

ROYAL CONSERVATORY THE HAGUE

BACHELOR THESIS

**Agent-Based Modelling as a Method for
Parameter Modulation**

Author:
Chris Vermeulen

Supervisor:
Bjarni Gunnarsson

Declaration of Authorship

I, Chris Vermeulen, declare that this thesis titled, “Agent-Based Modelling as a Method for Parameter Modulation” and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

Date:

“This is an exercise in fictional science, or science fiction, if you like that better.”

Valention Braitenberg

ROYAL CONSERVATORY THE HAGUE

Abstract

Institute of Sonology

Bachelor of Arts

Agent-Based Modelling as a Method for Parameter Modulation

by Chris Vermeulen

This thesis presents a method for designing sound morphologies and spatial textures (Smalley, 1997) with the help of a graphically controllable agent-based model (ABM), which is in functioning similar to boids (Reynolds, 1987). As such a model aims to explore the connection between local- and global behaviour, in a similar manner the relationship between material and form will be discussed, thereafter directly compared to the internal structure of agents and their collective behaviour. Furthermore, formalisations of sonic swarming behaviours are made and put into context of the texton (Nyström, 2011), which results in an elaboration how individual sound particle behaviour affects the group within an acousmatic framework. As a practical conclusion, compositional strategies towards obtaining unified spectral, spatial and temporal relationships are discussed with respect to three of the author's works.

Acknowledgements

I would like to thank Bjarni Gunnarsson for the immense support in writing this thesis ...

Contents

Declaration of Authorship	iii
Abstract	v
Acknowledgements	vii
1 Agents and the Formation of Material	1
1.1 A brief introduction to ABM	1
1.2 Emergence and Complexity	2
1.3 Simulating Nature	4
1.4 ABM as a Medium	6
1.5 Agent as Material	8
1.6 Agent as Morphology	8
2 A Method for Designing Sonic Behaviours	11
2.1 Introduction to SonicSwarms	11
2.2 Scale	13
2.3 Unification	14
2.4 SonicSwarms	15
2.4.1 Agents	15
2.4.2 Environment	19
2.5 Behaviours as Spectro- and Spatiomorphologies	21
2.5.1 Behaviour and Contiguity	21
2.5.2 Individual Behaviours	23
2.5.3 Group Behaviours	24
2.6 the Agent vs. the Texton	25
3 Compositional Strategies	27
3.1 Mathematical Gestures	27
3.2 Sound Genetics	29
3.3 Sonic Ecosystems	30
4 Conclusion	33

Chapter 1

Agents and the Formation of Material

Sounds which construct our cultural and natural sonic environment have always been an important influence on the process of composing music. A permanent feedback is present between what we hear and what we compose, catalysed by inspiration and carried out by creativity. Most of the music we hear today is evolving through our cultural environment and classified through a tremendous amount of genres and styles. However, for some, diverging from cultural expectations often results in looking at information which is not already present within music at the time. Such information can exist in many forms, mediated through hearing, vision, touch or smell and then accumulates into an interpretation of some entity or concept, which is significant enough to incorporate into a piece of music. Personally not only sounds, but also their underlying structure inspires me. For example a group of flocking birds, filling space with a cloud of stochastic impulses, rapidly fluctuating in shape and density. Not only is this a spectacular phenomenon, it is a very important and inspiring visualisation of a complex structure. Similarly a swarm of bees diffusing filtered triangle-wave signals, whilst unpredictably traversing space. Together they form an imposing group and chaotic source of sound. Likewise, a colony of ants, millions of microscopic waves emerging from the tip off their legs and communicating to each other through scent, touch, body language and sound, together forming a highly efficient super organism. Such phenomena form the basis for an approach to creating sound which attempts to unify timbral and behavioural aspects of a collection of virtual sound generating objects. This thesis discusses a method focussing on realising such unification, which involves the use of autonomous agents existing within a virtual space, together forming an Agent-Based Model (ABM) (Macal and North, 2010).

1.1 A brief introduction to ABM

The purpose of this paragraph is not to give an extended theoretical outline of ABM. However, a basic understanding on both conceptual and technical level is needed in order to shape a clearer view on what musical value it might hold. A typical ABM has three elements (Macal and North, 2010):

1. A set of *agents*, their attributes and behaviours.
2. A set of agent *relationships* and methods of interaction.
3. The agents' *environment*.

Furthermore, an agent is defined by the following characteristics (Macal and North, 2010):

1. An agent is a *self-contained*, modular, and uniquely identifiable individual.
2. An agent is *autonomous* and self-directed.
3. An agent has a *state* that varies over time.
4. An agent is *social* having dynamic interactions with other agents that influence its behaviour.
5. An agent may be *adaptive*.
6. An agent may be *goal-directed*

Characteristic five and six are optional, dependent upon the phenomena that is being simulated by the model. This varying has to do with the difference between adaptive and non-adaptive complex systems (Mitchell, 2011), which is discussed in more detail later. Furthermore, analogous to our own modes of perception, an agent's perception is limited, meaning that only *local* information is available to an agent. Also, according to characteristic four, agents interact with other agents, that are called its *neighbours* and located within an agent's *neighbourhood*. This neighbourhood can change rapidly over time, possibly affecting an agents behaviour and consequently others surrounding it, also resulting into the varying state as mentioned in characteristic three. The way relationships between different agents are defined is generally termed as the *topology* or connectedness of an ABM, which also has implications for the kind of space these agents are located in. Generally this space is called the environment of a model, which can function differently depending upon the phenomena it aims to simulate. Consequently, the environment has implications for the kind of *movement* that is being exerted by the agents, which also affects the way we perceive their behaviour. This movement is the result of a discrete sequence of *decisions* that an agent has made, all dependent on the current neighbourhood of that particular agent. Note that the state of the entire model is defined by the collection of states of all agents, combined with the state of the environment they are located in.

Take for example an ABM simulating the movement of cars within a traffic-jam. A behaviour is simulated which is defined in terms of 'real' movement through physical (euclidian) space, meaning that the topology is constructed by distance measures and velocities as we are used to. Every vehicle has its own neighbourhood, containing all local information it can observe. Based upon this information, decisions are made regarding steering, braking and accelerating, resulting in a particular behaviour. In total, a collective behaviour is constructed by all individual behaviours and it is the relationship between these two behaviours that an ABM analyses by simulation. Thus, ABM's do not only function as tools to simulate particular group behaviours, but also contribute to the fields of science that try to encapsulate properties different behaviours might share, which are placed within notions of *emergence* and *complexity* (Mitchell 2011).

1.2 Emergence and Complexity

The term emergent was already coined in 1875, described as followed (Lews, 1875)

"Every resultant is either a sum or a difference of the co-operant forces; their sum, when their directions are the same – their difference, when their directions are contrary. Further, every resultant is clearly traceable in its components, because these are homogeneous and commensurable. It is otherwise with emergents, when, instead of adding measurable motion to measurable motion, or things of one kind to other individuals of their kind, there is a co-operation of things of unlike kinds. The emergent is unlike its components insofar as these are incommensurable, and it cannot be reduced to their sum or their difference."

Based upon this statement, it could also be said that when a collection of communicating objects form unforeseen structure that is not already present within individual objects themselves, they exert emergent behaviour. We call this collection combined with imposed interactions upon its parts a *complex system*. More formally, the following properties are proposed for the definition of a complex system (Mitchell, 2011):

- Complex collective behaviour
- Signalling and information processing
- Adaptation

Even though such characteristics are made to separate complex from non-complex systems, there still exists ambiguity in measuring these three properties. Any communicating collection of objects can exert collective behaviour, but when is this behaviour also complex? Are there boundaries to complexity? Author of the book *Two's Company, Three is Complexity* states the following: "*Complexity Science can be seen as the study of the phenomena which emerge from a collection of interacting objects*". (Johnson, 2007) This means that, according to this definition, we can interpret the first property as the fact that a complex system must exert emergent behaviour, meaning that structure that could not easily be foreseen shines through the systems totality, self-organised by its parts. This fact is reinforced by the following alternative definition of a complex system (Mitchell 2011): "*a system that exhibits nontrivial emergent and self-organising behaviours*".

The second property puts specific constraints on the objects which constitute the whole. Mitchell states that such objects must have a sophisticated way of processing either *internal* or *external* information, which could be compared to an internally 'felt' state and an externally 'observed' state. By this restriction, systems which are constructed out of particles that do not have the ability to process information, are not considered complex. In some sense this places a borderline between systems that can only change by force of nature, and systems that can also change by will. Out of this stream of thought follows the third property naturally. Adaptation is the result of an inherent will to survive within all living species, constructed by our conscious and unconscious decisions and coupled with mutations occurring within our DNA. Still, this does not mean that a system with no adaptive features cannot exert complex behaviour. Mitchell addresses this by making a distinction between *adaptive* and *non-adaptive* systems. Some examples of phenomena which are either adaptive or non-adaptive are given below.

1. Adaptive complex systems

- Immune system

- Brain
- Ant Colony
- Economy

2. non-adaptive complex systems

- Earthquakes
- Climate
- Solar System
- Flock

Upon connecting agent-based models, capable of simulating such complex systems, to the domain of acousmatic sound production, there is always a constructed mapping between the possible states of the model and the sonic output. Through this fact, connections can be drawn between the previously defined properties of complex systems and relevant musical notions within acousmatic sound production, namely space, communication and time respectively. It is possible that the way in which individual agents self-organise has direct implications for the way sound is being organised within a space, which could be either spatial or spectral. Then, different modes of communication between individual elements could result in various degrees of coherence within the constructed morphology. Finally, the complete sonic evolution of such a sonic entity or mass could be directed by adaptation. In order to arrive at these three notions, a specific agent based model is chosen as a basis for implementation. The following paragraph aims to introduce this model and discuss its possible applications in an acousmatic framework.

1.3 Simulating Nature

“A flock exhibits many contrasts. It is made up of discrete birds yet overall motion seems fluid; it is simple in concept yet is so visually complex, it seems randomly arrayed and yet is magnificently synchronised. Perhaps the most puzzling is the strong impression of intentional, centralised control. Yet all evidence indicates that flock motion must be merely the aggregate result of the actions of individual animals. ” (Reynolds, 1987)

Complex animal group behaviour is one of many nature’s delights. Flocks, herds and schools show a beauty which does not only lie within global motion, but also in local structure. There exists no centralised control within these structures, making them very interesting objects to study. However, due to their lack of centralised control, a visual simulation of those behaviours seemed very difficult, until the famous paper: Flocks, Herds, and Schools: A distributed Behavioural Model, by Craig W. Reynolds was published in 1987. In this paper a method is described to realise autonomous movement of such groups. Only by defining three simple rules, *separation*, *cohesion* and *alignment*, individual decisions based upon these rules result in fluid and complex motion across its totality. This also opened up the possibility of implementing this behaviour within an ABM, creating a system which could be used for a wide range of applications. Within the context of such a specific implementation, agents are considered to be *boids* (Reynolds, 1987), which is a shortened version of bird-oid object, but is also the name of the tool that Reynolds created to realise such flocking behaviours virtually.

Analogous to real birds in flight, the topology underlying this model is made up of euclidean space, which means that interactions between these boids are mainly dependent upon distance measures. The environment can be differently defined, depending on the kind of flock, herd or school that is being simulated. However, since this model considers movement through simulated physical space, this environment is bound to physical laws, including the boids themselves. With respect to this, the environment becomes a medium through which boids move, which can have effects on their speed, mobility and vision, all dependent upon implementation.




FIGURE 1.1: Reynolds' boids model in Processing, implemented by Daniel Shiffman (Shiffman, 2012)

Reynolds draws close connections between *particle systems*, which have been used to model phenomena such as fire, smoke and clouds, (Reeves, 1983) and the boid flock model. Both are models consisting of objects with their own internal state, which consists of color, opacity, location and velocity within an animated context (Reynolds, 1987). However, different from particles within a particle system, boids have geometric orientation (Reynolds, 1987). This means that the direction in which they move is influencing the movement thereafter. Secondly, not only do boids contain an internal state, they also have an external state (Reynolds, 1987), which implies that boids have the capability of observing their neighbourhood, possibly containing other boids.

The agents (i.e. boids), topology and environment, which constitute this described ABM, could serve as a framework for connecting boids' behaviour to acousmatic

sound. The rules which define the behaviour occurring inside the model are variable, and could be exploited for means of experimentation. With this experimental approach, such a method gives way to exploring sonic behaviours that result from slight variations upon either existing or newly created rules. From here a more experimental and even playful approach to modelling sonic behaviours is possible. By doing so, a model is created which is not meant to analyse or even simulate already existing behaviours, but potentially could be used to create structures that transcend their natural origins, conveying behaviours that are derived from physical and natural principles.

1.4 ABM as a Medium

“As emergent properties are by their very definition unpredictable in nature, the processes needed to generate complex behaviour are hard to determine. To negate this problem, composers working in this scenario often employ processes known to produce complex results, for example flocking or swarm algorithms. Such processes, however, are extra-musical, often being drawn from Artificial Intelligence, or engineering based research. There is no guarantee that such algorithms applied to music are going to make musical or aesthetical sense.” (Davis, 2010)

Even though until now ABM's have been approached and described as a tool, its principles could be extended to describe a more general way of expressing art. When implementing such a model, it is almost impossible to remove the inherent particle based aspects it brings along. That is why considerations regarding aesthetical and conceptual values in relation with the work should be made. As an example, consider the installation *Cross Pollination* by Tom Davis.

“It aims to guide the ear of the listener to the emergent musical structures that arise from interactions between the installation's constituent parts. Its design and construction are guided by a study of Complexity Theory. . . . In order to create a society of interacting agents from which emergent structures can arise, this installation is constructed out of twenty physically distinct agents. Each agent is constructed from a one metre balloon, four metres of piano-wire, a piezo transducer, a three-Volt motor, and a motor controller.” (Davis, 2010)

He continues by explaining that, when activated, an individual balloon (or pollen) will start off a chain of feedback, possibly leaking out to neighbouring balloons as well. Altogether, the whole system can be disrupted and influenced by people interacting with individual or multiple balloons.



FIGURE 1.2: Cross Pollination by Tom Davis

Even though this example does not contain any computational elements, it still conveys the same principles that an ABM has. Balloons have a locally observed neighbourhood, interact with neighbouring balloons and are contained within an environment, which functions as a medium. It could be argued that even the observer interacting with the system is in fact an agent, contributing as a decision making entity to the complete sonic environment being created. Aspects of this example work very well with the thoughts described in *Music as a Gradual Process* (Reich, 1968):

“I am interested in perceptible processes. I want to be able to hear the process happening throughout the sounding music.”

“Though I may have the pleasure of discovering musical processes and composing the musical material to run through them, once the process is set up and loaded it runs by itself”

Still, within every process of creating art, there are decisions to be made regarding material or content. When a work of art has conceptually more relevance to structure than to particular sonic qualities, material that constitutes to the whole should fit effectively to the structure that is to be heard. Similarly described by Reich:

“Material may suggest what sort of process it should be run through (content suggests form), and processes may suggest what sort of material should be run through them (form suggests content). If the shoe fits, wear it.”

The latter statement suggests that form depends on choice of material and vice versa, but there is no comment about a possible hierarchy between the two. However, if sound is generated from emergent structure, temporal or spatial ordering of events is naturally more important than the substance distinct events are made of. Imagine integrating the behaviour of flocking birds into a musical context. Structure which

defines such behaviour is not so easily adjusted to have a better fit with material. Though, choice of material can be easily adjusted to have a better fit with the behavioural context it is placed in. Another reason for this hierarchical relationship results from the unpredictable nature of emergent behaviour. Structures that are not pre-composed form spontaneously and subsequently generate sonic material which cannot always be contained within the composer's aesthetical boundaries. Thus, not only choices of material are subject to form, but also decisions regarding particular mappings between the two.

1.5 Agent as Material

Placing the notion of material into a musical agent based context does not seem logical in some ways. A classical definition of material suggests a set of predefined sounds, prominently existing in either physical or virtual form. This means that processes which influence such material can only be applied after the material has been generated. However, digital and analog systems gave way to a class of material which is outside of time, defined by a set of values within particular sound generating structure. Even though material in this sense is not pre-defined in any way, for the process of composing sounds both types of material function similarly, namely as a space of sonic possibilities that is to be traversed through by means of structure, resulting into form. However, instead of classifying such a virtual parameter space as material, we call such a space an *instrument object*, which boundaries are defined by the sonic limitations of the object itself. From here, an important relationship can be drawn between two state-spaces, namely that of an agent and an instrument object. By considering these two spaces to be the same, an agent moving through its virtual environment would correspond to a point moving through a space of sonic possibilities created by the instrument object. This relationship addresses conceptually how the domain of an ABM could be connected to the domain of sound, in particular acousmatic sound. As a natural consequence, any behaviour that is being exerted by agents within the connected model is directly translated into sounds behaving in a similar manner. Furthermore, not only emergent behaviours resulting from sophisticated rules can be mapped into sound, but also simpler ones, which might have more general uses within the field of electronic sound production. A list of those sonic behaviours will be given, formalised and reflected upon in the following chapter.

1.6 Agent as Morphology

“Both material and form are composable and can be articulated. From a Vaggionian perspective, there is no difference of nature or ontology between the two terms of this duality, but a difference of (time) scale. This is why the duality could be put aside: one could then use a single word rather than two words, one word declinable according to the scale. This is the role that the concept of ‘morphology’ seems to play in Vaggione’s texts. “ (Solomos, 2005)

As a continuation on the discussion regarding the relationship between form and material, in the context of a digital system, Vaggione’s musical thoughts aid to unify the two. Digital systems allow for precise manipulation on temporal domains small enough such that causality between separate events cannot be heard, though it can

still be composed. Form seems to take another role once it's placed within smaller time-scales, affecting a spectral domain instead of temporal one. As the quote above describes: *"This is the role that the concept of 'morphology' seems to play in Vaggione's texts."*, and also within this thesis I seek to approach a morphology likewise. However such morphological thinking will be extended towards an agent based domain, coupled by a sonic one. This connection will be drawn by considering three things the morphological approach has to offer (Solomos, 2005):

"To begin with, it postulates that material (sound) is not neutral. In other words, its basic tenet is that minimum units ('blocks') do not exist to be assembled at will into some combinatory play that produces abstract 'forms' (i.e., totally autonomous in relation to the material) (Vaggione, 1999) "

In this regard, material connects very closely to agents within a model. Both have the capability to autonomously move within a network of possibilities, further developed by the process they are bounded to and the context they are situated in. As a result of this thinking, Solomos continues by stating:

"This leads to an approach different from the parametric one. If a sonic form must be analysed (the morphological approach is not necessarily holistic), we will speak of its characteristics, aspects or components, and not of its parameters. However, Vaggione does not object to parametric treatment. He postulates that the two approaches are complementary. (Budón, 2000)"

This approach, complementary to a parametric one, reinforces the notion of material being an object. If such an object were to be extended to an agent, its characteristics, aspects or components would be equivalent to an agent's behaviour or perhaps its internal state. Furthermore, not only does an agent have similar descriptions, it is able to build a bridge to a parametric domain, meaning that those two different approaches are not only supposedly complementary, but also unified into a single object. Finally, Solomos states:

"... the morphological approach allows sonic forms to be thought of as dynamic movements, as processes. In Vaggionian terminology, this approach is 'transformational', which implies that evolutions of characteristics or parts of a sonic morphology are envisaged with regards to their context (Vaggione, 1998)."

Again, this seems to apply well towards an agent based stance. A morphology is to be conceived within a particular context and further evolved over time, likewise an agent is spawned and evolved by its neighbourhood and environment. Building upon Solomos' three stances, it is possible to construct a similar morphological approach towards agents. The morphology of an agent would then be defined by its behaviour and state, both evolving over time and breaking similar conceptual boundaries to those of material and form, analogous to local and global behaviour. Furthermore, if such morphological thinking is applied to an individual agent's behaviour, it could be possible that adaptation directs the morphology that is being exerted by an individual agent's behaviour, shaping a form which is constructed out of characteristics from the material (i.e. agent) itself.

Though a morphological approach to an agent has been given, it has been put into

relationship with only temporal aspects. However, applying the same thinking towards agent based models with an euclidean topology, the possibility of spatial morphologies arise. Not only do steering behaviours manifest itself as events which are changing over time, they also form specific trajectories through space. Thus a multitude of agents also forms spatial morphologies, which is inherently connected to their behaviour. Vaggione also addresses the notion of space in relation to morphology Solomos2005:

Let us briefly examine the treatment of space as an example of the morphological approach. Like numerous other composers, Vaggione postulates that space is composable. However, unlike musicians coming from the serial tradition, he does not view it as a parameter. Space is part of the morphology of sound, and, if it does have relative autonomy, this is as a morphology 'which will modulate and be itself modulated by other morphologies' (Vaggione, 1998).

Once an instrument object has been connected to a collection of agents, modulating its parameters, each agent radiates their own sonic morphology at their own location in space. Together, they form a network which is connected spatially, temporally and therefore spectrally, constructed through modulation which is dependent upon their locally observed neighbourhood. How this network is created will be discussed in more detail in the following chapter, coupled with a thorough description of the agent based model that attempts to do so. After this, relationships regarding spatial and spectral morphologies between sonic objects within this model are evaluated within a framework of *spectromorphology*, *spatiomorphology* (Smalley, 1997) and also *textons* (Nyström, 2011) which in particular supports an evaluation of the relationship between local and global sonic behaviours.

Chapter 2

A Method for Designing Sonic Behaviours

This chapter aims to provide an explicit method of generating acousmatic sound material by controlling an ABM with euclidean topology, which is in fact very similar to an earlier described model, namely boids (Reynolds, 1987). This involves describing a virtually defined instrument object and how it is connected to the internal and influenced by the external state of an individual boid. Furthermore behaviours of those boids will be formalised within a spectromorphological framework, therefore analysing the spatial textures they are capable of forming (Smalley, 1997). The *texton* (Nyström, 2011), a conceptual acousmatic particle that constitutes a *spatial-texture*, will be compared to an individual agent and aids for the description of particular emergent textural states (Nyström, 2011) that multiple agents can form, depending on their behaviour.

2.1 Introduction to SonicSwarms

Sonic-Swarms is custom-build software which makes precise design of real-time, interactive, particle based sound environments possible, resulting in digitally generated multi-channel audio material, and is implemented in Processing (Shiffman, 2012) and SuperCollider (Wilson, Cottle, and Collins, 2011). The model which has been programmed within Processing has many similarities to Reynold's boids model described in 1.3. However, many additions and changes have been made in order to enhance its purpose, which is generating acousmatic sound material. One of these additions is an intuitive visual interface, that consists of multiple controls for the behaviour of agents contained within the model and also direct controls for synthesis parameters within SuperCollider. Furthermore, in the centre of this interface lies a circular space in which the agents are moving, thus functioning as visual feedback regarding positioning of sound sources and relationships between them. These relationships can be precisely designed by a third aspect of this interface, which is called the *interaction-matrix*. This aspect allows for a precise mapping between an agent's state and its corresponding instrument object.

Open Sound Control (Wright, 2005) mediates all information regarding the agents, environment and interface from Processing to SuperCollider. Through a variety of objects and functions, information regarding agents' and their environment is being mapped and routed to the arguments of a SynthDef (Wilson, Cottle, and Collins, 2011). An agent moving through its environment then corresponds to an instance of a SynthDef moving through its bounded parameter space, at each point in time outputting a either continuous stream- or discrete portion of sound at a specific location.

This means that even though the implemented ABM is not internally adjustable, the SynthDef connected to individual agents is, which implies that in fact SonicSwarms only modulates parameters of pre-defined instrument objects and does not generate waveforms internally. In total, a completely synthetic and spatialised sonification of a behaving group of agents can be designed in real-time, which may be a useful tool for the process of composing electroacoustic music.

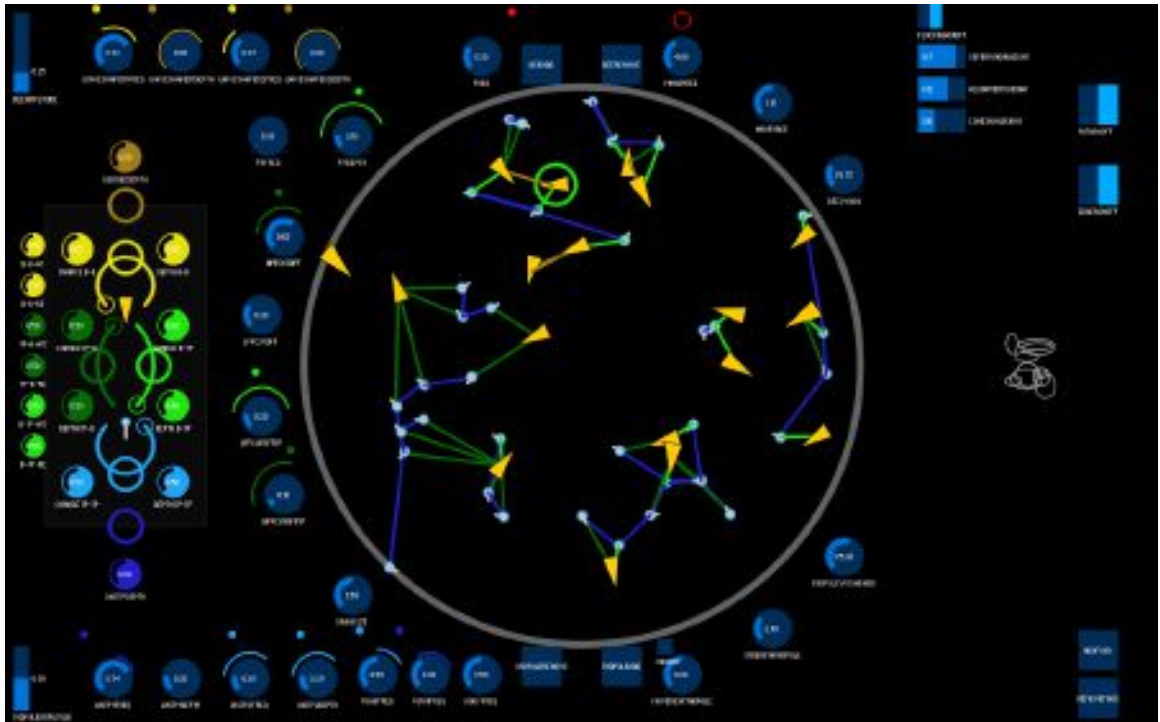


FIGURE 2.1: GUI of SonicSwarms

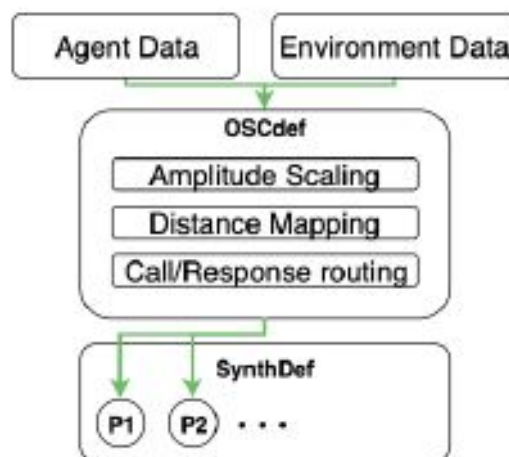


FIGURE 2.2: SuperCollider flow of data

Before continuing with further descriptions of the model and musical possibilities that lie within, the following two paragraphs aim to reflect upon the temporal scope of this tool, but also to discuss important resultants from unification among different time scales.

2.2 Scale

The book *Microsound* (Roads, 2004) starts off by describing 9 different time scales of music, including the following, ordered from large to small:

- *Supra*, A time scale beyond that of an individual composition and extending into months, years, decades, and centuries.
- *Macro*, A time scale of overall musical architecture or form, measured in minutes or hours, or in extreme cases, days.
- *Meso*, Divisions of form. Groupings of sound objects into hierarchies of phrase structures of various sizes, measured in minutes or seconds.
- *Sound object*, A basic unit of musical structure, generalizing the traditional concept of note to include complex and mutating sound events on a time scale ranging from a fraction of a second to several seconds.
- *Micro*, Sound particles on a time scale that extends down to the threshold of auditory perception (measured in milliseconds).

Roads also states the following:

“A central task of composition has always been the management of the interaction amongst structures on different time scales.”

In my opinion, not only composing is subject to this view, but also listening. Either conscious- or unconsciously we process changes happening within a composition differently, dependent upon the scale these changes are happening in. It is how the composer manages to create inter-dependencies between different time scales that affects the overall perception of the music. Furthermore, with these scale definitions in mind, it opens up a possibility of defining a temporal scope of a sound generating tool.

With the help of digital systems, it became possible to implement structure upon the micro scale, possibly echoing over other timescales as well. Creating relationships between sound-objects within a space regarding their location and movement, would fall within the same timescale category as arranging a collection of notes into a melody, which would be the meso-scale. Here I draw an analogy between the modulation of synthesis parameters of multiple sound objects within a space, accumulating into an *acousmatic image* (Smalley, 1997), and the potential harmonic progression of a melody. However, one could argue that it is not justified to make such a connection, since the definition of a note, or sound object, within the context of electronically produced sounds, is not so well defined when compared to the traditional note. But sound objects have been widely discussed, first coined by Pierre Schaeffer. To him, a sound object was a sound whose origin a listener could not identify (Schaeffer, 1977). Roads redefines this by stating: “Any sound within stipulated temporal limits is a sound object”. Still, a notion of micro structure within the sound object itself misses within these definitions. The ‘Vaggionian’ object, which has been discussed earlier as a morphology, is being put into comparison with the Schaefferian object as followed (Solomos, 2005):

“We know that in the field of electroacoustics it refers to the tradition of *musique concrete* as conceptualised by Schaeffer. Yet Vaggione’s approach is quite different for two reasons: Schaefferian ‘sound objects’ are

only located in the domain of macro-time, while Vaggionian objects can be found in any time scale; furthermore, the first are 'opaque' since they are produced on a magnetic-analogue tape, whereas the second, being digital, are always composable. (Vaggione, 1998)"

The reason for bringing up the discussion of time scales lies within the fact that a fundamental aspect of SonicSwarms can be described in this perspective. In my opinion, managing interactions between different time scales in music is not only a central task, but should also result into a sense of unification, especially in the context of electro-acoustic music. In other words, this means that any musical relationship defined upon one of those scales, echoes over different temporal formats. An ABM is the perfect gateway to achieving such interdependencies among time-scales, since individual behaviour of agents within the model, affects the collective behaviour of the group, which is bound to a different temporal domain. Connecting then such a model to sound, results in a virtual tool acting as a catalyst for the production of convincing electro-acoustic sound material. The word convincing has a very subjective implication, but the following paragraph will explain what is meant by this.

2.3 Unification

In the context of composing electro-acoustic music, there are two reasons I consider the idea of unification among different time scales as a central property to SonicSwarms, which have to do with instinct and resemblance, relating to communicative and practical notions respectively. In order to elaborate these two reasons further, I extend upon the ideas of Denis Smalley related to the transmodal perception of acousmatic music. He states the following:

"Although acousmatic music may be received via a single sensory mode, this does not mean that the other senses lie dormant; in fact they spill over into sonic experience. Our sense of texture is learned through vision and touch as well as sound; our experience of the physical act of sound making involves both touch and proprioception; spectral motion, and the movement and distribution of sounds in space relate to our own experience of physical motion and cultural and natural environments." (Smalley, 2007)

Even though this does not only apply to acousmatic music, being able to compose both the spatial- and spectral aspects of sound simultaneously allows a precise way of creating and arranging sound objects upon the meso-scale, which potentially could accurately resemble the behaviour of sounds surrounding us in our cultural and natural environment. It is exactly these kind of artificial sounds that will spill over into senses other than our hearing, potentially resulting in a *quasi-visual* listening experience (Smalley, 2007). Here the latter reason for unification among time scales, namely resemblance, comes into play. Most sound producing structures surrounding us inherently have a coupling between changes happening upon different temporal scales. Language would be the most prominent example of this. Relationships between paragraphs (macro scale), defines how sentences are constructed and arranged (meso scale), which have direct consequences for the words within (sound objects) and consequently influencing the produced vowels (micro scale). This applies as well to moving sound sources within our natural environment. For

example consider a fruit fly as a sound source, whose motion of wings control its velocity that consequently creates a trajectory in space upon a macro scale, modulates its amplitude within the meso scale, but shaping the internal waveform as well, thus manipulating the micro scale. Embedding properties of such sonic structures, present in our cultural and natural environment, into decisions related to composing electronic sound, is what I mean by creating convincing electro-acoustic sound material. Thereby striving towards a closer resemblance of the behaviour of sounds that construct our acoustic environment.

2.4 SonicSwarms

The following paragraphs aim to describe SonicSwarms in detail and evaluating its musical value within a spatio- and spectromorphological framework, accompanied by sound examples.

2.4.1 Agents

We have seen in 1.1 that there are three main elements which define an ABM, including a collection of agents. As mentioned in the introduction of this thesis, inheriting principles of complex animal group behaviours into the creation of acousmatic sound material is a focal point within this research. That is why agents that constitute SonicSwarms are designed similarly to those that constitute the boid model, which is described in 1.3. However, different from boids, within each agent in SonicSwarms there exists a duality, spreading across sonic and visual domains. Furthermore, by choice, there are two complementary classes of agents implemented within the model, namely discrete and continuous agents. Continuous agents are in fact completely similar to boids, resembling the smooth movement of birds in flight. On the other hand, discrete agents move in pouncing manner, coupled with a discrete temporal way of processing information. The reason for this class division can be found within a fundamentally different approach to synthesizing sound, which is exemplified later.

Putting this duality and distinction, between domains and classes respectively, into relationship with the principal characteristics and properties of agents, as described in 1.1, shines light upon important differences between those two classes and domains.

“An agent is a self-contained, modular, and uniquely identifiable individual.”

Both classes (or kinds) of agents within Processing are visually defined as distinguishable objects within a circular space, as shown in the figure below. Within the sonic domain, each agent is connected to an instrument object (i.e. synthesis structure), thus represented as an individual, though not always distinguishable, sound source. To be more precise, this synthesis structure is defined through a SynthDef within Supercollider, where its parameter boundaries are predefined. Continuous agents are commonly represented by synthesis structures that generate continuous sound material, inheriting principles of *additive synthesis* (Roads, 1996) when formed as a group, whereas discrete agents are represented by single grains of sound, emitted by every movement these agents make. Thus complex control over any *granular synthesis* structure (Roads, 1996) is possible by behavioural control over discrete agents.




FIGURE 2.3: Continuous agents as neighbours




FIGURE 2.4: Discrete agents as neighbours

“An agent is autonomous and self-directed.”

Similar to the boids model, agents within SonicSwarms have the ability to steer autonomously. However, there is the possibility to morph and switch between different types of autonomous steering, resulting in different spatial and spectral morphologies. More specifically, this means that sound sources are autonomously moving through the virtual acoustic space created by the loudspeaker configuration, which means they are able to collectively form emergent spatiomorphologies that are directly connected to individual agents’ spectromorphologies, thus resulting in emergent spectral behaviour as well. Note that only continuous agents have this variety in steering behaviours, due to the nature of their implementation. Discrete agents are purely an additional element which allows for the design of granular clouds, which potentially can interact with continuous sources as well. The following behaviours are implemented for continuous agents, which will be sonically dissected in paragraph 2.5.2:

- *Wandering* (Shiffman, 2012)
- *Path following* (Shiffman, 2012)
- *Seek attractor*
- *Attraction-Repulsion*
- *Forcefield-Velocity*
- *Flocking* (Shiffman, 2012)

“An agent has a state that varies over time.”

Both continuous and discrete agents have two properties that define its internal state, namely a *location* and *velocity*, which again applies equivalently to both sonic and visual domains. The external state is defined by an agent’s neighbourhood, bounded by its vision and possibly constituted by other neighbouring agents, objects and forces. Combined, the internal and external state of an agent influence its behaviour and vice versa. It is clear that an agent’s state varies over time as they move through their virtual environment, resulting in a spectromorphology which is inherently defined by an agent’s changing external state, but also a spatiomorphology influenced by an agent’s internal state.

“An agent is social having dynamic interactions with other agents that influence its behaviour.”

By the previous property, agents have the ability to observe and therefore influence each other. The euclidean topology allows for *distance* measures between agents, which serves as the main ingredient in implementing steering behaviours. Not only affects this distance measure their steering, it can also be mapped onto a specific set of parameters. Sonically this results in each sound source uniquely being modulated by the distance measure of its corresponding agent. Furthermore, trivial *call-response* principles are implemented and can be mapped similarly. Note that the model does not aim to simulate the full social complexity that certain animal group behaviours can exert, but only inherits the most basic components of such systems. Each agent has the ability to call out to their own kind (i.e. class), but also to respond to a calls from another kind. This means that four types of active interactions are possible between agents, that are sonically represented as envelopes, and can be designed and mapped within the interface in Processing. Those mappings are made possible by the interaction matrix, shown in the following figure.




FIGURE 2.5: All possible interactions: Continuous->Continuous (Brown), Continuous->Discrete (Light Green), Discrete->Continuous (Dark Green), Discrete->Discrete (Dark Blue)



FIGURE 2.6: Interaction Matrix. Including controls for attack/release, density of calls, depth of modulation

“An agent may be adaptive.”

Recalling back to the conclusion of 1.2, allowing agents to adapt would have direct implications regarding the evolution of the spatial and spectral morphology they create, which occurs upon the macroscopic time scale. By choice this is not implemented, due to the technical difficulty that such implementation brings along. This means that morphologies constructed upon the macro scale are resulting from user defined changes and not from agents adapting to each other and their environment. Hereby performative aspects start to arise as well, even though the tool is not designed to do so.

“An agent may be goal-directed”

Even though adaptation implies that agents must have a goal, which is in most natural related cases staying alive, the converse is not true. Implementing a goal can give the behaviour of agents direction and consequently the morphology they form upon a macroscopic scale, with much less technical difficulty. Here an important difference between discrete and continuous agents arises, which finds its roots within the type of synthesis that is applied to each of those classes. Similar to the comparison between boids and birds, discrete agents can be compared to growable worms, since they have the ability to grow by eating food which can be added into the system. In addition to the location and velocity vectors that make up a discrete agent’s internal state, they now also have a body size, which potentially could be used to evolve elongation of individual grains produced by the granular synthesis structure in SuperCollider.

2.4.2 Environment

Besides discussing the agents and their capabilities, it is important to examine their environment as well. The topology underlying the model consists of a two-dimensional euclidean space, meaning that distance measures between agents function as the basis for their steering behaviour. On top of that, physical laws are implemented in order to construct a medium through which agents are moving, incorporating friction into the model. However, this medium is not audible, meaning that the friction which an agent produces cannot be heard, unlike real physical movement through a medium. Different techniques are used to mimic an audible medium, which involves layering acousmatic material with a transformed version of itself. This is discussed in more detail in chapter 3.

Even though this paragraph does not aim to resolve specific technical problems, that arose while programming the model, one of them is very important to address. This has to do with unifying the fluidity of movement within both sonic and visual domains. As a continuous agent traverses its two dimensional euclidian space, a continuous sound source moves through the virtual acoustic space created by the loudspeakers, such that its location in relation to the centre is similar in both domains.

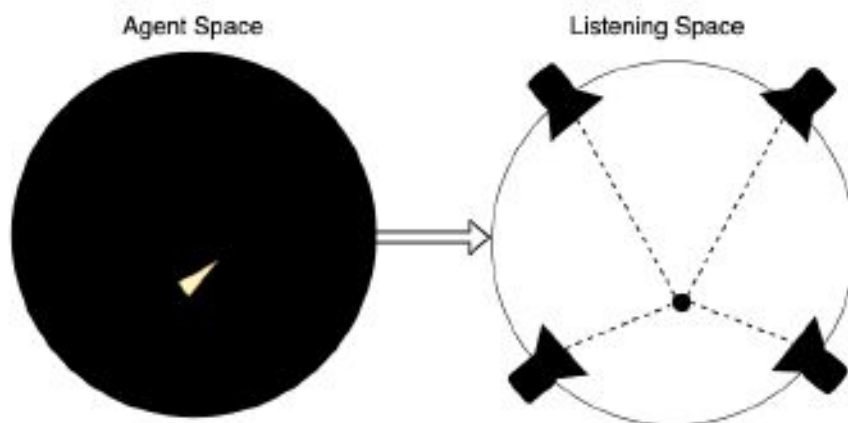


FIGURE 2.7: Agent Space to Listening Space

Visually, it is no problem to implement the agents' environment as an infinitely large space, meaning that the boundaries of the visible disc are not treated as physical walls. This is easily implemented by wrapping around an agent to the other side of the disc once it has crossed an edge. However, sonically this would mean that a sound source must be panned from one location to another instantly, resulting into discontinuous amplitude changes that are heard as clicks and pops. This problem is fixed with a technique called *acousmatic wrapping*. By placing a two-dimensional amplitude envelope, which is shaped as a Gaussian function, over the environment, agents closer to the edge are softer in amplitude than the ones near the centre. By doing so, it could be imagined that the agent's environment is not a disc anymore, but a sphere, where the listener is located on top and agents traverse its surface.




FIGURE 2.8: Two dimensional Gaussian function

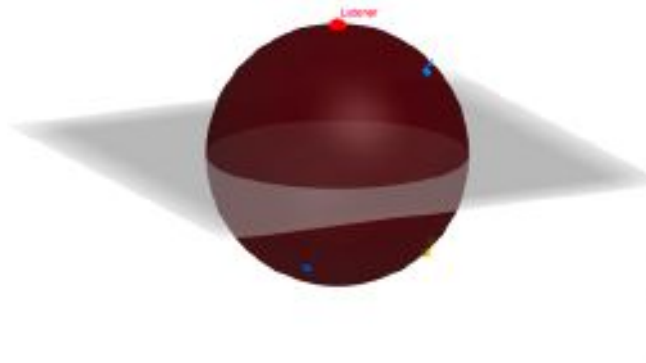


FIGURE 2.9: Spherical representation of the environment with listener on top and agents A,B and C.

In this regard, the boundaries of the visual disc in Processing correspond to the lowest point on the sphere's surface, exactly where the amplitude of all sources are zero. By this solution, an individual agent moving over this sphere's surface results in a continuously moving sound source without discontinuities.

It seems that the problem of continuity now has been solved completely, however this is not the case. Remember that agents have the ability to observe their surrounding locally, meaning that they can 'see' other agents approaching. Once an agent has been placed near the edge of the disc such that its circle of vision overlaps the boundary, it must be able to observe what is on the opposite side, meaning that a non symmetrical subset of the disc must be checked. This is solved by constructing a clone of each agent present within the system, depicted in the following figure.

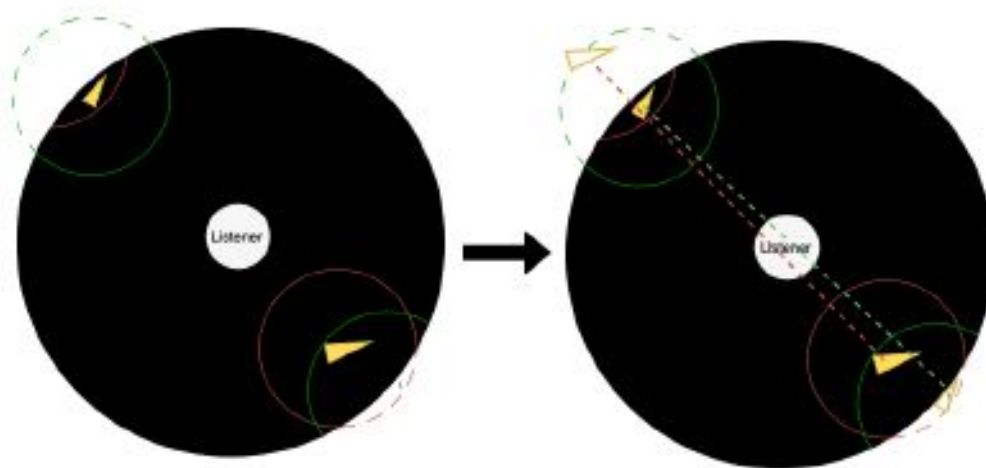


FIGURE 2.10: Solution to the boundary problem

As an extra functionality, there is the possibility of placing an algebraically defined force-field over the environment. Not only can this influence an agent's steering behaviour, it can also be directly mapped onto synthesis parameters, thus creating a direct relationship between an agent and its location in space. It is explained in paragraph 2.5.2 how this can be put into practice.

2.5 Behaviours as Spectro- and Spatiomorphologies

Until now, both agents and their environment have been discussed, as well as a brief examination of possible relationships between them. This paragraph aims to give a more thorough description of the possible sonic relationships that can be constructed by agents as a group, also putting it into the spatiomorphological framework proposed by Denis Smalley, and compare individual agents to textons, constituting spatial texture.

2.5.1 Behaviour and Contiguity

What has not been discussed yet is the influence of agent density on the perceived sonic environment, directly related to the perceived *spatial fill* and *spatial focus* (Smalley, 1997). Within this context, we will consider only continuous agents, due to their ability to form trajectories and occupy space more substantially, which means that their density has a greater sonic influence than discrete agents. SonicSwarms automatically scales down amplitudes of all agents once more agents are added to the system, such that the perceived summation of their amplitudes remains constant. This allows for a smooth morphing between dense textures and isolated trajectories, which could be compared to a morphing between *contiguous* and *non-contiguous* space. Smalley describes those notions as followed (Smalley, 1997):

“Spatial texture is concerned with how the spatial perspective is revealed through time. This is a question of contiguity. Space is contiguous when revealed, for example, in continuous motion through space (such as in a left-right gestural sweep), or when a spectromorphology occupies a spread setting (without spatial gaps). Non-contiguous space is revealed when spectromorphologies are presented in different spatial locations

such that two successive events are not considered near neighbours: there is no sense that a spectromorphology occupies or moves through adjoining sectors of space.”

As more agents are added to the system, the ear starts to attend towards different non-contiguous positions, even stronger if individual modulations are occurring, caused by agents approaching and calling to each other, overall creating a non-contiguous space. On the contrary, these interactive properties also unify the spatial texture that is created, forming a coherent whole once listened upon another structural level. Smalley points this out as well, stating that contiguity can differ depending on the structural level that is being considered. Furthermore, not only density of agents within the system has an important influence upon the perception of contiguity, but also movement. The behaviours described in 2.4.1 can be subdivided into two categories, namely individual and group behaviours. Each category has a different implication for the perceived spatial and spectral coherence of the generated texture, affecting also the contiguity upon different structural levels. Terminology for describing contiguous or non-contiguous properties is given below, aiding for a musical connection between agent behaviours and the spatial texture they form.

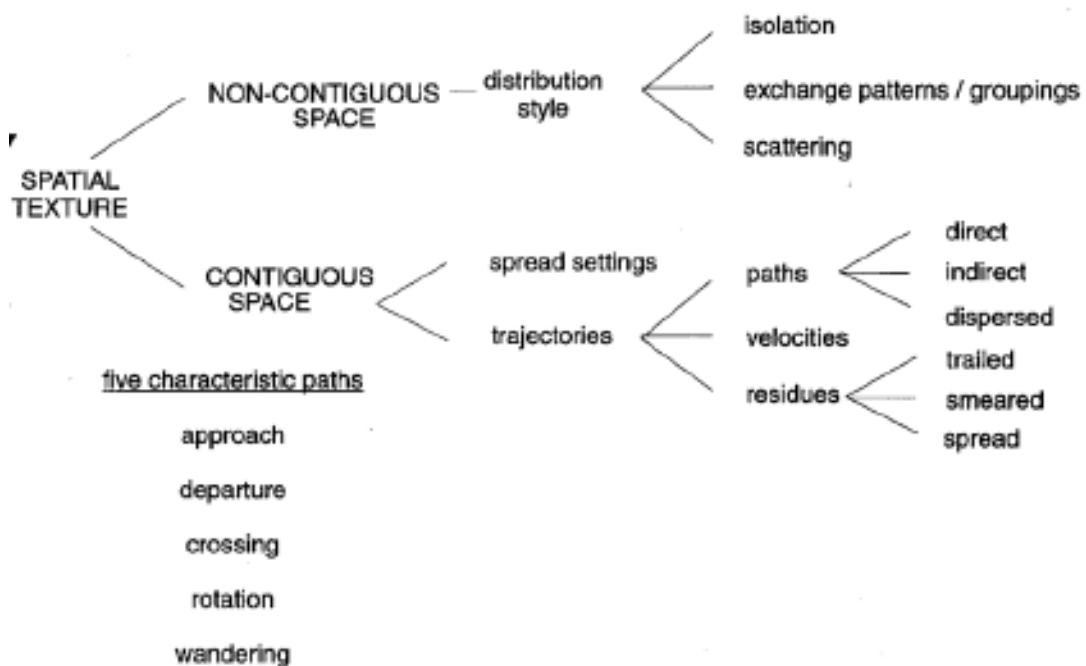


FIGURE 2.11: Terminology for describing spatial texture Smalley, 1997

The following paragraphs are analysing each behaviour of continuous agents sonically, included with sound examples. The implemented synthesis structure uses a granular FM sine wave generator, which has multiple parameters that allow for a wide timbral scope. For clarification, each agent is an individual oscillator with its own setting and spatial location. Furthermore, for practical reasons, all sound examples are in stereo format, meaning that musical value regarding multi-channel spatiomorphologies is hereby lost. This results in a lack of spatial transparency as well, meaning that the borderline between perceived texture and isolated sources is shifted in comparison to multichannel formats.

2.5.2 Individual Behaviours

“Wandering is a type of random steering which has some long term order: the steering direction on one frame is related to the steering direction on the next frame. This produces more interesting motion than, for example, simply generating a random steering direction each frame.”

(Reynolds, 2004)




FIGURE 2.12: Wandering mechanics (Shiffman, 2012)

An individual agent *wandering* through its environment results in a smooth brownian movement of a sound source, which is given as a characteristic path within Smalley’s terminology of spatial texture. Throughout sound example 1.1.1, an agent is slowly wandering through its environment, which this does not affect any spectral aspects of the source yet, since there are no relationships between its velocity and synthesis parameters constructed. Furthermore, no force field is placed over the environment and no other agents are present within the system. The second example is a collection of six agents, resulting in an amalgam of identical sound sources, distributed randomly through space, but together in similar random motion. This is heard in example 1.2.1. We can increase the spatial perception of example 1.1.1 by constructing a relationship between an agents velocity and a synthesis parameter, such that changes in velocity can be heard. Furthermore, a force field is placed over the environment, such that agents are receiving a force, i.e. a float value between zero and one, at each location in space. This force is then mapped to synthesis parameters, resulting in a continuous modulation inherently dependent upon a changing location in space. An individual agent, with velocity and force field mappings, is heard in example 1.1.2. The same is done for a group of agents and heard in example 1.2.2. All of these examples could be considered as contiguous spatial textures or movements. However, by constructing a distance mapping the contiguity can be broken up by modulations.

Since the manner in which an agent wanders is not affected by other agents, a similar random modulation can be obtained by constructing a distance relationship between agents, via the interaction-matrix. By doing so, new spectromorphologies are constructed, only heard when neighbouring agents are *approaching* and *leaving* each other. In example 1.3.1, agents with distal relationships are added one by one into the system, together approaching a non-contiguous space. Localisation of individual sources is enhanced again with the use of a forcefield. Throughout this example there are two main audible elements, namely agents with- and without a neighbour. We can remove agents with an empty neighbourhood from the acousmatic image by

coupling amplitude to distance. Consequently approaching agents are also increasing in amplitude, and agents without a neighbour are silent. This is heard in sound example 1.3.2.

A complementary behaviour to wandering is *path-following*, which gives way to designing periodic spatial movements. Since agents are not affecting each other's movement, it is classified as an individual behaviour. Furthermore, by placing a force field over the environment and coupling this to parameters, a periodic modulation can be achieved as well. Unfortunately within the stereo format a circular path cannot be heard, though this is the shape which agents follow within sound example 2.1.1.

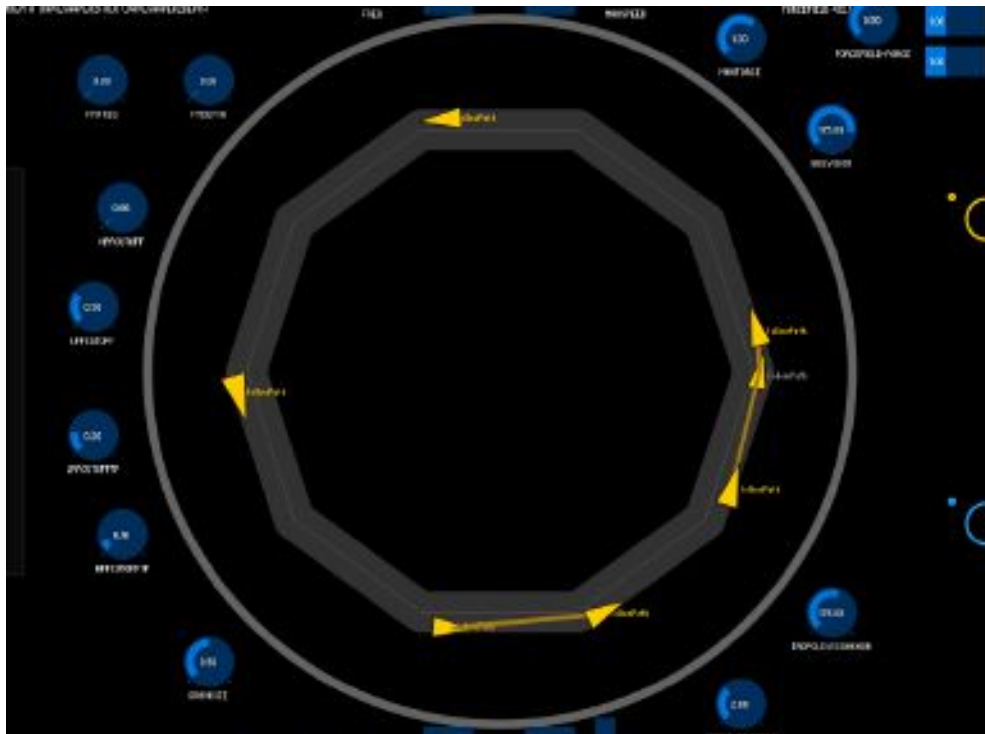


FIGURE 2.13: Terminology for describing spatial texture (Smalley, 1997)

2.5.3 Group Behaviours

Another implemented behaviour is called *attraction-repulsion*, which is in fact categorised as a group behaviour, due to the influence that neighbouring agents have on each other. This behaviour controls how much agents are attracted to- or repulsive from each other, by creating a relationship between an agent's speed and distance to neighbour. Purely spatial aspects of this behaviour cannot be heard easily, even less the more agents are present. However, by applying a distance mapping to synthesis parameters, elasticity of the resulting spectromorphology can be controlled. Such a sonic behaviour could be compared to one of Nyström's described *emergent textural states*, namely an *entitative state*. He describes this as followed: "*entitative states give the impression of living forms mapping out texture as they travel across spatial territories.*" (Nyström, 2011). Sound example 4.1.1, which contains agents that are attracted to each other, has a similar setup as example 1.2.2, only behaviourally they are different. Note that because of the attractiveness of individual agents, they tend to stick

together longer, resulting in a temporal elongation of modulations. Repulsive behaviour is audible in example 4.1.2.

Even though attraction-repulsion is a behaviour which is dependent upon an agent's external state, thus influencing the behaviour of a group, it only affects aspects regarding velocity. However, flocking is a behaviour that also influences direction and has a greater impact upon the created spatiomorphology, but also the occupancy of each agent's neighbourhood over time. Within example 5.1.1, a transition is heard from a randomly distributed wandering group of agents, to a united flock, slowly displacing themselves over time. Three simple rules constitute this behaviour, which are: *cohesion*, *separation* and *alignment* (Reynolds, 1987). Once a flock is formed, it has immediate consequences regarding the contiguity of space, due to distances that align and spatial locations which cohere. Sound example 5.1.2 is a collection of behaviours arranged into a small etude. All of the behaviours that are mentioned above are incorporated and layered to morph between different synthetic ecosystems.

2.6 the Agent vs. the Texton

Even though continuous agents and their behaviours have been placed within another framework regarding contiguity of space, we haven't focussed yet on musical properties that lie within agents themselves. Concerning a group of agents forming spatial texture, the idea of a texton fits well as comparison to the individuals. The following description of a texton is given (Nyström, 2011):

"Where spatial texture is concerned, however, we need to be able to describe elements in both spatial and temporal terms. Here, the textons would be a phenomenon of spatiotemporal granularity with locations, spectral distributions and aggregations in time."

Spatiotemporal granularity seems to fit perfectly within an agent based context, in particular those with euclidean topology. Each agent is coupled to a unique location, spectral state and behaviour, analogous to the description above. Almost every aspect of a texton can be compared to an agent, though there also exist differences regarding their capabilities of processing information, further reinforced by the following statement (Nyström, 2011):

"... textons are like small islands in space and time, clearly distinguished or diffusely radiating as perceived elemental morphologies. Distributed in aggregates, they are perceived as fleeting microshapes with locality in perspectival space and vertical elevation in spectral space. As textons propagate in time, they animate the acousmatic image and feed music with new space."

This description brings complementary aspects within the relationship between textons and agents to light. Textons are sonic elements that only radiate information, which is not influenced by surrounding textons. On the other hand, agents are decision making individuals, whose states are consistently dependent upon other agents within their neighbourhood. In other words, spatial texture constituted by textons is equivalent to the sum of its parts, while spatial texture constituted by agents is more than the sum of its parts. By this statement, I am not implying that one of those

two methods is superior to the other, but rather reflecting on exactly what is being heard in both cases. As some of the earlier sound examples show, spatial textures formed by agents are mostly constituted by audible relationships between agents, rather than the agents themselves. That is what's meant by 'more than its parts', implying that a spatial texture constituted by sonic agents is the summation of them, combined with their constructed relationships.

Chapter 3

Compositional Strategies

Within the final chapter of thesis, there will be three of the author's works discussed, composed within the first, second and third year of his studies at Sonology. As this thesis proposes a method for unifying morphologies upon different time scales, the following works attempt to do the same. It has been discussed that unification created by SonicSwarms does not extend towards the macro scale, that is why choices of the composer must be embedded into the form itself. Hereby compositional strategies are put into use, resulting in an overall form that is shaped accordingly to the structure within the material itself. Clearly, the compositional strategies discussed in the following pieces are work specific, thus there is no implication that those discussed strategies are appropriate approaches in general. However, they do tend to generalise the author's compositional approach to creating a macro form out of material, which is already full of structure up until the meso-scale. In every case this involves creating temporal relationships between long sections of material, attempting to enhance the overall listening experience, but also to create meaningful relationships towards the original method of synthesis itself.

3.1 Mathematical Gestures

Mathematical Gestures is a fixed media composition, composed for a quadraphonic loudspeaker setup, though sound example 6.1.1 contains a stereo mix of the piece. All of the materials are produced algorithmically within SuperCollider, though temporally ordered within Ableton (Manning, 2013). The algorithm that has been used for the creation of material was highly inspired by the opening section of Xenakis' work *Metastasis*, in which a mass of glissandi is played by large string section. Each glissando has its own trajectory within the frequency domain, together forming a mass of sound with visual and audible curvature even though constituted by linear segments. It was the author's goal to implement a program within SuperCollider that could transcribe any circular or quadratic function into a mass of glissandi, by drawing tangents on the specified function at regular time intervals. Not only would each tangent have its own spectral trajectory, but also a spatial trajectory. Furthermore, each intersection between tangent lines would be audible as an increase in energy upon the frequency at which they intersect, allowing for an exploration of harmonic relationships as well. An example of a function converted into acousmatic sound is given in the following figure.

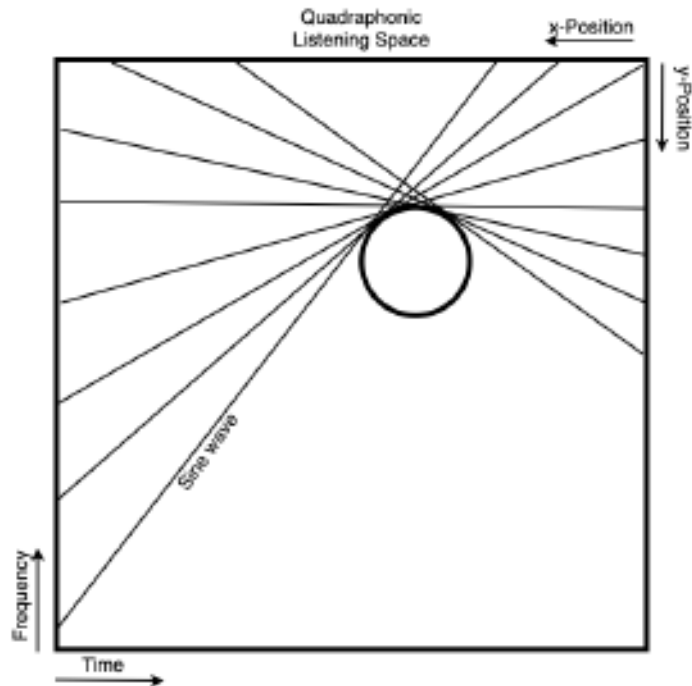


FIGURE 3.1: Process of translating a circle into glissandi and their trajectories in 2-dimensional space, created by the loudspeaker configuration

By specifying the function, its frequency boundaries, offset, amount of tangents, duration ect, it was possible to generate materials that functioned as either slow or long gestures of sound masses constituted by glissandi. Each sound fragment had a specific function coupled to it, meaning that there was an inherent connection to the visual domain. Because of this fact, it became easier to come up with interesting visual forms, than audible forms. By combining different circles and quadratic functions, the following five sections were constructed.




FIGURE 3.2: Mathematical Gestures Form

This composition is an example of a pure bottom-up approach. Placing no attention on any choices regarding form up until the moment that material has been generated. This same approach applies to the following pieces, though a stronger narrative approach will be used to place the applied synthesis technique in a meaningful framework.

3.2 Sound Genetics

Though complex systems have been discussed, this thesis did not put too much focus on adaptive aspects, besides discussing some temporal consequences it might bring along. Sound Genetics is a four channel, fixed-media composition that completely revolves around sonifying adaptation, or more precisely evolution, composed in 2017. With the help of an evolutionary algorithm, or more specifically a genetic algorithm (Mitchell, 2011), pre-recorded sounds are reconstructed out of noise. This algorithm is implemented in SuperCollider, though Spear (Klingbeil, 2009) is used to perform a partial analysis, functioning as the core of information upon which the genetic algorithm is applied.

A genetic algorithm aims to find an optimal candidate within a specific context. It does so by using the principles of natural selection, an important factor for adaptation among species. These principles are incorporated within the following steps of the algorithm (Mitchell, 2011):

1. Generate an initial population of candidate solutions. The simplest way to create the initial population is just to generate a bunch of random programs (strings), called “individuals”.
2. Calculate the fitness of each individual in the current population.
3. Select some number of the individuals with highest fitness to be the parents of the next generation.
4. Pair up the selected parents. Each pair produces offspring by recombining parts of the parents, with some chance of random mutations, and the offspring enter the new population. The selected parents continue creating offspring until the new population is full (i.e. have the same number of individuals as the initial population). The new population now becomes the current population.
5. Go to step 2.

By applying these steps combined with a sufficient method of storing spectral data, it became possible to gradually reconstruct sounds out of random sound. First a population of random sounds, which are stored as random partial descriptions (Klingbeil, 2009) within SuperCollider, would be constructed. Then, a fitness score would be assigned to each sound, based upon a comparison with the partial description of the destined sound. The two sounds which have the best match are chosen and recombined into a new sound, after which the process repeats. For each generation, the fittest candidate would be chosen and resynthesized into a waveform, meaning that a collection of separate audio files functions as the material for this piece. However, not only is there the problem of ordering distinct collections of generations into a form, but also audio files within an evolving population itself. A solution had to be found for inducing a gradual, but audible, evolution of sound.

This is done by introducing a continuous spatial movement among different generations, combined with granular aspects of the material itself, aiming to maximize the amount of contiguity between separate generations.

Similar to Mathematical Gestures, an algorithm functions as a tool to generate audio materials which are between 2 and 60 seconds. However, creating a form visually does not work as efficiently this time. Instead, a narrative format is introduced, directly coupled to the first chapter of the book of Genesis. Each section within the piece represents a day of creation, thus resulting in six distinct sections followed by silence.

1. Creation of Day and Night
2. Creation of Sky and Sea
3. Creation of Land and Vegetation
4. Creation of Stars, Sun and Moon
5. Creation of Sea Creatures
6. Creation of other Creatures and Mankind

Connected to this narrative, I attempted to generate a continuously morphing texture, evolving chronologically. At each day, the texture resembles the things that are being created, though transitioning smoothly instead of discretely. Note that the described algorithm is only used in section five and six, upon which elements of life start to gradually shine through chaotic structures. A stereo version of the piece is given in sound example 6.1.2.

3.3 Sonic Ecosystems

Sonic Ecosystems is a quadraphonic work composed in 2018. Materials are created with the help of SonicSwarms and control-voltage techniques. Such techniques are not only used to create a synthetic soundscape, functioning as a habitat of the sonic agents within, but also to transform the multi-channel material that is derived from SonicSwarms. Ultimately, such a transformation functions as an audible medium through which agents are moving, since SonicSwarms lacks this functionality as mentioned in paragraph 2.4.2. When only a single agent is present within the model, this transformation is done as followed. First, all four channels of audio are send to an amplitude demodulator, which creates four distinct control voltages that contain all information regarding the agent's location. Then, each of these voltages are used to modulate other parameters of synthesis structures, in particular amplitude. In this way, an external source is distributed exactly where the agent is located. Furthermore, each of those voltages can be send to different structures, resulting in a zoned transformative space through which an agent is moving.

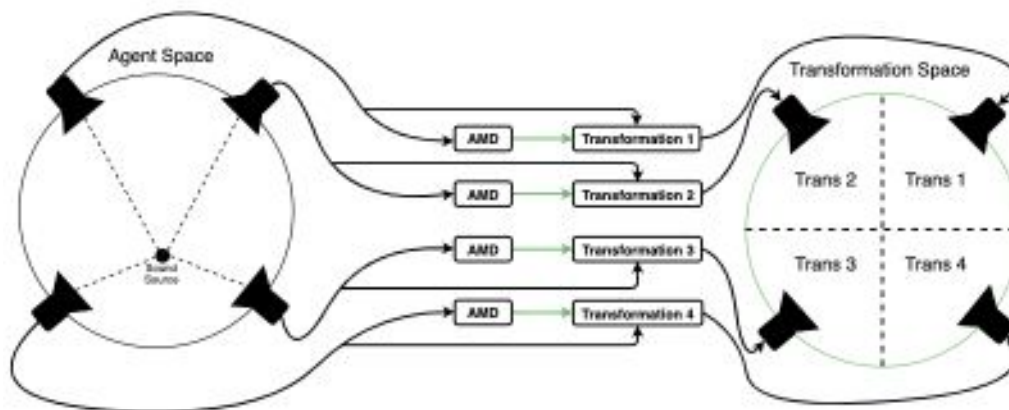


FIGURE 3.3: Multi-channel transformation to shape an audible medium

Moving from a microscopic to a macroscopic domain, decisions regarding form had to be made as well. Similar to the Mathematical Gestures and Sound Genetics, long segments of audio material had to be placed in time, such that their arrangement had a meaningful connection with procedures applied on smaller time domains. There are two main movements constructed within this arrangement, namely a transition between recorded (real) and synthesized (virtual) behaviour, but also a moving between a variety of behaviours. The first is accomplished by slowly layering a field recording of mating frogs, with similar synthetic material. Once a climax of this layering has been reached, a world is opened full of space and movement, of flora and fauna, created with analogue devices and SonicSwarms respectively.

One of the most difficult tasks within this work was transitioning between different behaviours. Because of the amount of spatial information within the material itself, fading amplitude-wise between different materials containing different behaviours did not work out well. Also, abrupt cuts and edits of the material resulted into a very contradicting sound, on one hand extremely generative, fluent and coherent, but on the other hand disrupted by human decisions. Instead of using amplitude based methods to accomplish smooth transitions, I attempted to fade between behaviours, with the help of SonicSwarms' visual interface. Even though the chosen instrument object to which agents are coupled varies discontinuously throughout the piece, their behaviour can be varying smoothly over time. The fact that decisions regarding the design of an instrument object and design of agent's behaviour are separate, is analogous to a distinction between timbre and movement. However, independent from those decisions, there always exists a unification between these two notions, due to the fact that relationships between agents and their environment are constantly formed and changing. It is up to the composer to intertwine timbre and movement via such models, forming a quasi-visual narrative which is a resultant from both the composer's and agents' decisions. A stereo version of the piece is audible in sound example 6.1.3.

Chapter 4

Conclusion

Agent based models and their potential musical value has been discussed and an explicit method of implementation has been given, combined with an analysis of possible behaviours within an acousmatic framework. This method aims to help with the process of creating spatial and spectral relationships between multiple sound sources, overall functioning as a tool for composing acousmatic sound material. Even though it offers much spectral and spatial variety, there still exists a particle based aesthetic between distinct materials, that are generated by the model. Aesthetical variation is difficult to obtain, since a large collection of sound sources distributed within a space, always sound like a dense mass, no matter their relationships. That is why a controllable density within the model is important, which opens up ways of shaping specific sonic trajectories and gestures as well. Another observation is the importance of velocity, not only does it affect the energy of the entire system, but also properties regarding localisation and spectral fluctuation, constructed by distal relationships. Furthermore, correlating a sound source to its space via a force field, enhances the ability to localise it within the perceived sonic environment. Ultimately, the most important observation is how derived material from an agent based model, such as SonicSwarms, is put into compositional perspective. Even though time scales from the micro to the meso are unified within the material itself, it still forms a macroscopic puzzle. By either pre-composing behavioural morphing within- and between materials, choosing synthesis structures which fit the exerted behaviour, or placing agents within a well-suited audible environment, coherence among all temporal scales and spatial domains can be approached. Such coherence results in logical relationships between separate events occurring within the produced acousmatic environment, ultimately shaping a sonic ecosystem which can be perceived as something more than the sum of sound sources within.

Bibliography

- Budón, O. (2000). "Composing with objects, networks and time scales: An interview with Horacio Vaggione". In: *Computer Music Journal* 24.3, pp. 9–22.
- Davis, T. (2010). "Complexity as Process: Complexity-inspired approaches to composition". In: *Organised Sound* 15.2, pp. 137–146.
- Johnson, N. F. (2007). *Two's company, three is complexity*. Oneworld.
- Klingbeil, M. K. (2009). "Spectral Analysis, Editing, and Resynthesis: Methods and Applications". PhD thesis. Columbia University.
- Lewis, G. H. (1875). "Problems of Life and Mind". In: *First series: The Foundations of a Creed* 2.
- Macal, C. M. and M. J. North (2010). "Tutorial on agent-based modelling and simulation". In: *Journal of Simulation*.
- Manning, P. (2013). *Electronic and Computer Music*. Oxford University Press.
- Mitchell, M. (2011). *Complexity: a guided tour*.
- Nyström, E. (2011). "Textons and the propagation of space in Acousmatic Music". In: *Organised Sound* 16.1, pp. 14–26.
- Reeves, W. T. (1983). "Particle Systems – A technique for Modeling a Class of Fuzzy Objects". In: *Computer Graphics* 17.3, pp. 359–375.
- Reich, S. (1968). "Music as a Gradual Process". In: *Perspectives of New Music* 19.1, pp. 373–392.
- Reynolds, C. (1987). "Flocks, Herds, and Schools: A Distributed Behavioral Model". In: *Computer Graphics* 21.4, pp. 25–33.
- (2004). *Steering Behaviours*. URL: <https://www.red3d.com/cwr/steer/Wander.html>.
- Roads, C. (1996). *The computer music tutorial*. The MIT Press.
- (2004). *Microsound*. The MIT Press.
- Schaeffer, P. (1977). *Traité des objets musicaux*. Paris: Editions du Seuil.
- Shiffman, D. (2012). *The Nature of Code*. Shiffman D.
- Smalley, D. (1997). "Spectromorphology: Explaining Sound Shapes". In: *Organised Sound* 12.2, pp. 35–58.
- (2007). "Space-form and the acousmatic image". In: *Organised Sound* 12.1, pp. 35–40.
- Solomos, M. (2005). "An Introduction to Horacio Vaggione's Musical and Theoretical Thought". In: *Contemporary Music Review* 25.4, pp. 311–326.
- Vaggione, H. (1998). "L'espace composable: sur quelques catégories opératoires dans la musique électroacoustique". In: J. M. Chouvel M. Solomos (Eds), *L'espace: musique-philosophie*, pp. 153–166.
- (1999). "The morphological approach". In: *Actes de l'Académie Internationale de musique électroacoustique* 4, pp. 234–239.
- Wilson, S., D. Cottle, and N. Collins (2011). *The SuperCollider Book*. The MIT Press.
- Wright, M. (2005). "Open Sound Control: an enabling technology for musical networking". In: *Organised Sound* 10.3, pp. 193–200.