also many external variables that play a significant role in the perception of sound. As mentioned in the introduction, I wanted to find out whether or not an emergent system can produce interesting audible patterns. When listening to a sound process which is the output of a complex system, would I then be able to identify aspects of that sound result not only as emergent but as characteristic of that system in particular? Designing a system based on the principles behind complex systems is not really a novel thing nor is it some activity leading to groundbreaking findings. But it is interesting and is a rather simple and flexible way of generating sound structures.

2.1 Looking at Nature

I am by no means the first one using the output of complex systems as either a sound source or as a control signal for sound processes. I think, just as many others have, that there are certain aspects and qualities of processes which display emergent behavior that lend themselves for artistic exploitation. Morimoto (2010) and Manousakis (2006) have explored sound synthesis possibilities using Cellular Automata and Lindenmayer systems respectively. Further, interesting and well-documented, applications using complex adaptive systems for live improvisation can be found in Blackwell and Bentley (2002) and Impett (2001). Through models or simulations of natural occurring phenomena we have come to understand much more about different processes in the natural world and their possible application in an artistic context. Abstractions of certain aspects of these processes have made it possible to translate them to meaningful data suitable for sonic exploration.

Flocks of birds, ant colonies or a swarm of insects can produce interesting patterns. Those patterns can be the result of their behavior as a group, or can be found in the structures they build. A computer simulation of any such behavior can successfully replicate the visual effect of the phenomenon as it is experienced in nature. However, our visual perception of patterns produced by emergent behavior is very different from our aural perception of them. That is 'due to our ability to perceive all elements of the scene at once, it is unclear how the ear might be sensitive to emergent patterns that depend on this overview' (Davis & Rebelo, 2005a, p. 2). They argue that in a visual representation of a complex system, our perception of the patterns or behavior depends and is restricted by our field of view. Factors such as distance from the subject matter, perspective and familiarity with the phenomenon play a significant role. We can decide how long we want to be exposed to it, but we still get the "whole picture" after a few seconds of exposure. In the sound domain it works differently since 'aural stimuli are mapped around our own body'.

Designing a physical model or a computer simulation that can accurately render most of the distinctive behavior of a complex system is an interesting way of exploring those systems. Many successful models have been developed. See Schacher, Bisig, and Neukom (2011) and Jones (2008). Since my interest is in the perception of certain aspects of emergent behavior, a level of abstraction is needed in order to highlight those qualities I am after while addressing those aspects of the process which produce them. Dorin (2001, p. 47) states that 'the programmer is in the unique position of being able to describe and manipulate abstract processes which may be used as a unique means of artistic expression'.

In order to be able to describe or manipulate abstract processes, is important to not only identify those aspect that will or could be translated. Davis and Rebelo (2005a) refer to the work of American scientist and Professor John Holland, known for his contributions in the study of complex adaptive systems. According to them Holland suggests the following:

> for a model to be "successful" it should provide a metaphor for a system that enables us to see new connections with, or add new meaning to, processes in the already existing system. He goes on to say that 'deeper extended metaphors [should] allow for a profound re-conception of the subject matter' (Davis & Rebelo, 2005a, p. 1).

Holland introduces the idea of metaphors as a means for translating and highlighting aspects of an existing model to meaningful data that can be interpreted by the subject of cognition. They go on to say that such 'metaphorical relationship is dependent on the identification of concepts that will come into focus when considering two different systems. The model of the emergent system needs to act as a metaphor of the original system'. Technical accuracy or mimicking are, in my opinion, of second importance as long as the system being designed can successfully translate those aspects I want to address. After all, it is how these things can be translated to sound in a meaningful way that makes them interesting in this context.

Using metaphors as a means to translate aspects of an existing system makes it a lot easier to figure out ways in which one could further explore characteristics of complex systems that could be used within a sonic context. Although I was not trying to model a particular biological system, by keeping this approach in mind I was able to mold my conception and understanding of complex systems. They became more than just a bunch of agents producing interesting and hard to predict patterns. What I used to look at and find fascinating became the source of a very interesting way of thinking about, listening to and creating sound. The same way in which these systems make effective application of the design principles and elements in a visual context had suddenly, in my mind, the potential to be used as a model for the creation of a system that could generate sound structures. It may sound logical to some but it took me some time before I concluded that. It was very rewarding to know that I could use similar concepts from the visual domain, which I was familiar with, in the sound domain. All that was needed was the right abstraction. A way to describe the process I had in mind in the simplest way possible so that I could use the limited amount of programming skills in an effective and somehow meaningful way.

> We are similarly inspired by biological metaphors. However, rather than replicating the visible forms of nature directly, we aim to capture the processes that generated them, such that the mechanisms themselves become the object of appreciation as much as the forms they generate. Due to its flexibility and programmability, digital technology is an obvious and suitable medium to explore this methodology fully (McCormack, Eldridge, Dorin, & McIlwain, 2009, p. 356).

2.2 Properties

Before I started this research, my interest was to create a process which would only require me to press a key or compile a line of code and the result would be a sound process with a duration relative to the amount of events I would like to have. Ideally, the result would always be slightly different than previous ones, although the function generating the sound result would be the same. It took some time before I realized that in order to explore the potential of emergence as a sonic construct, constraining the amount of properties that I would explore was much more fruitful than trying to investigate as many aspects as possible at once. Understanding the idea of a system from a 'perspective defined by mechanisms rather than materials' (McCormack et al., 2009, p. 355) seemed a reasonable thing to do. 'It suggests understanding the world in ways that favor process dynamics'. This approach allowed me to focus on the interactions, as was suggested previously in section 1.4.

I felt the need to identify and select some aspects which were artistically attractive. They are presented in arbitrary order:

- Large network of components: The simple design of each component makes it easier to focus on its interaction with other components and with the environment they share. If the rules dictating their behavior are clear, the components can effectively and accurately follow those rules.
- Development in space and time: Through grouping, emphasizing and selecting between states, patterns and structures are formed over time. These systems make effective use of the space were they exist.
- Multiple levels: The dynamic interaction between the components of a system usually occurs in multiple levels. It is a characteristic of complex systems that the result at one level tends to become obvious one level above that where the interaction is taking place. This opens the possibility of zooming in and out of the process in order to reveal aspects of the processes which are not easily perceived but still play a role in the overall character of the process.
- States: Oscillating between states is a common property of all the components. A state could be something as simple as being active or inactive, but can be extended to encompass more dynamic behavior.
- Balance: Complex systems which display emergent behavior are always looking for balance. In order to achieve a state of balance the components usually switch between states of disorder and incoherence and states that seem more organized and balanced. Those transitions offer some variety that is not really present in a state of balance or disorder. They are states of transition that happen over a certain period of time. Exercising some type of control over the time in which the transition takes place, reveals new details about the process.

So how can these elements be preserved and translated to a sound process? I am interested in exploring the process in which those properties and the components of the system interact with each other. Before describing how I did that, there are some concepts that need be explained in order to have a better understanding of the context in which these ideas developed.

2.3 The Natural, Abstractions and Rules

Human creativity has had a peculiar relationship with the concept of nature. It has served as a source of inspiration for centuries. ‘The self-conscious human pleasure in “nature" is often traced back by historians to the eighteenth century’ (Huws, 2000, p. 33). More recently, scientists and engineers have been exploring and exploiting nature, trying to introduce aspects of the processes found in it to the technology which runs our lives. Most of the creative work involving nature as a source of inspiration, until the use of the computer as an art medium, tended to have a timeless quality. A frame, a snapshot of a process seen in nature was captured and translated to different media. Huws (2000, p. 35) explains that 'because of the limitations of the media available, most of the visual representations of nature available to us are static.' It is possible to aesthetically enjoy the natural beyond such 'frozen' depictions. Our fascination comes 'from the continuing surprise of watching how each moment evolves into the next. The very evanescence of the event forms part of the attraction', she adds. We are fascinated with the wide range of behaviors that can be found in nature. I'm interested in exploring and describing those moments of transition between states of order and disorder, a characteristic shared by most natural processes.

Computer simulations as tools to explore the potential applications of biological systems were first used by scientists. Commercial availability of the computer provided people outside the research facilities access to the means necessary to apply similar concepts but with a different intention. 'With the advent of computers, even more elaborate methods, or algorithms, have been used. Various techniques from different research areas, most notably computer science and mathematics, have proven efficient as creative tools for the artists' (Dahlstedt, 2005, p. 1). It is important to note that the use of systems for the generation of art has little to do with the existence of the computer. 'Generative art preceded computer art, and in fact is as old as art itself' (Galanter, 2003, p. 15). The computer has extended the possibilities to formalize methods for the production of creative work, but generative art is not exclusive to the use of technology.

> Generative art refers to any art practice where the artist uses a system, such as a set of natural language rules, a computer program, a machine, or other procedural invention, which is set into motion with some degree of autonomy contributing to or resulting in a completed work of art (Galanter, 2003, p.1).

In the composition of music, abstractions and systems based on or derived from natural processes are well known. A classic example is the use of Cellular Automata by Iannis Xenakis.

> They are very simple rules which can create structures on very large surfaces. It's related to the nature of fluids, for instance. For me the sound is a kind of fluid in time – that's what gave me the idea to transfer one area to the other. I was also attracted by the simplicity of it: it's a repetitious, a dynamic procedure which can create a very rich output (Solomos, 2006, p. 5).

Using generative systems to create sound structures has been common at the Institute of Sonology since the 1970's. See Berg (2009) for a comprehensive summary of some of the systems developed in that period. The practice of creating rules and systems for the composition of electronic music or generation of sound structures requires an understanding of systems where 'simplicity, generality and abstraction are the keys' (Berg, 1996, p.25). A common denominator in generative processes using the computer is the use of algorithms for the description of abstractions. The purpose of the algorithm may be specific, i.e to be used only in one system or composition, but ideally the solution it provides should be general. Seeking general solutions to specific problems is one of the advantages of using algorithms. Berg (1996) suggests that 'the necessary compositional abstractions are not hierarchical descriptions of entire compositions, but generators of musical material or gestures.' In the visual arts we tend to regard the final product of a creative process as the representation of that process. We are not presented a process at the art gallery or art show. We are not so interested in the process that led to a piece but are rather interested in what is in front of our eyes. The final piece may be the result of a creative process where some abstraction could have been involved. But the process may be seen as product oriented, that is to say, the process is important as long as it produces a final, concise, piece. The result doesn't really have to reflect the process by which it came to be. In the case of electronic music using generative processes or rules of some sort, the process itself takes the leading role, this idea will be further discussed in section 3.2. This, again, allows one to focus in the way the components of the process interact rather than specifying how each component should behave.

2.4 Listening to Nature

In a previous section I mentioned some aspects of complex systems that I think have some potential for exploitation in a sonological context. I also posed a question regarding the possible ways in which those properties could be translated to sound without affecting the integrity of the process that produced them. A plausible answer is using algorithms.

> The advent of computers gave us the possibility to formulate, abstract and simulate 'natural' and 'synthetic' environments in radically novel manners. The quest for abstract modelling and organization of complex musical behaviours has manifested itself in the various algorithmic procedures that have been broadly used throughout the history of electronic music (Manousakis, 2006, p. 43).

I wanted to design a process from the bottom up. That means that the generative process itself should have some characteristics that could define it as a complex system. Whether or not it would result in some sound process exhibiting emergent behavior became the new question. After some attempts at designing the generative process I was after and some interesting discussions with Paul Berg, some elements regarding the algorithm's design came up. The algorithm driving the process should be distinguishable. It should have some characteristic behavior and most importantly it should do something. Namely, it should influence the process and introduce change.

> The concepts of process and algorithm are closely linked with those of dynamism and change, with becoming. When a process creates a new entity or brings about novel circumstances, it is a generative process with respect to the change(s) it brings about. Why not explore this concept of change through algorithmic means? (Dorin, 2001, p. 49)

Di Scipio (1994) describes the common concept between the different experiences in composition and sound synthesis since the late 50's until 1994 as follows:

> The common point is the fact that the composer's model of sonic material behaves like a model of *micro-time sonic design* [Di Scipio, 1993b]. To say it with different words, models of sound material and of musical design become inseparable: the compositional process is applied at the scale level of the microstructure of sound and yields the structured organization of myriads of minimal units (each of which, taken *per se* could not be perceptually significant) (Di Scipio, 1994, p. 205).

His "Theory of Sonological Emergence" was a departure point for the application of the ideas described in section 2.2. He says that:

> each approach of algorithmic micro-structural design proposes a Theory of Composition as a Theory of Sonological Emergence:*how to determine a ground-level system's or process' quantitative organization capable of bringing forth a meta-level system or process of peculiar qualitative, morphological properties* (Di Scipio, 1994, p. 205).

This relates to the idea of emergence being defined as the movement from low-level rules to higher-level sophistication given in section 1.2. Di Scipio's "Theory of Sonological Emergence" suggests that 'the emergence of a high level should happen through grains and samples' (Meric & Solomos, 2009, p. 59) with an emphasis on the micro-dynamics of the system by composing the interactions among it's components. Di Scipio (1994, p. 205) says that 'the task of microcompositional strategies can be described as one of letting global morphological properties of musical structure emerge from the local conditions in the sonic matter.'

Examples of systems which are made up of thousands of components giving rise to an "emergent listening environment" (Davis & Rebelo, 2005a) can be found in nature. A good example is that of the periodical cicadas, see Cooley (n.d.). After being underground for 13 or 17 years they emerge for their only mating season. While they are at it, they fill the air with their mating call and create an environment made up of multiple layers of sound. By getting really close to one of the cicadas, see BBCWorldwide (2008), one becomes aware of sounds that are by no means perceivable when exposed to the system as a whole.

Up to this point I had a clear overview of the concepts and ideas I wanted to focus my attention on when designing a generative process. I also had an idea of the synthesis method I would use for my initial tests, granular synthesis with frequency modulation. It provides a wide variety of sounds and small changes in parameters can yield very different results. I still had to come up with a set of rules or states, and a way of introducing them at different levels in the system I was going to design. Di Scipio (1994, p. 206) says that 'except for very few cases, it is quite difficult to find computer music systems which let the composer decide where - at which scale of structure - to enter the compositional process'. This encouraged me to introduce at least one parameter that could be changed per level. Since I was going to design a system made up of multiple levels, that seemed feasible. Another point I tried to keep in mind throughout the design process was that of simplicity. Simplicity in terms of design. In his paper 'Hacking Cellular Automata: An Approach to Sound Synthesis' Morimoto (2010, p. 3) concludes that 'adding complexity to the underlying rules for a system did not ultimately lead to more interesting musical result.' I decided to use this conclusion as advice. Simplicity in the design of the system was the strategy behind the materialization of the abstract process I had in mind.

> Digital computers are simulation machines *par excellence*, ideally suited to exploring the innovative systems-theoretic idea that processes can be generalized through *abstraction*, decoupled from their material origins, yet retain the essential dynamic properties that make them iteresting. If flocks can self-organize, then why not sonic structure? (McCormack et al., 2009, p. 364)

Chapter 3

A Generative Process

Designing a generative process for the production of sound was the main motivation for the research that was described in previous sections. In this chapter I would like to give an overview of the thought process I experienced during the development of a generative process. As mentioned in section 2.2, I had a general idea about the type of process I wanted to design. At the time, I didn't have a clear idea about how the results would be presented. As the complexity and the richness of the output generated by the process increased, it became clear it should be played in realtime.

3.1 Expectations

In his article 'Music As a Gradual Process', Reich (2004, p. 304) states: 'I am interested in perceptible processes. I want to be able to hear the process happening throughout the sounding music'. This statement resonates with some of the ideas and expectations I had when I started designing a generative process using SuperCollider. It was my interest to design a process that would generate a particular output, with perceptual characteristics that could be recognizable although the output would never be exactly the same. The material would have a character, a certain way of articulating the rules specified, which would make it distinguishable. As far as the sound synthesis goes, it wouldn't be restricted to one method. Flexibility in that respect was one of my main concerns. I wanted to know if the process would be as effective with, for example, any type of granular synthesis as it would be with additive synthesis. It was also important that rules could be added or taken away without having to restructure the whole algorithm. Following a modular approach could facilitate this.

Finding a way for the process to manifest itself in the sound was a crucial point. Some sort of change would need to come from within the process. Any change would have to be significant and coherent. Davis and Rebelo (2005a, p. 2) say that 'the perception of emergent systems can be said to be intrinsically related to engagement', the product would have to be interesting enough to captivate the listener's curiosity. But neither too saturated nor too empty so as to bore him or her. Galanter (2003, p. 8) argues that 'working artists understand that an audience will quickly tire of both a highly ordered and a highly disordered aesthetic experience because both lack any structural complexity worthy of their continued attention'. A possible way of exploring the 'character' of the process would be by introducing drastic changes to the environment where it exists and listening to its behavior while it self-organizes. Testing the boundaries of the process in order to disclose its self-organizing character could also reveal aspects of the algorithm that may have not been anticipated in the design phase but that could still be creating unwanted effects. Those effects could be audible, e.g. clipping; or could be inaudible but having a direct effect on the sound result, e.g. CPU overload due to very short wait-time between events tends to produce audible clicks. This exploration could also lead to the discovery of features that were not part of the initial design but that are nevertheless audible and desirable. The states defining the behavior of the process should provide enough repetition of patterns so that the listener can become familiar with it and should also provide enough variation in order to avoid monotony.

Huws (2000, p. 37) states that 'it is only recently, with the introduction of information technologies, that it has become possible to envisage forms of art that make it possible to explore or model the way "nature" behaves (as opposed to how it looks) in any depth.' I wanted to design a process that would behave '*like an organism*' (McCormack et al., 2009, p. 355). It would grow and shrink, while oscillating between states of order and disorder and eventually pass away.

3.2 Design

As mentioned in section 1.4, complex systems are characterized for having a multi-level structure and can be made up of sub-systems, which can be the product of another complex system. There are time variations in the interaction among components at each level. Some actions may take longer than others before they are carried out. Transitions from one state to the next happen at different moments and their effects on the system vary per layer. I wanted to able to zoom in and out of each level, amplifying or magnifying the interactions and transitions between states. That, in order to explore aspects of the process that are not quite perceivable when listening to it without dissecting and examining the multiple layers. This also introduced the idea of having the same process running multiple times in parallel. Each level made up of other systems that are running either in sequence or in parallel. Each level is a different instance of the same process but having slightly different input values for its parameters and an independent clock. I wanted to have no influence over the parameters once the process got going, it should have some degree of autonomy, which is a common characteristic of generative processes:

> In philosophical terms, generative computer programs are examples of the derivative intentionality of writing (Searle 2004: 20), where the code is predetermined by a human author who then yields moment to moment autonomy of execution to the machine. Human intervention is thus reduced to initial conditions alone; control is usually choosing when to start the program, providing the seed for the random number generator (Collins, 2008, p. 238).

Coming up with a generative process was the main goal of this project and my programming activities pretty much oscillated around that goal. After some tests and fulfilling aural results, I realized that this aspect of the project, designing the algorithm, was an exercise exploring simplicity. Less is more. I wanted to keep the system's design as simple and flexible as possible. That proved to be a challenge, given the almost infinite possibilities that SuperCollider has to offer. Setting some constraints, such as reducing the amount of control one can have on the process, the amount of parameters that can be changed and limiting the sound synthesis to FM granular synthesis allowed me to become familiar with the material, both the process and the product of it.

In order to become acquainted with the different aspects of the generative process, gradually increasing the complexity of the interactions between the multiple levels of the system, while listening to the sound result, was crucial. In this way it was easier to keep track of the overall behavior of the process. This is especially important when working with multiple complex systems as the changes introduced at one level can yield unexpected results at other, interrelated, levels. Another aspect that is worth mentioning is how the exploration of satisfactory combination of values for the different parameters occurred. Instead of using presets I decided to explore changes of parameters by working with the idea of slightly different concentration of values around specific "reference values" using random distributions with a shared probability value. Multiple beta distributions producing values properly mapped to each parameter's range are used for different parameters, but they all receive the same value for the hi and lo probabilities. This introduces enough variation so as to keep the relation between the multiple levels of the process coherent and allows one to focus on the exploration of satisfactory combination of values rather than mapping each individual parameter at each level individually. It also allowed me to become more familiar with the material and to understand the relationship between the multiple components of the process and the sound structures they produce. As it is mentioned in C. W. Johnson (2006, p. 3) 'An effective and efficient design could not usually be achieved without a proper understanding of the relationship between the whole and its parts as well as the emergent properties of the system'.

Obtaining a rich output had nothing to do with complicated rules or intricate sound synthesis networks, not for me at least. As mentioned in section 1.4 'Ignorance is useful'. Instead of modeling complicated individual behaviors, my focus was more on programming the interactions between components whose design and default behavior was rather basic. This was motivated by the idea of 'better to build a densely interconnected system with simple elements, and let the more sophisticated behavior trickle up ' (S. Johnson, 2001, p. 78). According to Manousakis (2006, p. 43): 'The act of simulating in general requires representing certain key characteristics or behaviours of a natural or abstract system in a compact and simplest possible way'. My interest was in simulating an abstract system based on the ideas described in 2.2. A system made up of multiple levels of large networks of components developing in time and space, each component switching between states and executing the rules provided to it at each state. Such a formulation does not provide any clues as to how the system could be implemented, how it could behave or what kind of output it could generate. It is a general description of the system's functioning and its components. Before I could start thinking about a generative sound process it was necessary to think about the system that would be producing it.

The system is designed using SuperCollider's Pattern class and is made up of the following elements:

- **Prototype**: This is the most basic component of the system. It can be seen as a blueprint or a template from which multiple "clones" can be derived. Any change on any of its parameters will have an effect on all of its dependencies. The prototype, see 3.1, is independent of any synthesis method, it is simply a placeholder for parameters which are determined by the synthesis instrument being used. An important parameter is the *delta* value, i.e. the time in-between events in a sequence. It is built-in the SuperCollider Pattern class and is independent of any synthesis method. This is similar to what Koenig (1985) calls "entry delay". In this case the default event pattern is made up of three granular synthesis events. The duration of the whole event is about 7 seconds and the *delta* time is 3 seconds, see figure 3.2. The generative process is based on variations of this basic granular pattern.
- **Clone**: Copy of the prototype. It inherits the prototype's functionality and parameters. Parameters may be changed in real-time without affecting the behavior of the prototype or any other clone.
- **Loop**: In order to be able to change any parameter in real-time it is necessary to play the clone in a loop. After experimenting with multiple loops and changing parameters in real-time I realized that although the results were interesting, they were mostly influenced by me. I was the one to decide the combination of loops, their duration, their *delta* values and the values for other parameters.

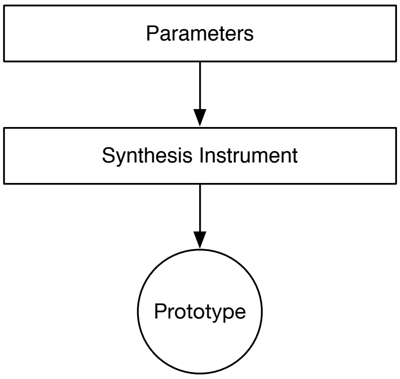

Figure 3.1: Prototype

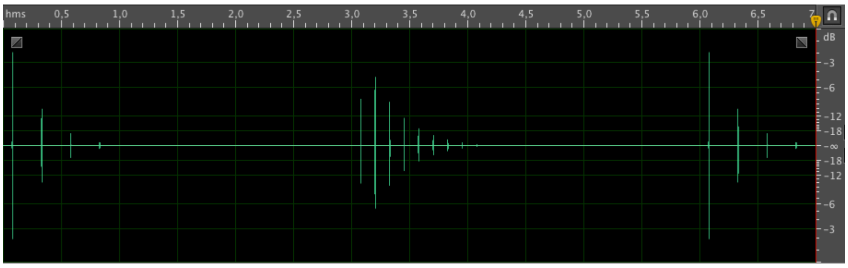

Figure 3.2: Default sound pattern

- **Spawner**: I needed to find a way to define some rules that would affect the behavior of each clone while looping, without my constant intervention. This is where I started working with the idea of a state engine, transitions between states at different times and manipulating the time in which those transitions take place. *Pspawn* is the SuperCollider class used for this purpose, it ‘spawns sub-patterns based on parameters in an event pattern’ (*SuperCollider Help*, n.d.). Some advantages of this approach are:
    - multiple instances of a clone, or different clones, can be looped and played either in parallel or in sequence
    - a *delta* value may be provided that determines the time in-between events produced by the spawner. This is independent from the *delta* value of each clone
    - it is a compact and economic way of grouping multiple clones. It is easier to refer to a group of clones based on their behavior than it is to recall each clone by itself
    - a spawner can be made up of other spawners. This was an important development in the system’s design. It relates to the idea of complex systems being made up of components which can themselves be sub-systems. In this case, multiple clones can constitute a single spawner and multiple spawners can be combined into yet another spawner. This makes it possible for rules to be introduced at different levels in the system. It allows a flexible exploration of the system’s states and combinations thereof.

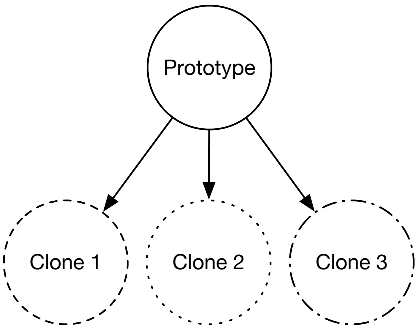

Figure 3.3: Clones

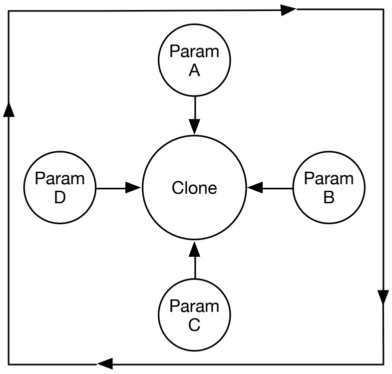

Figure 3.4: Loop

- **State engine**: A Deterministic Finite Automaton or Deterministic Finite State Machine (DFA or DFSM) is used as a state engine. 'DFAs naturally lend themselves to concisely representing any system which must maintain an internal definition of state' (Gribkoff, 2013, p. 1). In the system being described here multiple state machines are used at different levels and play two fundamental roles. Before explaining the roles and the states embedded in the DFAs, it is necessary to briefly introduce the concept of finite automata:

  > Finite automata are the simplest mathematical model of computers. Informally, a finite automaton is a system that consists of states and transitions. Each state represents a finite amount of information gathered from the start of the system to the present moment. Transitions represent state changes described by the system rules. Practical applications of finite automata include digital circuits, language design and implementations, image processing, modeling and building reliable software, and theoretical computing (Khoussainov, n.d., p. 1).

  A finite automaton usually consists of the following:

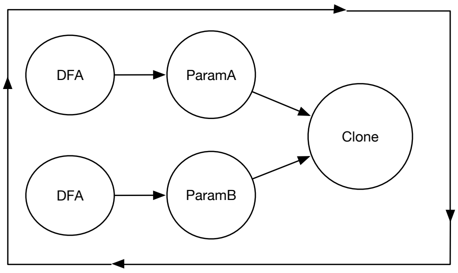

Figure 3.5: Spawner

- states: A finite number of states
- input alphabet: These are the input values, in this case the alphabet is a finite list of symbols. In each state a return stream or pattern must be provided for each input value. Otherwise either a default stream must be given or the process must be explicitly told to stop.
- transition function: Function which is evaluated before moving to a new state. Here is an example consisting of two states:
  State 0 = (A: [0, Beta distribution] , B: [1, Beta distribution])
  State 1 = (A: [1, Beta distribution] , B: [0, Beta distribution])
  where A and B are input symbols, 0 and 1 are the states and Beta distributions are used as the transition functions which provide values for the parameters. If, for example, the initial state is 0 and the input symbol is B, the beta distribution will generate n values and move to state 1, where a transition function will be chosen depending on the next input symbol.
- start state
- set of final states

Figure 3.6 is an example of a DFA with three states and starting at state 0.

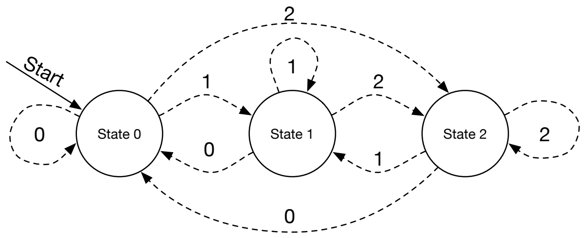

Figure 3.6: Transition diagram DFA

A deterministic finite automaton differs from a non-deterministic one in that for each state in the DFA there is one possible input while ‘a nondeterministic finite automaton has the ability to be in several states at once [and] the transitions from a state on an input symbol can be to any set of states'. I have chosen for a DFA because it allows one to define a list of input values to determine state transitions. In this way it is possible to use a DFA in multiple spawners but with a different initial state and an altered alphabet. Altering the alphabet is done by using a weighted random distribution with repetition check.

As depicted in figure 3.7 a DFA in this system is basically made up of two states. Each state contains two beta distributions which are chosen based on the input symbols. This is one of the two fundamental uses of DFA in this system, the generation of appropriate values for the different parameters of the sound synthesis instrument.

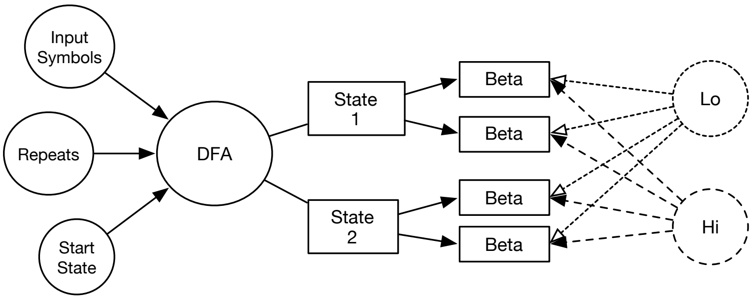

Figure 3.7: DFA

- **Parallel spawners**: The spawners are looped and played in parallel. Each process becomes a sub-system, a layer, of a larger system. At each layer values are provided for some parameters which can be changed, individually as well as parameters common to all the layers, in real-time, see figure 3.8. This allows one to introduce changes to each looping process, or to the parallel system as a whole, at different times and to group parameters per layer. Each layer or sub-system can be played individually in order to produce a sound structure, however, multiple parallel spawners are grouped in another spawner in order to have multiple process running at once.
- **Generator**: The parallel processes described previously constitute individual states in a larger DFA which is the actual sound-generating mechanism of the system being described here. The sound-generating DFA can be seen as a parallel spawner generator, where the structure of the resulting sound is determined by the input symbols, the number of times a given parallel spawner will be played and the starting state. This is the second fundamental use of DFAs in this system: selecting and playing the different parallel spawners embedded in its states. A state in the generator is depicted in figure 3.9. A state is made up of two parallel spawners played, again, in parallel.

Since the early stages of the design process I followed a bottom-up approach. The exploration of values and rules was done in real time. Collins (2008, p. 237) describes how time consuming it was to compose an algorithmic piece of music in the 1960's and concludes that 'it is not

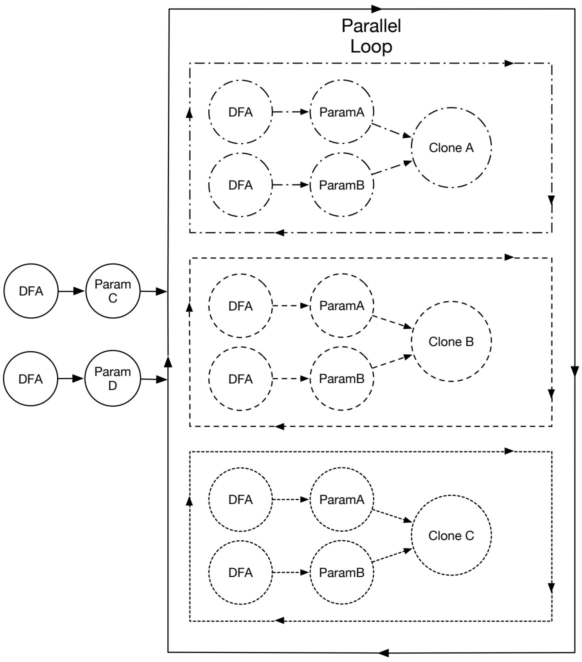

Figure 3.8: Multiple Spawners in Parallel

surprising that most contemporary explorations favor realtime systems for speed of feedback, a development that also supports complex interactive possibilities'. Although the system I designed does not have an interactive nature, it is a closed system, the design process did involve a high degree of interaction with the system. He suggests that 'anyone who has created CG-art knows the intimate negotiation between design of a program and feedback from program output. In this sense, during the design cycle CG-music is highly interactive'. Berg (2009, p. 76) points out something similar when summarizing the 'notable features' of the work of composers at The Institute of Sonology in Utrecht in the 1970's. 'This interaction underscored the importance of listening during the process of generating music using rules'. Being able to peek into the process as it is being designed offers a good overview of the state of the system and provides a direct and almost immediate impression of any change imposed by the rules.

After being exposed to the output of the process several times, familiarizing myself with some of the possible states that could be reached and being pleasantly surprised by some of the unexpected results, I was quite confident that the system did have some perceivable emergent behavior. Despite being able to predict some of the behavior based on the rules provided by the DFA, there are artifacts produced by the interaction of multiple instances of the process that I can, as of now, not really reduce to a single cause. At this point I also stopped making a difference between the process and its output. 'The concept of material was not limited to sounds but also included compositional methods and rules' (Berg, 2009, p. 77). The system being described does not represent a particular natural process or complex system. It is an

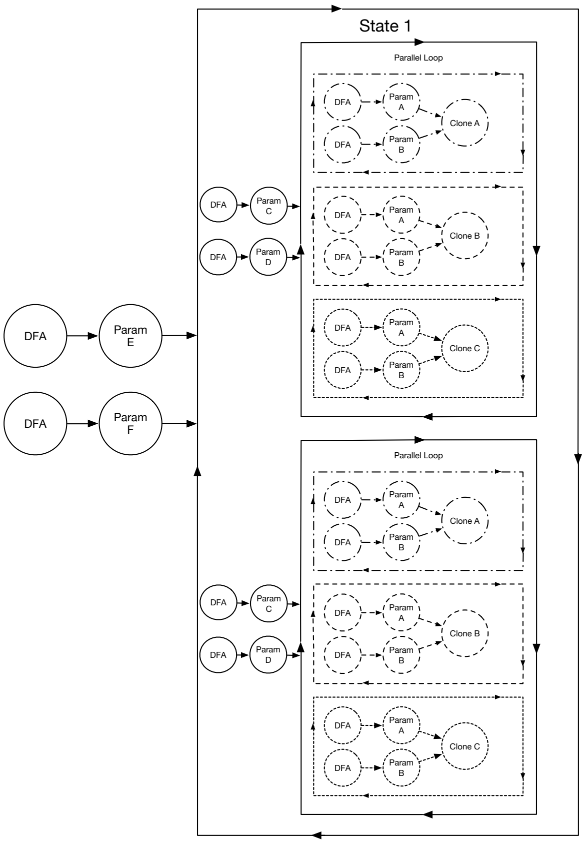

Figure 3.9: Example of a state in the sound-producing DFA

abstraction of a complex system that generates sound structures. I can introduce changes to the environment where the process exists and the interactions taking place, but I have no direct control over the sound process itself.

3.3 Mapping and Scale

Different approaches used to explore different aspects of complex systems provide very different results and yield new insights about the applications and artistic potential of such systems. Most of the models available and documented that simulate natural occurring emergent behavior can be said to depend on a visual representation of the phenomena and the sound aspect to depend on the visual representation.

> In the context of sound design, a potential problem with many of these systems is that the algorithms they are based on do not derive from sound but typically from a system that exhibits perceived emergence through the application of graphical (Boids) [13] or evolutionary (Genetic Algorithm)[1] models. The implementation of these algorithms to a sound world is then based on a more or less arbitrary

> mapping procedure between a graphical and a sonic model (Davis & Rebelo, 2005b, p. 1).

I have followed the idea proposed by Davis and Rebelo (2005a, p. 1) of the 're-thinking of the mapping process'. They propose using 'a translation of function rather than a mapping of results'. Translating aspects of the model in mind to meaningful information can be done by considering the listener and the space as components of the system. Since my interest was to explore the transitions between states of saturation and balance occurring in multiple process running at once, I wanted to be able to hear both states, saturation and balance, simultaneously on different loudspeakers. In order to achieve this I use reversed values in the beta distributions for the sub-systems (the parallel spawners) of the sound-generating DFA, that means that for every set of parallel spawners the list of parameter values is the same but for one of the parallel spawners the list of input values is reversed. This produces a richer output and it also makes it possible to hear the process in both a state of saturation and minimum activity at the same time. Because of the time differences introduced at multiple levels of the system by altering the different *delta* values, the system will rarely get to a long-lasting state of complete saturation and overload or a state of complete silence. Mapping parallel processes in this way to different loudspeakers in a room surrounds and immerses the listener in the material, hopefully instigating exploration of the material and the space.

Considering that 'one of the most fundamental challenges in Generative Art relates to the establishment of meaningful and traceable mapping relationships between the underlying algorithmic processes and the resulting aesthetic output [2]' (Bisig, Schacher, & Neukom, 2011, p. 260), it is important to first decide which aspects of the system one wants to highlight and then find a way to establish a coherent relationship between the sound material and how it is affected by the algorithmic process. If this is taken into consideration from the beginning of the design process, it is then easier to anticipate, to some extent, how the system might behave when implementing new behaviors or rules.

3.4 Saturation and Balance

I am interested in exploring how the senses can be gradually overloaded with patterns. How our brain keeps on looking for patterns in whatever information it receives through the senses. When designing the generative process described above I also became interested in finding out to what extent its output actually made sense, if it made sense at all. If there were patterns being picked up by the ear, either coming from the material itself or from its acoustic interaction with the room and the listener; when do they stop making sense? Can they be picked up again after they stop making sense? 'Hearing is just too important a sense to disregard when dealing with the task of rendering complex information' (Ballora, 2014, p. 30). There is a limit to the amount of information our senses can manage when being exposed to patterns. Testing the boundaries of the process was a way of exploring the tension that usually precedes or follows states of stability or balance. Shifting between saturated states and states where very few things happen was a way for the process to manifest itself. Turing (1952) describes the importance of states of instability in the symmetry of embryos. This way of thinking about transitions between states of saturation and balance and its influence on the overall behavior of the system applies to many processes that display emergent behavior:

> It was assumed that the deviations from spherical symmetry in the blastula could be ignored because it makes no particular difference what form of asymmetry there is. It is, however, important that there are *some* deviations, for the system may reach a state of instability in which these irregularities, or certain components of them, tend to grow. If this happens a new and stable equilibrium is usually reached, with the symmetry entirely gone. The variety of such new equilibria will normally not be so great as the variety of irregularities giving rise to them (Turing, 1952, p. 42).

My intention was to explore the formation of sound structures based on variations made on the prototype's pattern. By mapping at least one parameter at each level of the system to a beta distribution producing a different range of values, but with a shared input value for the distribution values, I could explore the states of the system with different combination of values per layer. Pushing the system to the limit and appreciating how it self-organizes started as a way to test any change exercised on the weights for the beta distributions which produce values for the parameters in the synthesis instrument. It was not until I started working with multiple processes running in parallel, each with a different clock, that I heard how sometimes more than two of these parallel processes would reach a state of saturation. A state of saturation could be the product of, for example, a lot of short-duration clouds being played very soon after each other and overlapping, or combinations of clouds with a long decay in the amplitude envelope and overlapping of the spectral content. By changing the clock's tempo, it was possible to either accelerate or slow down these transitions between states. When perceiving dramatic transitions between states, sometimes happening very fast on one layer and a bit slower on another, a lot of psychoacoustic artifacts are introduced. Such an effect is a lot more convincing when using pulse-like material. After becoming familiar with this aspect of the process I decided to make it a feature of the system. 'Nonequilibrium must be embodied at the highest structural level, as well as the lowest, if the system is to be able to self-organize with its environment'(Impett, 2001, p. 112).

3.5 An Installation

Keller (2000, p. 58) suggests that ‘once a sound model is defined, a range of behaviors can be explored’. When the first satisfactory results were obtained, the character of the material suggested that it would be better perceived if presented as an installation. This decision was motivated by the statement ‘Given the programs, enjoy the result’ (Berg, 2009, p. 82). It felt as if trying to organize, arrange or edit the output of the process in a different way, would take away its essence. The effort invested in composing the interactions and rules, as well as giving away most of the control over the parameters would have been in vain. I wanted the process to just be and immerse the listener. I was very much influenced by the work of Davis and Rebelo (2005b), Bisig et al. (2011) and Felix Hess. They explore the idea of emergence in sound through installations. Although their work and methodologies are quite different they implement abstract models of natural occurring emergent phenomena in order to translate aspects of the process to an audience.

Davis (2008, p. 1) writes that ‘installation artists by the very nature of their work are working in the full complexity of reality and thus have to consider not only the direct experience of sound but also its relation to space and its personal relationship to the individual; architectural and cultural’. I experienced a lot more engagement with the material, and the creative process giving rise to it, when fine-tuning aspects of the process to suit this presentation format than I did when manually editing sound material in order for it to better suit my needs and purpose. Davis (2008, p. 1) argues that ‘working in the field of installation also highlights a level of personal engagement with the work’. Since there is no editing involved in the sound result, it is necessary to familiarize oneself with as many aspects of the material as possible, the space where it will be experienced and the perception of the combination of the space and the material. It is then inevitable to separate the material, the room and the listener. This can be interpreted as a coupling of multiple complex systems. Where the room and the listener, seen as complex systems with their unique degree of complexity, can potentially influence the perception of the material by the listener. The process should consider them as active components of the system.

> In applications of music creation that utilize models based on abstract algorithms there is a need for a tighter linking of the algorithms with environmental and cultural context. Not only to make these algorithms more open ended in nature, not closed off from the complexity of reality, but to make them more accessible to the perception of participants in the work in such a way that there can be a co-evolution of interaction and understanding (Davis, 2008, p. 1).

The installation makes use of the generative process as it is described in section 3.2. Although the installation asks for a 8 loudspeaker setup, this feature was added at the very end of the system's design. The dynamic character of the material and the movements occurring in it were more perceivable when using more loudspeakers. A sense of space and movement is suggested by the interaction taking place at different levels in the multiple processes running in parallel but each on a different clock. Exploring the formation of patterns by having multiple instances of the same process running at once is the main objective of the installation. Keller (2000, p. 58) adds that 'in a general sense, pattern-formation refers to the emergence of higher-level forms or behaviors resulting from the interaction of two or more systems'. Changes in the *delta* value of the clones and the spawners introduces interesting time differences that add some variation to patterns that occur simultaneously on the multiple running processes. This exploration of time differences hopefully encourages the listener to engage with the material and the space where it is being perceived.

In repeated occasions when listening to the process in order to fine-tune the rules and to perceive it's behavior in a new space, I was reminded of an ecosystem. An ecosystem as described by McCormack et al. (2009):

> An ecosystemic metaphor considers components and their interactions in a potentially noisy environment. Ecosystems generally establish a network of tightly positive and negative loops. The system must organize to offer viable configurations, otherwise it breaks down and ceases to exist. In a well-designed system, viable dynamic configurations, are explored and exploited. The system flips between stable states, each of which represents a different aesthetic possibility. The appropriateness of these possibilities is a side effect of a well-designed system (McCormack et al., 2009, p. 376).

3.6 Working with the WFS

Working with Wave Field Synthesis seemed a reasonable way of further exploring the dynamic nature of the generative process I described in the previous section. Although by the time I started working on this section of the thesis I had been working with the WFSCollider, see Snoei, Ganchrow, Truetzler, and Negrão (n.d.) application and The Game of Life system, see *The Game of Life* (n.d.), for no longer than 2 months, it is important to discuss this aspect of my work for it has made way for new ideas and has uncovered aspects of the generative process I had not anticipated.

After the generative process was running smoothly, without clipping or overloading the CPU and crashing SuperCollider, and I had become familiar with the material; working with the WFS seemed the most natural thing to do. Given the character of the the material, it seemed reasonable to have multiple processes running at the same time and spreading them out as point sources in space. This idea of creating spaces has been explored by Natasha Barrett using ambisonics. Although I am not using ambisonic spatialization, her remarks about creating sonic spaces apply in this context as well:

> When you consider very tiny sounds building up an impression of space, then you can begin to imply a space within which these sounds should live. If you use abstract sound material, it can be difficult for the listener to find the spatial context. If you don't want to use reverberation, you don't have a clear spatial context to start off with, either. But gradually, as the sound material unfolds, its behavior, its motion behavior the relation between many things happening at once imply space, even though you are not using reverberation or clear sound identities (Otondo, 2007, p.13).

Instead of having events being arbitrarily mapped to loudspeakers, I could now define various aspects about their localization and the trajectories to be followed. Being able to define how the material could develop in space and time became a new feature of the system that called for new abstractions. In section 3.1 I mentioned how flexibility was a main concern during the design of the generative process. That effort paid off when I started using the WFSCollider application. It only took adding the parameters for the type of source (static or dynamic) to be used to the existing prototype and adapting the existing code to include the classes required to be able to use the WFSCollider application. By doing that, I was able to assign any type of source to the clones. The spatial behavior of each component, or of the whole system, may be defined using the deterministic finite state machine, the beta distribution or can be defined in the clone's parameters and left unchanged. This added a new level of control I had not anticipated. Being able to group parameters in space based on rules added a new way of relating the processes to the space where they happen. Using states not just as a way to articulate, but also as a way to relate certain states to certain movements or places in the room allows the listener to explore the space created by the material.

The WFSCollider imposed some constraints on the way I had been working with the process. Because of memory limitations and my lack of experience with the software, I could no longer have multiple processes running in parallel in the same way I did with the version of the process used in the installation described in section 3.5. In order to save some processing memory and to familiarize with the material produced by this new way of working, I decided to reduce the possible trajectories a source could follow to a single trajectory but with random speed per component. This decision had a strong effect in the perception of the movement of sources in space, having just one trajectory being followed at different speeds makes it easier for the ear to relate to the patterns being produced. It also significantly improved the real-time performance of the process. Which was a main concern of mine, I wanted to still be able to assess and explore the process in real-time. In this version I also decided to work with a different sound synthesis method. Sample based granular synthesis. I wanted to find out whether or not the process was indeed as effective with different granular synthesis methods.

The result up until now have been satisfactory. The constraints imposed by the WFSCollider actually turned my focus to aspects of the process I had overlooked and made me simplify the design of the prototype and spend more time specifying the rules.

Chapter 4

Conclusions

In section 2.2 I mentioned that the process I had in mind should only require me to press a key or execute a line of code and the process would run for a duration relative to the amount of events one would like to have. During the design of the generative process this idea changed quite a bit. I can execute a line of code and the process will produce an amount of sound structures whose entire duration depends on a combination of factors such as the rules specified in the states of multiple DFA used, the amount of times the list of input symbols will be repeated, the *delta* value of the prototypes and the spawners, and the speed of the clock on which the process is being run. This introduces some unpredictability as far as the duration of each sound structure goes, but it forced me to carefully plan the rules of the system and the interaction among its components. The advantage being that the engagement and familiarity with the material increases.

Exploring the idea of emergence in a sonic context has exposed me to concepts about interaction, systems and spatiotemporal relationships in generative processes in fascinating and inspiring ways. It has shifted the way I interpret and understand natural processes and has given me enough theoretical tools to translate aspects of those processes in a meaningful and coherent manner to sound. A couple of valuable lessons I have learned as a result of the research I have been conducting and the creative application of some of the ideas I came across, is the use of abstractions and generalization of concepts, and that of setting constraints in order to explore complicated or not-so-clear ideas. Schacher et al. (2011, p. 100) write that 'the almost unlimited number of possibilities for model customization and transformation into music constitutes some of the main challenges of simulation-based computer music'. By generalizing ideas and by setting constraints on the amount of ideas that will be explored at once, the results that may be obtained in the initial stages might provide new insights and sometimes simpler ways of approaching problematic aspects of a certain idea or plan as the complexity of its design increases.

Studying complex system and emergent behavior within an aural context has exposed me to a new realm of artistic expression where the multiple entities involved in the system, the generative process, the listener and the space, come together and form a type of ecosystem. According to McCormack (2012, p. 129) 'the generalized concept of an artificial ecosystem, which adapts concepts and processes from a biological ecosystem at a metaphoric level, is an appropriate generative system for creative discovery'. I have come to the conclusion that designing complicated and intricate systems for the production of sound structures does not guarantee an interesting output and that complicated models for the exploration of sound tend to result in 'lots of wheel reinventing' (McCormack et al., 2009, p. 376). They add that 'the creative ability of the artist is reflected in their design of system components and interactions - the narrative of its generative aesthetic'.

In my experience, designing a generative process using some of the ideas and characteristics of complex systems is a rather simple and fruitful way of exploring the aural potential of emergence. It is my hope that this thesis will spark the curiosity of the reader and encourages him or her to consider emergent behavior as a viable tool for the generation of sound structures. Some of the ideas described in this thesis may be a way to come up with new compositional concepts for the creation of computer music systems.

> The challenge today and in the near future is to develop new compositional concepts that are suitable for computer music systems. These aesthetic abstractions should be as resilient as the concepts from the 1960s have proved to be. They should represent evocative simplifications of compositional activity. Without the development of some new generators of parametric material, the idea of composing with a computer has little future(Berg, 1996, p. 25).

References

Ablowitz, R. (1939). The Theory of Emergence. *Philosophy of Science*, *6*(1), 1-16.

Ballora, M. (2014). Sonification, science and popular music: In search of the 'wow'. *Organised Sound*, *19*(01), 30–40.

BBCWorldwide. (2008). *Amazing cicada life cycle - Sir David Attenborough's Life in the Undergrowth - BBC wildlife [video file].* Retrieved May 5, 2014, from `http://www.youtube.com/watch?v=tjLiWy2nT7U`

Bedau, M. A. (1997). Weak emergence. *Noûs*, *31*(s11), 375–399.

Berg, P. (1996). Abstracting the Future: The search for musical constructs. *Computer Music Journal*, 24–27.

Berg, P. (2009). Composing sound structures with rules. *Contemporary Music Review*, *28*(1), 75–87.

Beurier, G., Simonin, O., Ferber, J., et al. (2002). Model and simulation of multi-level emergence. In *Proceedings of IEEE ISSPIT 2002* (pp. 231–236).

Bisig, D., Schacher, J. C., & Neukom, M. (2011). Flowspace–a hybrid ecosystem. In *Proceedings of the International Conference on New Interfaces for Musical Expression* (pp. 260–263).

Blackwell, T. M., & Bentley, P. (2002). Improvised music with swarms. In *Proceedings of the World Congress on Computational Intelligence* (Vol. 2, pp. 1462–1467). Los Alamitos, CA, USA: IEEE Computer Society.

Collins, N. (2008). The Analysis of Generative Music Programs. *Organised Sound*, *13*(3), 237–248.

Cooley, J. R. (n.d.). *Periodical Cicadas.* Retrieved May 5, 2014, from `http://www.magicicada.org/about/species_pages/species.php`

Corning, P. A. (2002). The re-emergence of "emergence": A venerable concept in search of a theory. *Complexity*, *7*(6), 18–30.

Dahlstedt, P. (2005). Defining Spaces of Potential Art: the significance of representation in computer-aided creativity. In *Description & Creativity Conference, online proceedings 2005.* Cambridge, UK: King's College.

Davis, T. (2008). Cross-Pollination: Towards an aesthetics of the real. In *Proceedings of the International Computer Music Conference 2008.* Belfast: International Computer Music Association.

Davis, T., & Rebelo, P. (2005a). Hearing emergence: Towards sound-based self-organisation. In *Proceedings of the International Computer Music Conference 2005.* Barcelona: International Computer Music Association.

Davis, T., & Rebelo, P. (2005b). Natural Selection: A Stethoscopic Amphibious Installation. In *Proceedings of the 8th Generative Art Conference.*

Di Scipio, A. (1994). Formal Processes of Timbre Composition Challenging the Dualistic Paradigm of Computer Music. In *Proceedings of the International Computer Music Conference 1994* (pp. 202–208).

Dorin, A. (2001). Generative Processes and the Electronic Arts. *Organized Sound*, *6*(1), 47–53.

Galanter, P. (2003). What is Generative Art? Complexity theory as a context for art theory. In *In GA2003–6th Generative Art Conference.*

The Game of Life. (n.d.). Retrieved May 5, 2014, from `http://gameoflife.nl/`

Goldstein, J. (1999). Emergence as a Construct: History and Issues. *Emergence*, *1*(1), 49–72.

Gribkoff, E. (2013). *Applications of Deterministic Finite Automata.* Retrieved May 5, 2014, from `http://www.cs.ucdavis.edu/~rogaway/classes/120/spring13/eric-dfa.pdf`

Heylighen, F. (1989). Self-organization, Emergence and the Architecture of Complexity. In *Proceedings of the 1st European Conference on System Science* (Vol. 18, pp. 23–32).

Huws, U. (2000). Nature, Technology and Art: The Emergence of a New Relationship? *Leonardo*, *33*(1), 33–40.

Impett, J. (2001). Interaction, Simulation and Invention: a Model for Interactive Music. In *Proceedings of ALMMA 2001 Workshop on Artificial Models for Musical Applications* (pp. 108–119).

Johnson, C. W. (2006). What are Emergent Properties and How Do They Affect the Engineering of Complex Systems? *Reliability Engineering & System Safety*, *91*(12), 1475–1481.

Johnson, S. (2001). *Emergence.* Penguin Books.

Jones, D. (2008). AtomSwarm: A Framework for Swarm Improvisation. In *Applications of Evolutionary Computing* (pp. 423–432). Springer.

Keller, D. (2000). Compositional processes from an ecological perspective. *Leonardo Music Journal*, *10*, 55–60.

Khoussainov, B. (n.d.). *Deterministic Finite Automata.* Retrieved May 5, 2014, from `http://www.cs.auckland.ac.nz/courses/compsci220s1t/lectures/lecturenotes/old/DFA.pdf`

Koenig, G. M. (1985). *Programmed Music.* Retrieved May 5, 2014, from `http://www.koenigproject.nl/Programmed_Music.pdf`

Manousakis, S. (2006). *Musical l-systems.* Retrieved May 5, 2014, from `http://www.sonology.org/NL/SOmain.html`

McCormack, J. (2012). Creative Ecosystems. In *Computers and Creativity* (p. 39-60). Springer Berlin Heidelberg.

McCormack, J., Eldridge, A., Dorin, A., & McIlwain, P. (2009). *The Oxford Handbook of Computer Music* (R. Dean, Ed.). Oxford University Press, USA.

Meric, R., & Solomos, M. (2009). Audible Ecosystems and Emergent Sound Structures in Di Scipio's Music: Music Philosophy Helps Musical Analysis. *Journal of Interdisciplinary Music Studies*, *3*.

Mitchell, M. (2009). *Complexity: a guided tour.* Oxford University Press.

Morimoto, Y. (2010). Hacking Cellular Automata: An Approach to Sound Synthesis. In *Proceedings of SuperCollider Symposium 2010* (pp. 1–10).

Otondo, F. (2007). Creating Sonic Spaces: An Interview with Natasha Barrett. *Computer Music Journal*, *31*(2), 10–19.

Pierce, J. (1983). *The Science of Musical Sound.* Scientific American Books.

Reich, S. (2004). *Audio Culture: Readings in Modern Music* (C. Cox & D. Warner, Eds.). Continuum.

Schacher, J., Bisig, D., & Neukom, M. (2011). Composing with Swarm Algorithms—Creating Interactive Audio-Visual Pieces Using Flocking Behaviour. In *Proceedings of the International Computer Music Conference 2011* (pp. 100–107).

Snoei, W., Ganchrow, R., Truetzler, J., & Negrão, M. (n.d.). *WFSCollider.* Retrieved May 5, 2014, from `http://sourceforge.net/projects/wfscollider/`

Solomos, M. (2006). Cellular automata in Xenakis's music. Theory and Practice. In *Definitive Proceedings of the International Symposium Iannis Xenakis.*

SuperCollider Help. (n.d.). Retrieved May 5, 2014, from `http://doc.sccode.org/Classes/Pspawn.html`

Turing, A. M. (1952). The Chemical Basis of Morphogenesis. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *237*(641), 37-72.