# COMPOSING INTERACTIONS

Dissertation by

David Pirrò

Submitted for the Degree of
Doktor der Philosophie

Comitee:
Prof. Gerhard Eckel, (University of Music and Performing Arts Graz, Austria)
Prof. Robert Höldrich, (University of Music and Performing Arts Graz, Austria)

Institute of Electronic Music and Acoustics
University of Music and Performing Arts Graz, Austria
November 2017

# *Abstract*

This thesis investigates interaction in computer music composition, particularly in the context of performance-oriented generative music practice. The research follows three approaches of inquiry.

The first is a scholarly and theoretical analysis of the concept of interaction and its understanding in the field of computer music. The topic is discussed in relation to theories of *perception* and *cognition* in philosophy and cognitive sciences, in particular with the concepts of *embodiment* and *enaction*. This approach introduces an understanding of interaction as a temporal process of *mutual influence* taking place between agents. At this point, the concept of the *agent* becomes central to this dissertation.

The second direction of research is based on the mathematical theory of *dynamical systems*. The framework implies a *process-based* mindset and offers an *ecological* perspective that emphasises the role of interrelations between elements in a system. In the context of this work, dynamical systems are understood as the most apt language for formulating and understanding processes of interaction.

A third approach consists of personal *artistic engagement* in the development of interactive computer music environments. This thread interweaves with the former two and allows for continuous *aesthetic experimentation*: speculations and abstract intuitions are put into perceptible form and, in turn, concepts and formulation can be sharpened by experience. An essential part of this engagement relies on the software framework *rattle*, which has been developed for the formulation and the real-time simulation of dynamical systems.

The dissertation develops an approach towards interaction that employs the language of dynamical systems to address the agency of generative computer music processes. Eventually, *agency* is re-interpreted as an essential perceptual quality generative computer music systems should be afforded with to allow for a composition of interactions to emerge.

# *Zusammenfassung*


Diese Dissertation untersucht Interaktion im Kontext der Komposition von Computermusik im Allgemeinen sowie der Praxis performance-orientierter generativer Musik im Besonderen. Die Forschung verfolgt drei methodische Ansätze:

Der erste Ansatz besteht in einer wissenschaftlichen und theoretischen Analyse des Konzeptes von Interaktion und dessen Verständnis im Bereich der Computermusik. Dieses Thema wird in Relation mit Theorien von *Wahrnehmung* und *Kognition* innerhalb von Philosophie und Kognitionswissenschaften gestellt, insbesondere durch die Konzepte von *Embodiment* und *Enaction*. Eingeführt wird eine Auffassung von Interaktion als einem zeitlichen Prozess *gegenseitiger Beeinflussung*, die zwischen Agenten stattfindet. An dieser Stelle entwickelt sich das Konzept des *Agent* zu einem zentralen Thema der Dissertation.

Die zweite eingeschlagene Richtung der Forschung basiert auf der mathematischen Theorie *dynamischer Systeme*. Dieses Bezugssystem gewährt eine *prozessbasierte* Denkart und eine *ökologische* Perspektive, welche die Rolle von Wechselbeziehungen zwischen Elementen eines Systems betont. Im Rahmen der vorliegenden Arbeit wird dieser Ansatz als die geeignetste Sprache betrachtet, um Prozesse der Interaktion zu formulieren und zu verstehen.

Ein dritter Ansatz besteht in der persönlichen *künstlerischen Beschäftigung* mit der Entwicklung interaktiver Computermusik-umgebungen. Dieser Strang wird mit den beiden vorherigen verwoben und ermöglicht das kontinuierliche *ästhetische Experimentieren*: Vermutungen und abstrakte Intuitionen werden in wahrnehmbare Form überführt, und umgekehrt können Konzepte und Formulierungen durch die Erfahrung geschärft werden. Ein wesentlicher Teil dieser Beschäftigung stützt sich auf das Software-Framework *rattle*, das für die Beschreibung und Echtzeitsimulation dynamischer Systeme entwickelt wurde.

Diese Dissertation entwickelt einen Standpunkt hinsichtlich Interaktion, welcher die Sprache dynamischer Systeme gebraucht um, die Wirkmächtigkeit generativer Computermusikprozesse zu erfassen. Schlussendlich wird *Wirkmächtigkeit* (*agency*) als eine essentielle Wahrnehmungsqualität neu interpretiert, mit welchen generative Computermusiksysteme auszustatten sind, um das Komponieren von Interaktionen zu ermöglichen.

# *Acknowledgments*

To my partner Theresia for her encouragement and support.
To Martin Rumori and Hanns Holger Rutz for their friendship and the inspiring discussions.
To my mother Ella and my father Giuseppe for their assistance and their motivation.
To all my friends and colleagues.
To the Institute of Electronic Music and Acoustics in Graz for being a great place to work.
To Agostino Di Scipio for being a great source of inspiration.

To my supervisors Prof. Gerhard Eckel and Prof. Robert Höldrich for their guidance and the exchange of thoughts, ideas, and experiences.

# Contents

# *List of Figures*

# 1 Introduction

Introductions are usually written last and this one is no exception. This chapter has been written after the completion of a work traversing years of engagement with different research themes, projects, and artistic practices. Consequently, it feels more like a conclusion than a preface. A good introduction should frame what will come next, easing the entrance into the following text. But it should also transcend this aim, providing the reader with a few building blocks for understanding, without giving too much away.

To this end, this chapter first presents a personal introduction to this work and the motivation behind it. This will be followed by a historical reconstruction of the path that led to this work as well as a clarification of its methodologies.

## 1.1 Motivation

This dissertation has its origins in the practice of *Computer Music*: it could be described as 'practice-motivated'.

With a background in traditional, acoustic modes of musical performance, my initial encounters with computer music offered new and fascinating possibilities. I could plunge headfirst into the smallest details of synthesis while simultaneously organising sound at larger timescales. All temporal aspects of a musical composition seemed to be accessible at once. Most interesting to me was the *physical* engagement with sound that the computer made possible - a previously unimaginable prospect had become a Utopia of a bodily engagement with composition and sound.

However, with this fascination corresponded a frustration with the realities of performative practice in computer music. Practices like Live-Electronics are often shaded by incoherence or *dissonance* between potentialities offered by electronic media and the modes of performance through which these possibilities can be engaged with. As the role of the performer reduces to that of a mere controller of an often very complex computational machinery, the promise of a total engagement with sound could not be fulfilled. I know now that part of the problem lies in a misunderstanding of what musical engagement means, influenced as it is so strongly by traditional music practice. This mindset does not justice to the specificity of the computational medium.

When I began this thesis, I perceived the problem to be a lack of *interactivity* between the musician or composer and the computer, which I set forth to tackle. This theme is not unknown in computer music research and has been addressed in many different ways and from many perspectives.[1] The source of trouble is typically located at the split between the processes of sound synthesis and the interface they present for the user's or musician's interaction. The aim of this research was to understand how interfaces could be designed to allow for a more engaging, physical, bodily and intuitive relationship with computer music systems.

[1] Newton Armstrong. *An enactive approach to digital musical instrument design*. PhD thesis, Princeton University, 2006; Bob Ostertag. Human bodies, computer music. *Leonardo Music Journal*, 21:19-23, 2006; and F Richard Moore. The dysfunctions of midi. *Computer music journal*, 12(1): 19-28, 1988

Even if, over the years, my research questions have evolved, *interaction in computer music* remains at the core of this dissertation. It signifies a line of inquiry that seeks to address the fundamental qualities of both interaction and of computer music.

## *1.2 Historical Path*

The research question the dissertation was set to answer, at least at its beginning, could be formulated as:

> *Can simulated physical models allow for greater bodily and intuitive interaction with computer music instruments?*

This question was influenced by two main factors. Firstly, I had previously studied Theoretical Physics with a particular emphasis on Computational Physics. I was therefore already acquainted with the theoretical and technical knowledge that would allow me to address the problem. Secondly, and most significantly, was that I was involved in the *Embodied Generative Music* research project at the *Institute of Electronic Music and Acoustics* in Graz. The project's main research theme centred around the dissociation of sound and bodily movement, pursued from both a scientific perspective and a performance-oriented computer music practice. Due to its thematic proximity, the project was extremely influential to this work.

The idea behind the above question was to develop a practice of interface design. These interfaces would tap into our implicit bodily knowledge of the physical world, allowing for interaction with computational and sound synthesis processes. We hold embodied knowledge of the world's 'mechanisms', the rules the physical world exposes us to through our bodily interactions. Interfaces directly modelled on these same environmental dynamics would, I hypothesised, elicit resonances in the user or the performer on a bodily level.

The theoretical basis for this perspective, as well as for the Embodied Generative Music project more generally, is the *Embodiment Theory* of cognition. The theory holds that our perceptual and motor systems are responsible for shaping fundamental aspects of human cognition. As the body and its interactions within the environment are essential in the formation of higher functions of our brain, cognition cannot be understood as a process detached from the world and the body in which it takes place. The theory posits that the materiality of the body transcends its function as physiological substrate, strongly shaping thought processes. The embodiment theory originates in philosophical discourse, but it has since spread to research fields as diverse as neuroscience, psychology, linguistics, neurobiology, and even robotics and artificial intelligence. It opens up a new perspective for scholars looking to include the body in their thinking. This is also how embodiment theory has found its way into the fields of interaction design and, in particular computer music, where it has become the basis for addressing concerns regarding the lack of bodily presence in composition and performance.

In the context of the Embodied Generative Project, our understanding of embodiment was that of an extension of the body into an unfolding, generative sound process. The dancers we worked with during the project would be able to extend their proprioception such that their body would be allowed to *inhabit* the sound. The metaphor we used to describe this situation was that of slipping into a dress, which then would follow the movements continuously adapting itself according to each action or movement: sound would be that dress. We were following what could be considered a classical *Human Computer Interaction* approach in searching for design strategies which would generate *transparent interfaces*. A transparent interface is not material, but rather acts as an ideally non-conditioning channel for the transmission of information between the performer, who is providing the input, and the actual system, which is to be fully controlled. The interface is purely functional in linking these two actors, connecting the user to what is 'behind' the computer music system.

These thoughts were also influenced by the artistic works I was engaging with, both within the *Embodied Generative Music* project and in my own practice. This practical, artistic and aesthetic engagement troubled our understanding of embodiment and our tools. We discovered that these were insufficient for understanding the relation between performer and computer music system. In particular, the interaction with real-time generative computer music processes seemed to be a qualitatively different problem to non-generative processes, needing other conceptual tools.

This was the moment in which a search for an alternative understanding and theoretical framing began: an inquiry that forced us to rethink the premise of this research and the concepts on which it was based.

- What are computer music instruments?
- How can interaction be defined?
- What are physical models?
- How do they resonate with perception?
- How does perception work?
- What does 'bodily' mean?

These are just a few questions that emerged during this period. They are broad and at this point the research project, too, dramatically expanded in the most unexpected directions. Looking for guidance, I ventured into neurophysiology, cognitive sciences, philosophy, interaction design, dynamical systems theory, and cybernetics, to name a few. That is, I strode into a nexus of different research streams, which would blow the theme of this dissertation into vast dimensions: the exact opposite of Umberto Eco's recommendation to narrowing down a thesis' subject.[2] Nevertheless, I was once again inspired. The research fell into a *pulviscular* state, in which many ideas, concepts and experiences were floating around with unclear connections. It also exploded my central research question.

[2] Umberto Eco. *How to write a thesis*. MIT Press, Cambridge, Massachusetts, 2015

The word *pulviscular* is not really an English word. It has been used as an approximate translation of the Italian *pulviscolare* which writer Italo Calvino uses in some of his texts. Aside from my personal liking for the word, its meaning - as indicating something that is constituted by fine and almost impalpable particles - aptly describes this moment of my research. Calvino uses the word to describe text which, rather then presenting linear narrative development, is instead *constructed* by the reader: he or she identifies particles or concentrations of interest and draws connections between them, generating a vibrating net through which the substance reveals.[3]

[3] Italo Calvino. *If on a Winter's Night a Traveler*. Houghton Mifflin Harcourt, 1981

> *Reading is a discontinuous and fragmentary operation. Or, rather, the object of reading is a punctiform and pulviscular material. In the spreading expanse*

> *of the writing, the reader's attention isolates some minimal segments, juxtapositions of words, metaphors, syntactic nexuses, logical passages, lexical peculiarities that prove to possess an extremely concentrated density of meaning. They are like elemental particles making up the work's nucleus, around which all the rest revolves. Or else like the void at the bottom of a vortex which sucks in and swallows currents. It is through these apertures that, in barely perceptible flashes, the truth the book may bear is revealed, its ultimate substance.*

I regard this thesis as the textual trace of an endeavour towards a reconstruction, a condensation of this 'dust' into a few gravitational centres. In some parts, its original pulviscular nature might still be felt; the narrative path I have attempted to construct should help the reader.

## *1.3 Research Methods*

The dissertation relies on an ensemble of research methods.

I have adopted *scientific methods* to analyse existing research literature in computer music research and to address specific ideas in cognitive sciences and philosophy. These investigations served as the basis for the formulation of ideas and hypotheses.

Further, I have employed *technological research* methods while developing software tools in order to test my assumptions. The realisation of these instruments involves multiple cycles of hypothesis building, experimentation, testing and step-wise improvement towards an aim. Moreover, in my experience, this processes cannot be reduced to purely functional activities, i.e. only defined by their final output. The process of development also informed the research process on a conceptual level: formulations of abstract ideas in terms of technological artefacts always affect how those ideas might be experienced and re-formulated. I therefore regard the development of software tools for my research not only a necessity, but also as a valuable method of research.

The most important method in this dissertation is however *artistic practice*. The praxis of computer music itself became one of the most important tools in addressing issues of interactivity. This went beyond the production of pieces or even of proof of concepts to a purely experimental sense. Artistic practice has been used as a tool for generating conditions in which specific attributes of *interaction* might be seen: an experimental condition 'whose outcome cannot be foreseen'[4].

[4] Bob Gilmore. Five maps of the experimental world. *Artistic Experimentation in Music: An Anthology*, pages 23-29, 2014

The artistic works I reference here do not centre around a 'result' they should provide or clearly defined arrival point in the research (although they may have one). These artistic artefacts should instead be understood as generators of experiences and reflections, and in this sense they are generative. They are artefacts through which speculations may be pursued, inspiration can be drawn, or new questions raised. These artefacts form a dialectical relationship with this dissertation's concepts and technologies. Further, in remembering that the issues of *interaction* are primarily *aesthetic*, artistic works are the most direct means of access to those questions.

Throughout this text I make use of the term *aesthetics* in diverse forms (e.g. as *aesthetic experience*), and at this point it seems necessary to elucidate the term. This is not a dissertation in philosophy and I am not trained in this discipline, so I will not be able to cover this concept in great depth, much less the numerous controversies over the last 300 years of philosophical discourse. Instead I will limit myself to clarifying the term as it is used in this text, making no attempt to be exhaustive.

I understand the aesthetic in the sense of the German *Rezeptionsästhetik*. This is the philosophical discipline that considers the sensuous and cognitive reception of artistic works and, in particular, how this reception is influenced by factors that may be located in the work itself. A good example of what I mean is Umberto Eco's understanding of 'Open Work': a work, in his case a text, which allows for multiple interpretations and meaning, whose primary value is to permit and elicit the readers' (or as Eco further explains with regards to Brecht, an audience's) interpretative action. It is a work that does not present a 'solution', but the affordances to construct one. It is therefore *in* the artistic work that qualities are placed which affect and evoke mechanisms of perception and cognition.[5]

[5] Umberto Eco. *Opera aperta*. Harvard University Press, 1989

The aesthetic I am interested in does not refer to questions of the formation of taste, or to the qualities that allow an object to be judged as work of art. More interesting are questions of perception or, more precisely, the *shape of perception*. An aesthetic object is, in this sense, not just a perceptual object, but an artefact whose nature is to bring the mechanisms of perception to light. An aesthetic experience is an experience that points to the function of perception, it makes it conscious. Philosopher Alva Noë holds that artistic practice is the most effective tool for making inquiries into perceptual consciousness, and consequently into aesthetics (in this sense, at least) more generally.[6] Noë refers here to the

[6] Alva Noë. Experience and experiment in art. *Journal of Consciousness Studies*, 7(8-9): 123-136, 2000

words of installation artist Robert Irwin:

> *To be an artist is not a matter of making paintings or objects at all. What we are really dealing with is our state of consciousness and the shape of our perception.*[7]
>
> *The act of art has turned to a direct examination of our perceptual processes.*[8]

I fully share this understanding of artistic practice. Therefore, as the phenomena this work focuses on are of a perceptual nature, following this reasoning, artistic praxis seems the best method to address them.

Thus, the role of the artistic works I cite in this dissertation, especially in the appendix, is a structural one: they are arguments which are instrumental to my research and cannot be relegated outside. These works are deeply rooted within this research process and, in turn, they also strongly influenced and shaped that same process.

[7] Robert Irwin. The state of the real. In Beatrice Hohenegger, editor, *Notes Towards a Conditional Art*, chapter 7, pages 49 - 53. Getty Publications, 1972a

[8] Robert Irwin. Re-shaping the shape iof things. In Beatrice Hohenegger, editor, *Notes Towards a Conditional Art*, chapter 8, pages 54 - 60. Getty Publications, 1972b

## *1.4 Structure of the Thesis*

The 'condensation' operation I've described above produced three major axes along which this dissertation unfolds: each is treated in one of the central chapters of the text. Every chapter follows its own narrative, which is why the transitions may appear a little abrupt: the themes they centre on are of very different nature. There is something like a red thread joining them, which will emerge through the process of reading.

The work is organised as follows:

- Chapter 2 Interaction departs from the historical development of electronic and computer music. The theme of interaction is introduced as it appears in these practices, which is then brought in context with embodied theory and, later, with the enactive approach to cognition, a core element of this work. In this chapter, definitions of most of the terms used in the text are given.

- Chapter 3 Dynamical Systems introduces the theory of dynamical systems, first from a mathematical perspective, and then as a general language and thought framework for processes of temporal evolution and interaction. An explanation follows showing how this language is used in cognitive sciences, the study of perception and computer music, and provides a justification for its use.

- Chapter 4 Case Studies describes a path through the actual technical and artistic engagement with the theme of interaction. The three case studies discussed in this chapter are the most significant in terms of their effects. It begins with the development of a software framework, passes through artistic research in the context of the *Embodied Generative Music* project, and ends with a case study featuring how thinking in terms of dynamical systems might shape the development of a computer music environment for interaction.

- The Chapter 5 Conclusions and Outlook offers a résumé of the thesis and highlights its central claims. Furthermore, prospective and ongoing research directions connected to this work are described.
- The Appendixes contain diverse materials. A catalogue of works in appendix A aggregates descriptions of some of the artistic works that were central for this dissertation. The remaining sections are concerned with detailed descriptions of the formulations used in the implementations of software tools I have realised.

# 2
# *Interaction*

This chapter attempts to bring together the most important concepts I use later in this work. It provides the context from which these concepts have been drawn, for myself at least - I make no claim to completeness. In collating these concepts, I hope to construct a narrative, beginning with the historical genesis of electronic and computer music, traversing through to live-electronics and interactive practices, touching on ecological psychology and embodiment cognitive theory, and ending with the theory of enaction and agency. *Generative music*, *live-electronics*, *interactive composing*, *affordance*, *ecology*, *embodiment*, *enaction* and *agency* are just a handful of terms and concepts to be found on this path.

## *2.1 Computer music: a generative art*

The central theme of this work, as its title suggests, relates to the theme of *interaction* specifically in the context of computer music composition.

*Interaction* seems to have become an ubiquitous term nowadays. Generally speaking, it indicates the ability of a tool, usually digital (i.e. a programme that is executed on some digital device or computer), but also simpler artefacts like sliding doors, to be able to accept or sense input and adjust its state according to some internal rules. This characterisation is, of course, very broad and may fall apart when confronted with more specific situations.

Questions related to interactivity have been significant for *Electronic Music*, either implicitly or explicitly, since its beginnings. Early electronic music researchers used instruments with a radically different kind of relationship between the bodily action of their operators and generated sound. Compared to traditional acoustic instruments, these electronic instruments have a much smaller (if it exists at all) dependence on the energy input of bodily gestures or movement. One could say that, in general, the ratio between energy that is introduced into the instrument's system by the operator's bodily movement and its perceived sonic effect is significantly smaller with a *Moog* synthesizer than a double bass.

This does not necessarily mean that those new instruments are more efficient. Rather, the cause for this dramatic shift lies in the injection of a second form of energy (beside that provided by the human body) into the instruments' system: *electrical* energy. This form of energy provided means of an *amplification*, or even an indefinite *sustain*, of an operator's actions, reducing the bodily effort required for sound generation, or almost eliminating it altogether. It is a (perceived) *infinite* source of power and, therefore, possibilities. That is, the use of a secondary energy source – electrical power – in conjunction with a means of manipulating it (analogue circuitry such as filters or oscillators), greatly expands the field of sound generation mechanisms and renders direct bodily action less and less necessary.

One exception to this, however, is the organ. This instrument, which long precedes any electronic instrument, also made use of an additional energy source to greatly amplify the player's actions. It is not surprising that many of the first electronic instruments, like the *Dynamophone* or Oskar Sala's *Trautonium*[1], as well as more recent synthesizers like the *Moog*[2] were modelled after the organ.

This unfurling of the connection between body and sound generation saw a gulf open between them where there was once continuity. This freed up space for a re-composition of this relationship, a space of possibilities to re-think the body-sound relationship: the most obvious examples being the *theremin*[3] or the *terpsitone*.

With the advent of early digital computing machines came *Computer Music*, developing out of *Electronic Music*[4] inheriting the same relationship between bodily action and sound producing devices, but also bringing into play new and distinctive qualities to compositional practice. The computer, the medium in which this kind of music is composed, allowed for a fundamental shift in music-making: this new 'instrument' facilitates the formulation and execution of *processes*, programmes or algorithms able to generate complex formal structures. *Algorithmic composition*, the praxis in which a set of rules devised by the composer generates musical scores, has been applied in compositional praxis to various degrees since Guido d'Arezzo[5] (early middle Ages), and saw a great upswing with the advent of the computer. Lejaren Hiller is one of the very first experimental computer music composers of that time to engage with the novel possibilities of the *Illiac* computer installed at the University of Illinois, Urbana-Champaign.[6]

[1] André Ruschkowski. *Elektronische Klänge und musikalische Entdeckungen*. Reclam, 1998

[2] Robert A Moog. Voltage-controlled electronic music modules. In *Audio Engineering Society Convention 16*. Audio Engineering Society, 1964

[3] Leon S Theremin and Oleg Petrishev. The design of a musical instrument based on cathode relays. *Leonardo Music Journal*, 6(1):49–50, 1996

[4] A clear distinction between *Electronic Music* and *Computer Music* might seem, at least nowadays, difficult, as almost every electronic device is in fact integrated with computer. The distinction I try to hold here is used as a rethorical tool to make specific characteristics of different practices as they emerged historically clearer.

[5] Gerhard Nierhaus. *Algorithmic composition: paradigms of automated music generation*. Springer Science & Business Media, 2009

[6] Lejaren Arthur Hiller and Leonard M Isaacson. *Experimental Music; Composition with an electronic computer*. Greenwood Publishing Group Inc., 1979

Computers gave sound synthesis new capabilities which, when combined with he subsequent widespread popularity of tape recording technology and sound projection devices, led to even further crucial developments. At this point, not only computer music offered an unprecedented range possibilities of the *simultaneous* composition of sound and music. Loudspeakers could be used for sound projection, meaning the complete process of composition, realisation and performance was in the composers' hands. The establishment of electronic music studios consolidated an emergent compositional practice where the composer would work in isolation and autonomy. There was no longer the need to rely on performers to have one's music generated and no need to cope with the indeterminacy and the subjectivity of human interpretation. This situation nourished a desire for total control in many composers.

Early composers who worked with these emergent tools not only realised groundbreaking musical works, but also generated and contributed to new discourse in and around composition, often by writing about their ideas and practice. Herbert Brün[7], Gottfried Michael Koenig[8] and, most notably, Iannis Xenakis[9] are just a few who were decisive for the future developments in computer music.

More generally, this musical practice shared some of its roots with *conceptual art*, which was concurrently influencing visual artistic practice. As summarised by Sol Lewitt:[10]

> *In conceptual art the idea or concept is the most important aspect of the work. When an artist uses a conceptual form of art, it means that all of the planning and decisions are made beforehand and the execution is a perfunctory affair. The idea becomes a machine that makes the art.*

Algorithmic art and computer art developed these ideas even further, moving away from materiality towards abstraction. Take, for example, the work and writing of Georg Nees and Frieder Nake: 'Computer art is concept art insofar as it describes an idea and does not show the material work'; 'Computer art shares with conceptual art [. . .] a neglect of materiality'.[11] In a similar vein, computer music at this time sought an (almost) complete disconnect from traditional modes of musical performance. The central elements of musical artwork were the processes of production, rules and algorithms programmed into, and performed by, the computer.

These movements resulted in what is today known as *Generative Music*. Pushed by the advent of personal computing during the 80s and by the development of high level

[7] Herbert Brün. *über Musik und zum Computer*. G. Braun, 1971

[8] Gottfried Michael Koenig. Kompositionsprozesse. In *Ästhetische Praxis*, volume 3 of *Texte zur Musik*, pages 191-210. PFAU Verlag, Saarbrücken, 1993

[9] Iannis Xenakis. *Formalized music: thought and mathematics in composition*. Pendragon Press, 1992

[10] Sol LeWitt. Paragraphs on conceptual art. *Artforum*, 5(10): 79-83, 1967

[11] Frieder Nake. Paragraphs on computer art, past and present. In *Proceedings of CAT 2010 London Conference*, pages 55-63, 2010

programming frameworks for sound synthesis and algorithmic control (e.g. Max was developed by Miller Puckette at IRCAM and made public around the beginning of 1990), generative music started to play an increasingly important role in music production - this remains the case, as the recent works by composer Brian Eno testify. In the composition of Generative Music, the composer/musician formulates a process that then generates the music, developing sound without any further human intervention. In the case of *generative computer music* the definition is more precise, being, as Nick Collins puts it, music 'produced by leaving a computer program to run by itself, with minimal or zero interference from a human being'[12]. Generative computer programmes could therefore be considered as examples of *derivative intentionality*[13], in which the code or algorithmic formulation is written by a human composer who then retreats and yields the autonomy of execution to the machine.

[12] Nick Collins. The analysis of generative music programs. *Organised Sound*, 13(3):237-248, 2008

[13] John R Searle. *Mind: a brief introduction*. Oxford University Press, 2004

Generative music is a praxis that resonates with the intrinsic characteristics of the computational medium: it provides the tools for the formulation of processes and, at the same time, the space for their realisation and actualisation. My understanding of computer music centres on its generative potentialities: this is the understanding I will implicitly use throughout this work.

## 2.2 Live-Electronics and interactive composing

*Process composing* or *composition of processes* is not only central to computer generated music. In particular John Cage's work was paradigmatic in this sense, as it transcended the boundaries drawn by the means employed in its execution. He said about his compositions: 'I was to move from structure to process, from music as an object having parts, to music without beginning, middle, or end, music as weather'[14]. Many of his works (as for example *Fontana Mix*) would consist of instructions for how the score could be generated. He provided performers with directions, initiating a process of putting together, literally composing, the materials of the piece in a certain fashion. Most of those processes included and depended on aleatory elements, operations of chance, whose function was to steer the construction of the work away from the idiosyncrasies of personal taste or subjective interpretation. Cage used these operations in order to free the composition and himself, the composer, from individual taste and memory and to initiate a process which would produce something he did or could not imagine: to be, in a way, surprised. In his words:[15]

[14] John Cage. John cage: An autobiographical statement, 1990. URL http://johncage.org/autobiographical_statement.html. Accessed on 29/10/2017

[15] John Cage and Roger Reynolds. An interview with john cage on the occasion of the publication of silence. *Generation - The University Inter-Arts Magazine*, pages 40-51, November 1961

> *What actually happened was that when things happened that were not in line with my views as to what would be pleasing, I discovered that they altered my awareness. That is to say, I saw that things which I didn't think would be pleasing were in fact pleasing, and so my views gradually changed from particular ideas as to what would be pleasing, toward no ideas as to what would be pleasing.*

The use of chance was also, for similar reasons, one of the fundamental ingredients in computer-aided algorithmic composition. Cage was not only aware of this, he was involved: together with Lejaren Hiller he composed the piece *HPSCHD* at Experimental Music Studios at the University of Illinois at Urbana-Champaign, a composition based on a random sampling of scores and pre-recorded tapes, controlled by computer processes.

But there was more to Cage's use of chance and *indeterminacy*. He wished to embrace the indeterminacy of the world and of performance, juxtaposing it with the traditional utopian narrative of a composition that exists outside of time and of the contingencies of its realisation. He placed indeterminacy at the very centre of his work, elevating it to an organisation principle, to a generator of experiences: 'I don't think we're really interested in the validity of compositions anymore. We're interested in the experiences of things.' As Joel Chadabe puts it, the use of indeterminacy in Cage's work points 'back', out of the electronic studio and into the liveliness of performance.[16] It's no coincidence that Cage is considered one of the initiators of *Live-Electronics*.

[16] Joel Chadabe. The history of electronic music as a reflection of sructural paradigms. *Leonardo Music Journal*, 16:41-44, 1996

Although antecedents can be traced to Cahill's Dynamophone and the subsequent development of electronic instruments in the period between the two world wars, live-electronics more specifically indicates a practice driven by the desire to bring production processes and technologies onto the stage which, at that point (the beginning of the 60s), were still relegated to the studio. Tape recorders, microphones, sine-wave generators and effects such as ring-modulators entered the stage in compositions by, for example, Stockhausen (*Mikrophonie*), Kagel (*Transition II*) or Lucier (*Music for solo Performer*).[17] This time also saw the emergence of ensembles like *Musica Elettronica Viva*, *Gruppo di Improvvisazione Nuova Consonanza* and *AMM* which were, to varying degrees, based on a practice of improvisation, incorporating performative habits common to other musical traditions such as jazz.

[17] Peter Manning. *Electronic and computer music*. Oxford University Press, 2013

In the context of computer music, algorithmic practices combined with sound synthesis put the composer in a unique position. In traditional composition, a composer needed to wait until the point of performance before they could hear the music (this excludes the composer's own *inner hearing* of the music during the compositional process). here would likely be a large temporal separation between the writing of the piece and its realisation and listening. The computer radically changed this situation. Although the first calculators at the disposal of musical experimenters meant a gap of some hours between the start of the score and synthesis generating programme and its output, the rapid development of digital technologies closed this temporal separation. Eventually the gulf between formulation (of the processes' rules) on one side and the realisation and the listening to a composition on the other shrank, allowing the composer to oscillate between these two states so rapidly that they almost conflated into one continuous action. Suddenly generative music practices, at least in the studio, began to be interactive.

At this point, computer music composers could begin to *perform* their music as they were composing it. This situation enabled composers to extend their actions and their bodies into the construction of a composition's structural aspects in a completely new way. An evolution that was pushed forward by the rapid development of tools allowing a high-level control of sound synthesis processes (e.g. sequencers) and later by a growing number of control interfaces which could be paired with digital computers (e.g. joysticks). Catalysed by a similar impulse that led to live-electronics, computer composers saw the opportunity to bring their studio practice onto the stage, where the act of composition would be the object of performance. Joel Chadabe was one of the first to engage with this potentiality, working on pieces that included an interactive control of the composition and by developing the idea of *interactive composing*. He sees this practice as 'a two-stage process that consists of (1) creating an interactive composing system and (2) simultaneously composing and performing by interacting with that system as it functions.' Regarding his idea of interaction:[18]

[18] Joel Chadabe. Interactive composing: An overview. *Computer Music Journal*, 8(1):22-27, 1984

> *The performer [...] shares control of the music with information that is automatically generated by the computer, and that information contains unpredictable elements to which the performer reacts while performing. The computer responds to the performer and the performer reacts to the computer, and the music takes its form through that mutually influential, interactive relationship.*

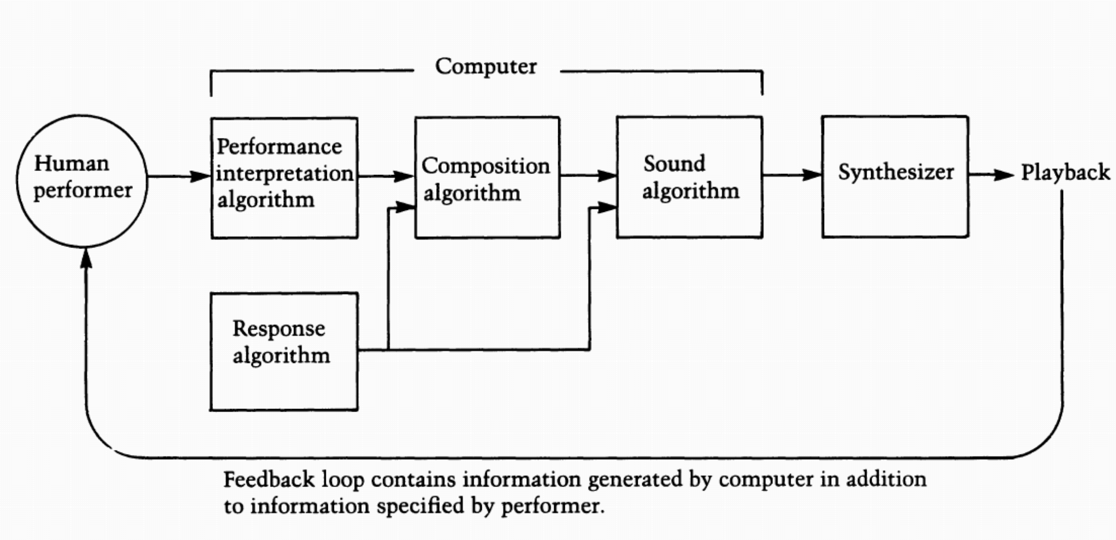


Figure 2.1: Organisation of on interactive composing system. From Joel Chadabe *Interactive Composing: an overview*, Computer Music Journal, 1984, pp. 22-27

In order to create this situation, the 'interactive composing system operates as an intelligent instrument'. Implemented on a programmable computer, this instrument should (see figure 2.1): 'Interpret a performer's actions as partial controls for the music. Generate controls for those aspects of the music not controlled by the performer. Direct the synthesizer in generating sounds'. I believe the above thoughts to be seminal for interactive computer music practice and the research that would follow.

A important detail implicit in Chadabe's words appears to be the mixture or, perhaps more appropriately, the collision of roles, practices and disciplines he addresses. Suddenly the boundaries of the composer's and performer's roles are blurred and even overlapping: the composer performs his composition and the performer composes while performing. The composition is simultaneously an instrument. Composers/musicians/performers must then also programme and construct their instruments, and could also be described as engineers. This mixture of different, and sometimes contrasting, concepts lends computer music its appropriative character, right from its very beginnings: themes, technologies and ideas from the most disparate disciplines find their way into the discourse around computer music keeping it diverse and lively. The downside is that such imbrication with other musical practices hinders the possibility of a clear definition, something still lacking today.

However, central to Chadabe's formulation is a fundamental tension that is generated by the juxtaposition of an understanding of interaction as *mutual influence* of performer/composer and the computer music system and the notion of the latter as an instrument. On the one hand, an instrument is a means to achieve a determined result by being employed in a specific way and is required to produce a consistent output. On the other hand, a situation of mutual or shared control between performer and computer music system presupposes an independence of the involved entities: both would be able to receive external influences and to take decisions according to some internal (possibly evolving) mechanisms: both actors are expected to be able to take and exert control and none is subordinate to the other. I believe that these two perspectives on computer music systems ascribe qualities to them which are qualitatively opposing. Chadabe tries to build a bridge between them when he attributes an *intelligence* quality to the instrument, but what he means is a capacity of the computer system to *sense*, interpret (correctly) and transpose the performer's input: intelligent means, in this context, reactive. This tension - between the 'instrumental' and the 'agential' perspective - shows how computer music systems posed challenges which are difficult to overcome with traditional musical concepts and roles: performer, instrument, score and composer do not fit with computer music, especially when it encompasses an interactive aspect. However, such uncharted space can be seen spanning between those opposing concepts, a space in which composition could extend into: the composition of interactive relationships.

With the development of the interactive dimension, computer music merged with already-existing live-electronic practice. The flexibility and growing possibilities offered by personal digital computers meant that they rapidly substituted analogue equipment in live-electronic contexts. Consequently, computer music systems, especially in connection with performative practices, have increasingly shifted towards an instrumental perspective. Such a perspective has been adopted by most of the researchers and composers, and continues to be the case today. Although the special dual nature of computer music systems as *composed instruments*[19], being both instrument and score, is mostly acknowledged„ the generative character of computer music systems is often overlooked in favour of more functional view which reduces computer music processes to sound synthesis. The relation between performer and computer music system is often thought of as a linear communication flow, an understanding strongly rooted in an instrumental perspective.[20] Information is sent by the performer towards the computer music instrument which interprets it and produces output accordingly. Such a schema (similar as in figure 2.1) can be found in different iterations in most of the literature pertaining interaction design in computer music. In those texts, the backward communication

[19] Norbert Schnell and Marc Battier. Introducing composed instruments, technical and musicological implications. In *Proceedings of the 2002 conference on New interfaces for musical expression*, pages 1-5. National University of Singapore, 2002

[20] Agostino Di Scipio. 'Sound is the interface': from interactive to ecosystemic signal processing. *Organised Sound*, 8(3):269-277, 2003

channel, going from the instrument to the performer, is underspecified and limited to the function of a feedback channel simply informing the user of the system's status, another indication of the widely adopted unidirectional and performer-centred perspective.

The instrumental condition of computer music systems bears an important characteristic which distinguishes them from traditional acoustic instruments. A distinction can be made between the sound-generating components and its parameter-regulating control mechanisms. A simple oscillator, for example, can be controlled and 'played' through a keyboard-like interface, or a 'machine-like' interface with knobs and faders, or a touchless interface as with the *Theremin* or the proximity-sensitive antennas Joel Chadabe uses in his piece *Solo*. Any musical instrument, whether acoustic or computer, can be separated into a sound generator and a *performance device* - the interface with which the instrument can be played - with a link between them.[21] With acoustic instruments, the performance device is a structural part of the instrument's sound generating mechanism, whereas in a typical computer or digital instrument this is not the case. As it is not integral to an instrument, any performance device, ranging from traditional keys and knobs to sensors of any kind, can be attached to any of those digital instruments. Therefore, the act of choosing and developing a mechanism (as an interface) linking movement to a sound generator process becomes part of the compositional process and has to be placed in the addressed musical context and the performative role given to the musician. That is, the relation between bodily performance and sound synthesis can be made subject to re-composition: questions of interface design can and have therefore to be addressed from inside compositional practice.

[21] Joel Chadabe. *Electric Sound:The Past and Promise of Electronic Music*. Prentice-Hall, Upper Saddle River, New Jersey, 1997

With the growth of possible interfacing technologies, so too emerged a new research field at the intersection of *HCI: Human Computer Interaction* research and computer music, now commonly defined as *NIME: New Interface for Musical Expression*[22] after the eponymous conference series. While instrument design and interaction design are equated, a plethora of new instruments are developed exploring different ways to connect performer and sound synthesis.[23]

[22] http://www.nime.org/, accessed on 03/11/2017

[23] Sergi Jorda. *Digital Lutherie Crafting musical computers for new musics' performance and improvisation*. PhD thesis, Department of Information and Communication Technologies, 2005

As new sensing and motion tracking technologies allow us to capture the performer's bodily actions and movements in new ways, today's interaction research is understood as an endeavour in finding solutions to the problem of how that sensed information enters the computer music instrument. Questions of *mapping*, a general approach employing functions, or maps, to connect or translate input information to the parameters of the digital instrument, become central to many applications centred on gestural input (or control).[24] Although increasing the complexity of the mapping function seems to result in more engaging situations,[25] it remains true that within such an approach, the computer music system remains a deterministic machine completely under the control of the user: the *generative potential* of computer music systems is here completely suppressed.[26]

An alternative approach is to assign to the computer music system the capabilities of analysis of the input information, inputting the result of that analysis into the sound producing process. That is, the computer is now understood and implemented as a sort of *listener* and *interpreter* of the sensed input.[27] Especially reading Rowe's work, it is clear that here interaction has typically 'been predicated on the technical acquisition of information about the momentary relationship of action (body) and reaction (system)': the proposed model of interaction 'draws from human conversation'[28]. A similar, highly sophisticated example in this direction is the work of George Lewis on his composition *Voyager*, which employs a 'virtual interactive computer-driven, improvising orchestra that analyses an improvisor's performance in real time, generating both complex responses to the musician's playing and independent behavior arising from the programme's own internal processes'[29]. Peculiar to this work, which is based on a very idiomatic conception of musical performance rooted in free-jazz practice and culture, is the author's understanding of the relationship between musician and system: 'there is no built-in hierarchy of human leader/computer follower: no "veto" buttons, pedals or cues'.

The above examples are only a few taken to exemplify a field of research which is very active, diverse and encompassing different concepts what interactivity could or should be. In the broader context, typically associated terms like control, instrument, influence, sensing etc. contribute to a greater loss of focus in of the meaning of interactivity: a step back is needed. A critical look at the majority of those practices reveal the dominant model of interactive systems as that of a reactive or responsive system.[30] If we were to take its meaning literally, from the Oxford English Dictionary (2000):

> *The prefix inter- [meaning] Between, among, mutually, reciprocally. Interact [meaning to], act reciprocally or on each other Interaction a noun, [meaning to] blend with each other*

[24] Claude Cadoz and Marcelo M. Wanderley. Gesture - music. In M.M. Wanderley and M. Battier, editors, *Trends in gestural control of music, Paris, IRCAM/Centre Pompidou*, 2000; Marcelo M. Wanderley and Philippe Depalle. Gestural control of sound synthesis. *Proceedings of the IEEE 2004*, 92(4):632 - 644, April 2004; Dylan Menzies. Composing instrument control dynamics. *Organised Sound*, 7(3):255-266, 2002. ISSN 1355-7718; and Andy Hunt, Marcelo M Wanderley, and Ross Kirk. Towards a model for instrumental mapping in expert musical interaction. In *Proc. of the 2000 International Computer Music Conference*, pages 209-211, 2000

[25] Tellef Kvifte. On the description of mapping structures. *Journal of New Music Research*, 37(4):353-362, 2008

[26] Joel Chadabe. The limitations of mapping as a structural descriptive in electronic instruments. In *Proceedings of the 2002 conference on New interfaces for musical expression*, pages 1-5. National University of Singapore, 2002

[27] Robert Rowe. *Interactive music systems: machine listening and composing*. MIT press, 1992; and Todd Winkler. *Composing Interactive Music*. MIT Press, 1998

[28] Garth Paine. Interaction as material: The techno-somatic dimension. *Organised Sound*, 20 (1):82-89, 2015

[29] George E. Lewis. Too many notes: Computers, complexity and culture in "voyager". In *Leonardo Music Journal*, volume 10, pages 33-39, 2000

[30] Garth Paine. Interactivity, where to from here? *Organised Sound*, 7(3):295-304, 2002

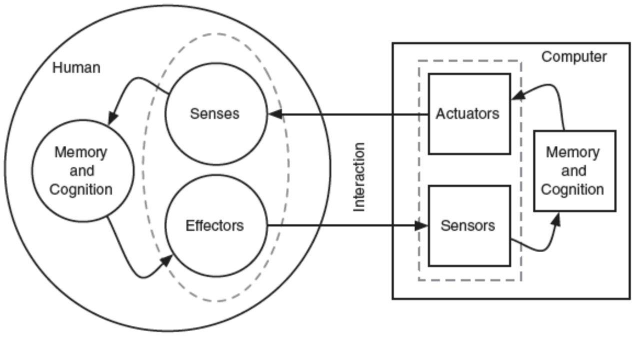


Figure 2.2: B. Bonger's model for Human-Computer interaction. Graphic from: Jon Drummond. Understanding interactive systems. *Organised Sound*, 14(2):124-133, 2009

Avoiding questions of etymology, the definition implied by the previously cited formulation by Chadabe, formed through praxis, refers to a different situation. Most approaches, when confronted with the actual qualities of the systems these use, fail in generating a situation of mutual influence of reciprocity and implicitly fall into paradigms of unidirectional control-effect flows.

Drawing ideas from *cybernetics*, the transdisciplinary scientific discipline which studies the structure and behaviour of regulatory systems, the communication between performer and computer music system may be formulated differently. Breaking the unidirectional relationship, this alternative formulation transforms that relationship into a mutually influential *closed loop* of communication, forming the essential basis for interaction. Bongers, in particular, sees the failure of typical approaches to interaction in realising this closed loop in the inattention given to the feedback channel from system to performer: this channel lacks a proper specification and is not actively employed.[31] Bongers sees the solution in providing 'a level of *cognition*'[32] to the computer system. Apart from the perhaps naive and schematic understanding of human perception and cognition, it is interesting to note how, in this model (see figure 2.2), the solution is to further *humanise* the computer. In order for the computer system to properly interact with a human it has to *act* and *think* like a human. What can be seen in Bongers' approach is an important change in the discourse about interaction in computer music: in order to build interactive systems that have those qualities it is necessary to understand how human cognition functions. *Cognitive sciences* now enter the field. The term cognitive sciences defines a highly interdisciplinary research field focusing on the mind and its processes: a field in which psychology, philosophy, neurosciences and artificial intelligence cross. These disciplines were already influencing *HCI* long before Bonger's publication. As an example, Don Norman's

[31] Jon Drummond. Understanding interactive systems. *Organised Sound*, 14(2):124-133, 2009

[32] Bert Bongers. Physical interfaces in the electronic arts. In M.M. Wanderley and M. Battier, editors, *Trends in gestural control of music, Paris, IRCAM/Centre Pompidou*, pages 41-70. IRCAM-Centre Pompidou, 2000

well-known influential works in HCI draw from the field of psychology and in particular from Gibson's theories of of ecological psychology.[33]

[33] Donald A Norman. *The psychology of everyday things.(The design of everyday things).* Basic Books, 1988

It seems that in order to create interactive situations or instruments, questions regarding what *interactivity* actually means and what it affords have to be asked. What do we mean by 'mutual' interaction? How could such mutuality be concretely addressed and generated? Are there other formulations of interaction and interactivity? In order to shape interaction between humans and computers a theory is needed that addresses what interaction means to us, what role it plays in our perception and cognition and the role of the body.[34] The embodied cognition theory is that theory.

[34] Dag Svanæs. *Understanding interactivity: steps to a phenomenology of human-computer interaction.* PhD thesis, Norges teknisk-naturvitenskapelige universitet, 2000

## *2.3 Embodiment*

The roots of *embodiment* theories may be traced back to philosophy and in particular to *phenomenology*. Phenomenology refers to the method developed by philosopher Edmund Husserl of finding the essentials of consciousness or of perception. While focusing on those subjective aspects, the method follows a path of *reduction* that would lead to the 'thing itself'. What is observed should be freed from all prejudice and pre-formed conceptions which could distort the image of the observed phenomenon. Phenomenology attempts an *objectification of the subjective*, doing so without eliminating the latter. Instead, it puts a systematic reflection about the world 'as we live it' - that is as we perceive and experience it - at its centre. This stands in contrast the Cartesian methodological tradition, in which the world is a collection of objects and relations between them, detached from the subject. The Cartesian subject gains knowledge about the world only through an abstract, immaterial reasoning process. At its core, phenomenology calls for an alternative worldview which surpasses dualistic body/mind, reason/matter, theory/practice perspectives, which establish the primacy of one over the other.

Heidegger's works develop phenomenology further. He does so with a critique of the tendency towards abstraction he sees implicit in Husserl's thought. For Heidegger, Husserl places experience *in the head* thus retaining a sort of 'mentalistic' model of perception and indirectly reaffirming the primacy of theory (the abstract) over practice. This contradicts the basis of the phenomenological method. Heidegger, rather, holds that experience is something that happens 'in the world'. In his major work *Being and Time*[35], he argues that we gain access to the world through our practical involvement with it and, further, the relation

[35] Martin Heidegger. *Being and time: A translation of Sein und Zeit.* SUNY press, 1996

with the world we construct must not be based on pre-formed *mental images* or *models* of it. We engage with the world through what he calls the *ready-at-hand* objects which do not appear to us as such but that we use without reflecting on them. The example he uses is that of our use of a hammer in the act of hammering. In that situation the tool is an extension of our body; it does not have a separate nature, but rather almost disappears as we are engaged in the action. It is only in the event of a *breaking down*, for instance the hammer doesn't work properly anymore, that the object is suddenly recognised and presents itself to our perception - it becomes *present-at-hand*. But this is not just a revelation of the object, it is the moment in which the object comes into existence. The very ontological structure of the world is therefore not pre-given, but arises through interactions implying embodied actions and their breaking down. The importance of his thought is acknowledged in the field of HCI: for example Winograd and Flores adopted these ideas, directly contradicting the dominant models that were influencing computational theories of cognition circulating in computer science at that time (the 1980s).[36]

But, the central philosophical work regarding embodiment is the *Phenomenology of Perception* by Merleau-Ponty.[37] At the beginning of his work he rejects the idea of a perception as passive reception of stimuli. In his view, perception is an active process *we perform*: 'sense-experience is a vital process, no less than procreation, breathing or growth'. Here sensations are not states of mind, but are rather emerging from an ongoing process of access to the world through movement and the active use of our senses. Further, 'the thing is inseparable from a person perceiving it. . . . To this extent, every perception is a communication or a communion', opposing the view (popular in psychology and computational cognitive sciences) that the brain functions as a processor of some 'data' passively received by the senses. Merleau-Ponty goes on to suggest that, firstly, there is no perception without action, and secondly, that perception and the 'information processing' (cognition) function of the brain cannot be considered separately: they are, in other words, intertwined. The nexus of this interconnection is the body, meaning that Merleau-Ponty's theory of perception is also a theory of the body. The body occupies, in his theory, a special place. It is neither simply an object 'among all others', in contrast to classical psychology, nor it is completely internalised into consciousness. Rather it is the body *as lived* or, as he calls it, the *phenomenal body* or *corps propre*. As identified in Dreyfus' analysis, three aspects of

[36] Terry Winograd and Fernando Flores. *Understanding computers and cognition: A new foundation for design*. Intellect Books, 1986

[37] Maurice Merleau-Ponty. *Phenomenology of Perception*. Routledge, 2002

*embodiment* contribute to the construction of the phenomenal body: first, the physical embodiment of a human subject's body having a specific shape; second, a set of bodily skills developed and acquired by the subject; and third, the cultural and social skills gained as the subject is embedded in a cultural world.[38] The *phenomenal body* is in itself not static, but rather a dynamic entity, equipped with the structural flexibility that allows us to learn and acquire new skills, to adapt to the external world through an active 'incorporation', as Merleau-Ponty writes, while describing how an organist learns to play a new organ: the new instrument becomes part of the experienced body, extending it.

[38] Hubert L Dreyfus. The current relevance of Merleau-Ponty's phenomenology of embodiment. *The Electronic Journal of Analytic Philosophy*, 4:1-16, 1996

Building on Merleau-Ponty's highly influential work, theories of embodiment have been an object of research and further developed in philosophy and cognitive sciences. In particular, Varela, Thompson and Rosch, state in their proposition of an *enactive* cognition theory:[39,40]

[39] Eleanor Rosch, Lydia Thompson, and Francisco J Varela. *The embodied mind: Cognitive science and human experience*. MIT press, 1991

[40] Here the term *sensorimotor* refers to the coupled sensory (the set of sense and their physiological mechanisms) and motor (the set of all movement actuators in the human body) systems.

> *By using the term embodied we mean to highlight two points: first that cognition depends upon the kinds of experience that come from having a body with various sensorimotor capacities, and second, that these individual sensorimotor capacities are themselves embedded in a more encompassing biological, psychological and cultural context.*

Emphasis is placed on the embeddedness or *situatedness* of cognition, descending from the assumption that perception and cognition are interrelated and determined by action and interaction with the physical world. Knowing is situated in the physical or social or cultural context where the action takes place. Andy Clark develops this even further: for him, cognition emerges from 'continuous reciprocal causation', that is from the 'continuous, mutually modulatory influences linking brain, body, and world'[41] He sees cognition and the mind-body as extending through its continuous interaction with its environment, an 'extended mind'[42].

[41] Andy Clark. *Being there: Putting brain, body, and world together again*. MIT press, 1998

[42] Andy Clark and David Chalmers. The extended mind. *Analysis*, 58 (1):7-19, 1998

Embodiment is therefore, following these thoughts, strongly linked to interaction. Interaction is a continuous engagement and exchange with the world, and is the basis for perception and cognition. Interaction is shaped, performed and sensed by our body: it is embodied. The importance of embodiment for the HCI field has been recognised in particular by Paul Dourish, who set out to define an 'embodied interaction'. Taking into account the phenomenological approach to perception and the developments in social and tangible computing, Dourish tries to develop an approach to interaction design based on essential physical and bodily aspects, but also addressing the social and cultural context interaction is embedded into.[43]

In the context of computer music, embodiment has been widely understood in relation to an acknowledged lack of bodily presence in electronic music and computer music performance practices.[44] This perspective is largely based on an idealised view of traditional music practice, found in classical music or instrumental music in general (jazz, rock, pop etc.). Comparing these practices with those common in computer music, a different, less bodily involvement of the musician/performer is evident, and this difference is seen as the origin of an *expressive* problem inherent to electronic music.[45] That is, simply put, since computer music performance lacks bodily engagement it is fundamentally less expressive. Expressivity therefore becomes a central theme in computer music interaction design: interactive instruments and interfaces are sought through which the musician could enter a more embodied relationship with the sound the computer music system produces: these instruments would therefore allow for an enhanced musical expression.[46] This strand of development is fuelled by music cognition research, in particular by the work of Rolf Inge Godøy on *motor mimetic cognition* and Marc Leman on *embodied music cognition* in systematic musicology.[47] Music production and performance (the model here is again traditional acoustic music) is seen here as intimately and necessarily connected to embodied perception, while music reception implies a bodily engagement as 'a process of incessant mental re-enactment of musical gestures'[48]. *Musical gestures* are, in this framework, the vehicle through which music is 'performed and perceived', and 'can be directly felt and understood through the body, without the need of verbal descriptions'[49]. As a consequence, gestures are therefore understood as a fundamental aspect of performance that computer music interfaces should be able to sense and transpose or map into sound synthesis parameters.[50] That is, bodily gestures enable a more embodied receptionof electronic and computer music, beyond a 'disembodied, mentalesque engagement'[51].

I am critical of this use of embodiment and the underlying assumptions these approaches imply. As the above statement by Marc Leman about a 'disembodied' musical cognition exemplifies, these approaches to embodiment assume, often only implicitly, that electronic music as 'mental' or 'abstract' and therefore, in general, affording a disembodied engagement. This thought actually reinforces a division between the mind and the body, and are therefore in opposition with the most important characteristic of embodiment (already seen in the foundations of

[43] Paul Dourish. *Where the Action is: The Foundations of Embodied Interaction*. The MIT Press, 2001

[44] Bob Ostertag. Human bodies, computer music. *Leonardo Music Journal*, 21:19-23, 2006

[45] Michael Gurevich and Jeffrey Treviño. Expression and its discontents: toward an ecology of musical creation. In *Proceedings of the 7th international conference on New interfaces for musical expression*, pages 106-111. ACM, 2007

[46] Jin Hyun Kim and Uwe Seifert. Embodiment and agency: Towards an aesthetics of interactive performativity. In *Proceedings of the 4th Sound and Music Computing Conference*, pages 230-237, 2007; and Garth Paine. Towards unified design guidelines for new interfaces for musical expression. *Organised Sound*, 14 (2):142-155, 2009

[47] Marc Leman. *Embodied music cognition and mediation technology*. MIT Press, 2008

[48] Rolf Inge Godøy. Motor-mimetic music cognition. *Leonardo*, 36(4): 317-319, 2003

[49] Marc Leman. Music, gesture, and the formation of embodied meaning. In Rolf Inge Godøy and Marc Leman, editors, *Musical gestures: Sound, movement, and meaning*, pages 126-153. Routledge New York and Abingdon, England, 2010

[50] Marcelo M Wanderley. Gestural control of music. In *International Workshop Human Supervision and Control in Engineering and Music*, pages 632-644, 2001; and Eduardo Reck Miranda and Marcelo M Wanderley. *New digital musical instruments: control and interaction beyond the keyboard*. AR Editions, Inc., 2006

[51] Marc Leman and Rolf Inge Godøy. Why study musical gestures. In Rolf Inge Godøy and Marc Leman, editors, *Musical gestures. Sound, movement, and meaning*, pages 3-11. Routledge New York, NY, 2010

phenomenology): the overcoming of the mind/body division. This paradox can be found also in Leman's most important work, *Embodied Music Cognition and Mediation Technology*. In particular, 'Leman's account for an action-oriented approach, based on the notion of corporeality is, in fact, supposed to overcome the problem of dualism, but its aim is to provide an epistemological foundation for bridging the gap between musical mind and matter intrinsically contradicts its own assumptions'[52] by falling back into a dualistic perspective. Further, there seems to be a strong underlying assumption that the paradigms of traditional acoustic instrumental practices should be equally applied to computer music instruments: an assumption which cannot be fully justified. There are, of course, good reasons in applying those paradigms: instrumental music has a long history, evolving over time, and therefore much could be appropriated or learnt from it. Computer music practice, however, can originate other, equally valuable paradigms leading to different performative situations. Further, even if musical *expression* is at the core of much of the research and development in interactive interfaces for computer music, it is still not entirely clear what its contents should be - e.g. communication, meaning, emotion, articulation or even 'style'. Is expression even relevant to contemporary music practices at large? It surely is not in the works of Cage, or Xenakis, or Brün or Morton Feldman.[53]

Another understanding of embodiment in computer music is that of an *embodying something*. This draws from Don Ihde's concept of *embodied relations*.[54] Ihde's ideas address relations between humans and technology. In particular, relationships in which artefacts become means for perceiving, encountering and interacting with the world. These relations may be embodied in that the technology does not become evident to perception: rather it is a transparent means through which the environment is explored. Technological objects such as glasses, hearing aids, a blind person's cane, or a hammer enter into symbiosis, in which our perception extends into the artefact itself. Ihde's work is clearly rooted in Merleau-Ponty, in particular in the idea of the embodied perception *extending* into the technological artefact, which at that moment becomes part of our body. Approaches employing this understanding to create interactive computer music environments attempt to create the conditions for an embodied interaction of this kind.[55] On the base of this interpretation of embodiment, researchers and composers of computer music try to develop interactive instruments that can completely dissolve in their interaction with the performers, being fully inhabited

[52] Andrea Schiavio and Damiano Menin. Embodied music cognition and mediation technology: a critical review. *Psychology of Music*, 41(6):804-814, 2013

[53] Michael Gurevich and Jeffrey Treviño. Expression and its discontents: toward an ecology of musical creation. In *Proceedings of the 7th international conference on New interfaces for musical expression*, pages 106-111. ACM, 2007

[54] Don Ihde. *Technology and the lifeworld: From garden to earth*. Indiana University Press, 1990

[55] Garth Paine. Interaction as material: The techno-somatic dimension. *Organised Sound*, 20 (1):82-89, 2015; Gerhard Eckel. Embodied generative music. In Deniz Peters, Gerhard Eckel, and Andreas Dorschel, editors, *Bodily expression in Electronic Music*, chapter 10, pages 143 - 151. Routledge, 2012; and Gerhard Eckel and David Pirrò. On artistic research in the context of the project embodied generative music. In *Proceedings of the 35th International Computer Music Conference*, pages 541-544, Montréal, 2009

and embodied in the perception of the musicians interacting with them. The utopian thinking behind such an approach is again that of total control of the instrument by the musician; an impulse which, in my opinion, still does not reflect the relation between mutually influential entities. In this case, one of the entities is 'absorbed' by the other, who then asserts complete control. If a state of mutuality is to be reached, the 'other' or counterpart should remain visible, perceivable and almost graspable.

So how is the concept of *embodiment* relevant to computer music? How could it be effectively employed in addressing the idea of a mutual interrelation that I am placing at the core of my proposed interaction model? The difficulty is that the theory of embodiment describes properties of human perception and cognition which are *inherent*. Embodiment is a *permanent mechanism* of our perceptual system: *disembodiment* could not really exist. It is difficult, therefore, to talk about an interface or computer music system, or even a kind of sound, which could be in some way disembodied. At the core of embodiment as a mode of thought lies a perspective that is deeply rooted in the perceiving subject. The perceiving subject is the central actor in this theory: there are no objects detached from this subject. It seems, therefore, inappropriate to identify an 'external' object that has an ontological quality of being disembodied *a priori* independently of the perceiving subject. Every object is defined in relation to the perceiving subject. Following this logic, there is actually no possibility for a disembodied relation to occur between human and technological artefact, could this mean that embodiment is actually entirely meaningless for interaction design? As Dourish puts it:[56]

[56] Paul Dourish. *Where the Action is: The Foundations of Embodied Interaction*. The MIT Press, 2001

> *If we are all embodied, and our actions are all embodied, then isn't the term embodied interaction in danger of being meaningless? How, after all, could there be any sort of interaction that was not embodied? What I am claiming for embodied interaction is not simply that it is a form of interaction that is embodied, but rather that it is an approach to the design and analysis of interaction that takes embodiment to be central to, even constitutive of, the whole phenomenon.*

The value of embodiment for interaction design is that it reveals aspects that should be put at the very centre of such design process, i.e. the qualities of the physical body and how we may enter in relation with the world through it.

At this point, before continuing, I think it's useful to look back and understand how we have arrived here. We found that, in order to create a situation of mutual influence between a performer and a computer music system, venturing into theories of cognition and perception have provided a better understanding of *how* that influence is exerted and received. In particular, a better and stronger specification of the 'feedback channel', the connection that links the system to the performer. We looked at how human cognition and perception function for the purpose of equipping a system with similar capabilities, to some extent imitating them in the computer system. We also encountered the theory of embodiment which seems apt as a basis for such an endeavour, in that it explicitly puts the body and its continuous interactions with the world at the centre of cognitive functions. What is needed at this point, though, is a perspective that focuses on the qualities of such interactions; qualities that could be exploited by a computer music system to enter in a continuous, mutual interaction with the performer.

An exchange between philosophers Deniz Peters and Alva Noë in the book *Bodily expression in Electronic Music* - published out of the Embodied Generative Music project - is especially interesting here.[57] The discussion revolves around the question of the possibility of a disembodied electronic music. Following similar pathways, both interlocutors conclude in the negative: embodiment always centres on the perceiving subject, but they stress the fact that this is not a just passive reception of information. The crucial point is, recalling Merleau-Ponty's *phenomenology*, that this perception is active, in the very real sense of performing an activity. Alva Noë, in particular in his book *Action in Perception*, stresses opposition to a projective idea of perception, in which our senses are subject to external stimuli that are then processed by our cognitive system. From this perspective the world we live in is also removed from us. He instead endorses the idea that we are in continuous contact with the world. In describing perception, he argues that 'seeing is like touching', where an object is picked up in the hands, moved around, its form, weight and surface felt. In other words, perceiving and bodily movement are interrelated: moving and acting on an object produces variations in that object, which we then reconstruct and integrate in perceiving it: 'to perceive is to exercise one's skillful mastery of the ways sensory stimulation varies as a result of bodily movement'[58]. The act of listening to a sound, therefore, requires an *action of perception* that is deeply bodily. Sound *is embodied*, in the sense that perceiving it requires bodily interaction.

[57] Alva Noë. What would disembodied music even be? In Deniz Peters, Gerhard Eckel, and Andreas Dorschel, editors, *Bodily expression in Electronic Music*, chapter 3, pages 53 - 60. Routledge, 2012; and Deniz Peters. Touch. real, apparent, and absent: On bodily expression in electronic music. In Deniz Peters, Gerhard Eckel, and Andreas Dorschel, editors, *Bodily expression in Electronic Music*, chapter 1, pages 17 - 34. Routledge, 2012

[58] Alva Noë. *Action in Perception*. The MIT Press, 2004

Perception is an activity, a process of bodily interaction with the world, and an *enactive* process.

## *2.4 Enaction*

The term *enactive approach* and the concept of *enaction* refers to the perspective towards cognition that Varela, Thompson and Rosch elaborated in their book *The Embodied Mind*[59]. It draws ideas from research areas as remote as biology, cognitive sciences, neurology, psychology and philosophy in order to construct an unifying theory of cognition. Nonetheless, its fundamental roots still lie in Merelau-Ponty's phenomenology, the idea of perception as action and, in particular, his theory of embodiment. In their words, the enactive approach views 'cognition as *embodied action* and so recovers the idea of embodiment'. From this perspective embodiment 'encompasses both the body as a lived, experiential structure and the body as the context or milieu of cognitive mechanisms'; in other words, as the locus both of the coupled sensory-motor system acting on the world and of higher cognition functions. But, *enaction* goes further still:

> *the enactive approach consists of two points:*
> *(1) perception consists in perceptually guided action and*
> *(2) cognitive structures emerge from the recurrent sensorimotor patterns that enable action to be perceptually guided.*

[59] Eleanor Rosch, Lydia Thompson, and Francisco J Varela. *The embodied mind: Cognitive science and human experience*. MIT press, 1991

As with Merleau-Ponty's phenomenology, perception here is an action performed by a perceiving subject. These actions then alter the local world the subject is embedded in, meaning it is impossible to refer to a pre-given and invariant external world. The sensory-motor system of the perceiver, and how that is embodied, is of utmost significance. The enaction approach asks how action is *guided* by a subject, with the aim of understanding the principles of the linkage between sensory and motor systems. How, for instance, is action influenced by continually changing sensory information? Important to this idea is that perception is not just situated and embedded in its environment, but it actively contributes to its enactment: 'the organism both shapes and is shaped by the environment'. It is a continuous process taking place between the subject and the environment: it is a circular relationship encompassing the actor's action, the change this produces in the environment and the sensing of this changes. Cognitive structures emerge from the recurring temporal patterns this process might elicit. Such cognitive structures are the basis for understanding (e.g. of cause-effect relationships) and allow, in turn, to act on the link between the sensory and the motor systems, modulating it and therefore guiding perception.

Enaction therefore proposes an approach to cognition which is fundamentally opposed to *representationist* approaches that view cognition as merely information processing, in which data is provided by a pre-given external world. Instead, enaction views cognition as a highly dynamic, temporal process arising from continuous embodied action in an environment. From the enactive perspective, perceiver and environment are two coupled and mutually interacting systems. Action is necessary, but more fundamentally so too is *being acted upon* - action from the environment is necessary in order to perceive and construct cognition structures. Essentially 'living beings and their environments stand in relation to each other through mutual specification or co-determination'[60]. This co-determination of organism and environment is central to the concept of enaction.[61]

The concept of enaction has received little attention in the context of computer music. Isolated examples that do make use of it in interface design acknowledge perceptual sensory-motor coupling, but remain mired in their instrumental perspective. A performer could more *intimately control* a computer music instrument using a *transparent* interface, an interface capable to tap into enactive qualities of perception.[62] As I explained above, the central idea of the enactive perspective is that of a structural coupling of two systems - in the present case, between the human performer and computer music system - which engage in a mutual interaction based on a continuous sensory-motor engagement. I believe that this correlates with the ideal of interaction I have been describing. In order to perceive the computer music system, the performer must both continuously sense it and act upon it. The variations of the system's responses to the performer's movements, if exhibiting recurring patterns, would allow cognition to 'resonate' into a coherent image. Cognition would, in other words, *attune*[63] to it. The paradigm of touching, as used by Alva Noë, might be useful here: to perceive the system the performer has to touch it, move it, weigh it. Another useful concept in this respect is *resistance*[64], a sensible quality felt while interacting with the system. This might be a system's resistance towards the performer's actions, signalling to the user that 'something is there to explore', thus revealing the form and characteristics of that system.

Another concept which could be useful, and one that pairs well with resistance, is that of *affordance*. The

[60] Eleanor Rosch, Lydia Thompson, and Francisco J Varela. *The embodied mind: Cognitive science and human experience*. MIT press, 1991

[61] Evan Thompson. *Mind in life: Biology, phenomenology, and the sciences of mind*. Harvard University Press, 2010

[62] David Wessel. An enactive approach to computer music performance. *Le Feedback dans la Creation Musical,Lyon: Studio Gramme, France*, pages 93-98, 2006

[63] The concept of *attunement* is here borrowed from Merelau-Ponty and then reprised by Thomson in *'Mind in Life'* as it fits well the idea of resonance. Though, I will not give a complete definition of the term here.

[64] Newton Armstrong. *An enactive approach to digital musical instrument design*. PhD thesis, Princeton University, 2006

concept was first introduced by psychologist J.J.Gibson in his *ecological approach* to psychology.[65] He elaborated a perspective on perception in which the role of the environment is central. In the ecological approach, the environment offers opportunities for interaction relative to the sensory-motor capabilities of the perceiving subject. Affordances are thus ecological features of the world of things that elicit actions. A mug's handle, for instance, affords a specific motor action for picking up the mug. Affordances are the 'motor sense' of objects in the environment.

[65] James J Gibson. *The ecological approach to visual perception: classic edition*. Psychology Press, 2014

Let's return to the question of how a computer music system should be designed in order to be interactive (as defined above), while also incorporating an enactive approach. We could say that the system should, at the same time:

- Present *resistance* against the performer's actions, making sensible the limits or constraints in which the interaction takes place. This also affirms its boundaries and identity as a perceivable object.
- Offer *affordances* by eliciting responses by the perceiver's sensory-motor system. Using the mug's metaphor, the system provides and exposes the handles having the correct dimensions and shape for being grasped and used.

So how can these qualities be realised? The computer music system should *attune* to the performer's perception and cognitive system's structure. And, in order to elicit a sense of resistance and affordance, should also be built around a *model* that, in its functioning, resonates with the human's enactive perception and cognition processes. Fortunately, within the theory of enaction, such a model can be found in the definition of *agency*.

The a year-long collaboration between Francisco Varela and Humberto Maturana culminated in their book *The Tree of Knowledge*[66], one of the most significant texts for this theory; this text shows how the enactive approach is strongly rooted in biology and neurosciences. Their research focused on finding the biological roots of understanding, that is, those basic mechanisms which are the foundations of knowing and cognition in living beings. Cognition is an essential quality of living organisms, and the presence of cognition can therefore be used to define a *living entity*. This perspective establishes a *circularity* between the concepts of 'living' (and living organisms) and 'cognition': that which lives, cognises and that which cognises, lives. In the words of Maturana[67]:

[66] Humberto R Maturana and Francisco J Varela. *The tree of knowledge: The biological roots of human understanding*. New Science Library/Shambhala Publications, 1987

[67] Humberto R Maturana and Francisco J Varela. Biology of cognition. In *Autopoiesis and cognition*, pages 2-58. Springer, 1980b

Figure 2.3: Diagram illustrating the definition of agency: redrawn from Xabier E Berandiran, Ezequiel Di Paolo, and Marieke Rohde. Defining agency: Individuality, normativity, asymmetry, and spatio-temporality in action. *Adaptive Behavior*, 17(5):367-386, 2009

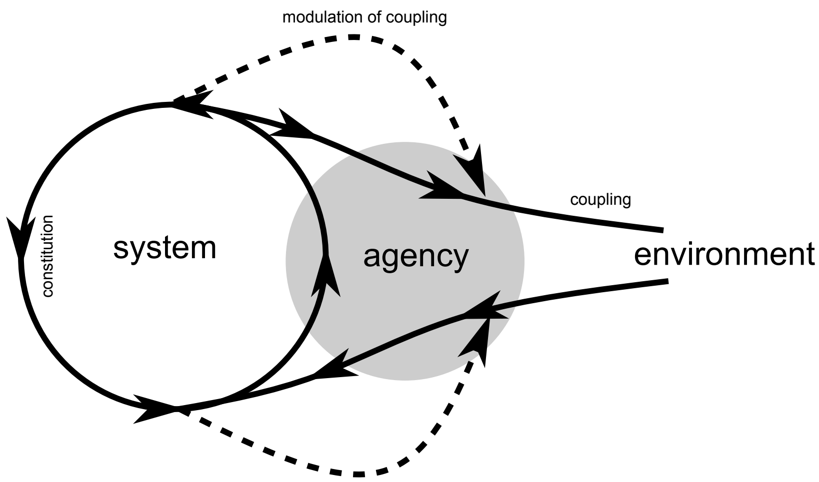


*A cognitive system is a system whose organisation defines a domain of interaction in which it can act with relevance to the maintenance of the self and the process of cognition is the actual acting or behaviour in this domain. Living systems are cognitive systems, and living as a process is a process of cognition. This statement is valid for all organisms, with and without a nervous system.*

This statement suggests that cognition is a process of acting, or *doing* - a perspective that resonates with the enactive theory. Living beings are characterised as being agents (to be understood in the Latin sense of *agens*, or doer) of cognition. For Evan Thompson, the concept of agency is at the foundations of enactive theory: the idea of enaction is 'that living beings are autonomous agents that actively generate and maintain themselves, and thereby also enact or bring forth their own cognitive domains'[68]. Further 'a cognitive being's world is not a prespecified, external realm, represented internally by its brain, but a relational domain enacted or brought forth by that being's autonomous agency and mode of coupling with the environment'. Agency and enaction are, therefore, very closely related and even interdependent.

[68] Evan Thompson. *Mind in life: Biology, phenomenology, and the sciences of mind*. Harvard University Press, 2010

In the first formulation of the enaction theory by Maturana and Varela, agency is strongly related to the concepts of *autonomy* and the *adaptivity* of the living system. Autonomy designates the the system's ability to *self-organise* and self-specify - what Maturana called the *autopoiesis*[69] of the living system. Adaptivity postulates a *coupling* between the system and its environment and defines the capacity of systems to regulate this connection with the environment. Adaptivity is a function that calibrates the agent's action and perception processes to external stimuli.

[69] Humberto R Maturana and Francisco J Varela. Autopoiesis: the organisation of the living. In *Autopoiesis and cognition*, pages 73-135. Springer, 1980a

Autonomy is certainly a necessary condition, but only in conjunction with adaptivity can it also become a sufficient condition for agency.[70] A more detailed characterisation has been provided by Barandiaran and Di Paolo:[71] following this work, agency can be defined as three different (but interrelated) conditions a system has to meet (see figure 2.3):

[70] Ezequiel A Di Paolo. Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the cognitive sciences*, 4(4): 429-452, 2005

[71] Xabier E Barandiaran, Ezequiel Di Paolo, and Marieke Rohde. Defining agency: Individuality, normativity, asymmetry, and spatio-temporality in action. *Adaptive Behavior*, 17(5):367-386, 2009

1. INDIVIDUALITY: For a system to be an agent, there must be some distinction between the system and the environment. The agent must have clear boundaries and there must be some clear relation between the system and its environment. It is an entity identifiable from the perspective of the environment.

2. INTERACTION ASYMMETRY / SOURCE OF ACTIVITY: This concept is related to *action*: the agent *does* something, and is a source of activity rather than a passive receiver of external effects. The type of coupling between agent and environment is not just one of *reaction*, but rather a coupling through which the agent acts by some internal, local and individual mechanisms. This becomes evident the moment in which the agent re-modulates its coupling to the environment from within, therefore breaking the symmetry of two coupled systems.

3. NORMATIVITY / ADAPTIVITY: The coupling with the environment is modulated in order to move the system's state towards a specific *goal*, here called a *norm*. This modulation might result in *success* or *failure* in achieving that goal: this is what is defined as the *normativity condition*. *Failure*, or the possibility of it, is a central characterising quality of agency - for instance, the planetary system cannot fail to follow the laws of gravitation and is therefore not an agent. The specification of a system's goal or aim might seem odd, dependent as it is on the perspective from which the dynamics of system and environment is observed. What is actually meant by this, is that an agent system should tend towards maintaining its *norm*, meaning it should work towards preserving its further activity. The norm is its continued existence.

These specifications clarify one important point. An enactive agent is not only both autonomous and coupled to its environment, but, most importantly, it has the faculty to adapt by re-calibrating this coupling. The system must be able to *observe its own state* and adapt its coupling with its environment in such way that their interaction would pull that state towards regions more apt for its functioning and thus securing its continued existence. I believe that this quality is fundamental and constitutes an essential and discriminating trait of an agent.

Now we have a characterisation of the fundamental and defining qualities of an agent. As Evan Thompson suggests in his writing about empathy towards other (human) living organisms: 'we perceive her [. . .] as a locus of intentional agency and voluntary movement'[72]. This suggests that agency is not only an essential ingredient for an enactive organism, it is also a perceptual quality by which agents recognise and identify each other. So, how can one 'equip' a computer music system with the same degree of agency for the performer to recognise the system as an agent with which an enactive interaction is possible. The aim is to realise a system whose behaviour correlates with the actions of the performer and has such perceivable characteristics, that *agency will be perceived or re-constructed by a performer through enactive interaction*. Essentially, agency as the process that possess the *resistance* and the *affordance* to be perceived as an autonomous and interactive process.

[72] Evan Thompson. *Mind in life: Biology, phenomenology, and the sciences of mind*. Harvard University Press, 2010

The pressing question is, now, how can such a system be realised? With which tools and using which formalism? For possible answers one may look to the mathematical formalism of *Dynamical Systems*: within the theory of enactive perception, dynamical systems appear in various forms; in some places as metaphors, in other as precise mathematical formulations. The language of dynamical systems will be the language I will use to deal with the questions I delineated above. The next chapter will provide a short introduction in the theory and use of dynamical systems.

# 3

# *Dynamical Systems*

*Dynamical systems* are mathematical formulations describing the time dependence of an object or an ensemble of objects in a geometrical space: they are constructs which describe temporal evolution. They do so by employing mathematical statements, or equations, that express how an object will change its state given its present state or the state of other elements in the system. These are largely rules of temporal evolution which act as a 'force' on those objects, pushing them to move on a path. The word 'dynamical' is, illuminatingly, etymologically rooted in the Greek *Dynamis* (δυναμισ), meaning 'force'.

The above description already hints at how dynamical systems could address an enormous variety of phenomena. Any process that has a temporal evolution can, in principle, be expressed in terms of dynamical systems. Since most (if not all) phenomena we are confronted with will have a temporal dimension, dynamical systems are therefore entirely ubiquitous. Such broadness and lack of specificity risks losing meaning altogether. If everything in the end is a dynamical system, how is the concept useful?

I believe that the phrase 'XY is a dynamical system' is so indefinite that it contains no information about the thing itself. Still, there is something implied by such affirmation which is important to underline. 'XY is a dynamical system' means that I am looking at XY as a temporal phenomenon, its *temporality*, and I understand its evolution as a fundamental quality. That phrase pertains less to what XY *is* than an indication of *how* I am looking at it, what kind of perspective I have chosen to examine it. This is a perspective based on time and interaction, indicating what kind of language I have chosen to use in formulating my thoughts.

The next section contains a short theoretical and lightly mathematical framing of dynamical systems theory. The mathematical study of dynamical systems is a large field, and much lies beyond the scope of this work: this is meant only as very brief introduction. The next sections will describe how dynamical systems have entered the research fields of perception and cognitive sciences and, further, which approaches to computer music have been influenced by this perspective. These sections are meant to justify my extensive use, later in this work, of formulations based on dynamical systems theory.

## 3.1 Theory

I will begin this section by attempting to clarify the definition of *dynamical system*, as I understand it, with a dissection of the concept into two.[1]

Here the term *dynamical* or *dynamic* has to be understood by its meaning most common in physics: as indicating the quality something has by being in motion, and therefore exhibiting some sort of temporal evolution. *Dynamic* is used to denote phenomena showing patterns of temporal evolution at one time which are interrelated to those at different times.[2] With this meaning *dynamic* becomes almost a synonym for time-evolution or pattern of change, referring to the unfolding of events in an continuous evolutionary process.

Most of the phenomena we are confronted with in our daily lives have some dynamic aspect, both physical and social. Simple physical systems include moving objects like a kicked football travelling through the air under the effect of the gravitational and other forces. Complex social systems could include those found in hierarchical organisations or evolving economical structures. Dynamics are a pervasive quality of what we perceive.

The term *system* is used in diverse contexts that furnish the word with different meanings. It can therefore be very unspecific. Returning to its Greek origin, συστημα, 'a whole composed by several parts', we may think of this word as denoting an identified set of elements linked by mutual connections and interactions. This network of interrelations is responsible not only for the appearance of the set as a whole, but also, in the case of *dynamical systems*, of its evolution. These interactions are responsible of the particular form this unfolding takes: the *behaviour* of the system.

*Dynamical systems* (or, more frequently, the reduced form *dynamics*) not only stand for the time-evolutionary phenomena the world presents us with, but also for the mathematical discipline that attempts to formulate and analyse such phenomena from an abstract, general perspective. This branch of mathematics has a history of almost 350 years: it has its origins in physics, but over time it has spread into various research fields, including chemistry, social sciences, biology, communications engineering and others.

[1] For a good and complete introduction to dynamical systems theory, I suggest referring to a more extended treatment as can be found in the book by Steven Strogatz. Some examples I report here are taken from that book.
Steven H. Strogatz. *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering.* Westview press, 2014

[2] D.G. Luenberger. *Introduction to Dynamic Systems: Theory, Models, and Applications.* Wiley, 1979

This progression led to the development of a multitude of different mathematical and conceptual tools. However, each is grounded on the work and ideas of its two most prominent fathers:

- NEWTON was one of the main developers of *calculus* (together with Leibnitz) and the inventor of *differential equations*, a means of studying and analysing physical processes. With this tool he formulated the laws of motion and gravitation, and laid the foundations of modern physics. In particular, Newton's fundamental second law of motion established a relationship between the force $F$ affecting an object of mass $m$ and the acceleration $a$ the object experiences as effect of the force:
$$F = ma \tag{3.1}$$
As the acceleration is the second derivative of the object's position $x$ with respect to time,[3]
$$a = \frac{d^2x}{dt^2} \tag{3.2}$$
Since then, dynamical systems have been expressed in terms of differential equations: from the perspective of mathematics, the study of dynamics become almost synonymous with the theory of differential equations.

[3] Throughout this text and from here on, we will use the overdots notation to denote differentiation with respect to time $t$: $\dot{x} = \frac{dx}{dt}$ and $\ddot{x} = \frac{d^2x}{dt^2}$

Differential equations express some (physical) quantity's dependence on time. More precisely they formulate *how* that quantity will change its value with respect to time, or how the path of its variation will be shaped.

By way of example, we could consider the ideal spring-mass system - ideal, here, meaning frictionless system[4]. The system consists of a mass $m$ attached to a spring (see the diagram in figure 3.1 ). *Hooke's law* states that the force $F$ the spring exerts on the mass is proportional to the elongation or compression $x$ of the spring:

[4] The term *system* here is used as introduced above. It denotes the set of two elements, spring and mass, bound together by their mutual interaction.

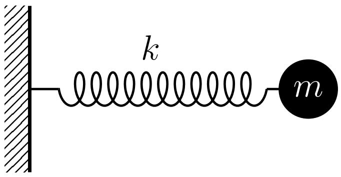

Figure 3.1: The diagram for ten classic spring-mass system

$$F = -kx \tag{3.3}$$

where the coefficient of proportionality $k$ is the *spring constant*. Following then Newton's law (equation 3.1) the system's dynamics can be mathematically expressed:

$$m\ddot{x} = -kx \tag{3.4}$$

The above describes how the acceleration of a mass $m$ is connected to the force exerted by the spring with the spring constant $k$. Or, in other words, how the velocity $v$ of the mass will change over time under the effect of the force exerted by the spring. Differential equations are the mathematical formulations of the laws of change and variations governing a system.

Differential equations were applied in order to mathematically formulate a variety of mechanical problems, as well as to find solutions to those problems - the explicit mathematical formulae expressing the motion of the involved elements. For example, in the case of the above spring-mass system, the solution describing the mass's motion:

$$x(t) = A cos(\omega t - \phi) \tag{3.5}$$

where $A$ and $\phi$ are parameters dependent on the initial conditions of the mass when the time began. What position was it in, for instance, or what was its velocity when observation began. The above equation, containing as it does a periodic function *cos*, states that the motion of the mass is oscillatory with a fixed frequency $\omega$ and is a 'solution' as it expresses this motion - that is, how the position of the mass changes over time, uniquely in relation to the passing of time and a few constant parameters.

The most interesting problems for Newton's contemporaries were those concerning the evolution of the planetary system, the motion of planets and satellites. After Newton, physicists tried to solve the so-called 'three-body' problem, to find the laws of motion for the three objects of mass $m_1, m_2$ and $m_3$ operating under reciprocal effects of gravitational forces. Despite the relatively simple formulation and clear problem, this turned out to be *impossible* to solve. Still today, even with the mathematical tools we have since developed, we remain unable to formulate solutions to this problem.

The three-body example shows what I think is an essential characteristic of dynamical systems theories. The complexity of a problem - meaning the difficulty of understanding it or of finding solutions - is not proportional to the complexity of its formulation. Or, said in another way, dynamical systems theory can provide very simple formulations for very complex problems.

- Poincaré was extremely influential in the development of dynamical systems theory. His greatest contribution was introducing, in the late 1800s, a completely new methodology. He developed way of thinking that considered the *qualitative* aspects of a system's temporal evolution, prioritising these over *quantitative* questions. For example, when considering the aforementioned three-body problem, he would ask 'Is the system *stable* as the earth-sun system or will the objects eventually fly off to infinity?' rather than seeking a mathematical formula expressing the position of the objects at any time.

  In order to answer these more qualitative questions, he developed a powerful *geometrical* approach through which a system's properties (e.g. how it would generate motion or its temporal evolution) would appear as figures or images of its behaviour. The qualitative method is grounded on a visual understanding of a system's behaviour. It is based on the concept of *phase space* - a geometric space containing all possible states for a system. The dynamical system 'acts' on this space in such a way that each point in this space, each state, is 'pulled' towards another state according to the system's rules (its differential equations) of evolution. That is, the dynamical system acts as a sort of *flow* in this space, dragging and pushing states around resulting in certain trajectories, *phase trajectories* which geometrically depict the system's behaviour.

The critical point here is that this representation of a system's dynamics is *isomorphic* to the system's formulation in terms of its differential equations. The two are equivalent and interchangeable: the visual, geometric approach is not just an approximation, it *exactly* corresponds to the system.

Before turning to some more detailed examples of *phase space* geometrical representations, some more notes about the kind of problem I am interested in may be useful.

In the set of problems dealing with differential equations there are two big families that can be distinguished: the first is that of *ordinary differential equations* which involve only ordinary derivatives with respect to time. See, for instance, the equation for a damped oscillator:

$$m\ddot{x} + b\,\dot{x} + kx = 0 \tag{3.6}$$

In this kind of problem, *time* $t$ is the only independent variable. *Partial differential equations* constitute the other category. Here is the equation for heat:

$$\frac{\partial u}{\partial t} = \frac{\partial^2 u}{\partial x^2} \tag{3.7}$$

In this equation, both time $t$ and space $x$ are independent variables. My focus, however, is only on ordinary differential equations, since I am only concerned with temporal behaviour.

The most general formulation for a dynamical system as a system of differential equation is as follows:

$$\begin{aligned} \dot{x}_1 &= f_1(x_1, x_2, \ldots, x_n) \\ \dot{x}_2 &= f_2(x_1, x_2, \ldots, x_n) \\ &\vdots \\ \dot{x}_n &= f_n(x_1, x_2, \ldots, x_n) \end{aligned} \tag{3.8}$$

Firstly, there is set of variables $x_1, x_2, \ldots, x_n$, which could be anything from chemical concentrations, to planet positions, to populations of different species sharing the same ecosystem. Their temporal evolution, how these variables will change as time advances, is expressed by the functions $f_1, f_2, \ldots, f_n$. The values of these functions are dependent on, in principle, all other variables in the same system. The number of equations $n$ is the *order* of the system, also called *dimension* (if $n = 2$ the system is second-order and so on).

From the formulation above, our understanding of a *system* starts to become clearer. A system is composed of a set of variables whose evolution is interdependent on each other's state. All variables are *coupled*, their evolution being affected and affecting each other. The concept of dynamical system reveals itself here as one deeply rooted in a perspective of a world as interconnected, *mutually interacting* entities or actors. A world of elements and their interconnections, whose evolution is brought forth through their interactions. I see this as very close to an enactive position.

It should be noted here that the systems in equations 3.4 or 3.6 are formulated differently to that in equation 3.8, as the differential equations involve second derivatives. Generally, however, it is possible to make use of a formulation in which only first derivatives are involved. This is typically done with a simple change of variables and introducing a new variable in the system. For instance, in the case of the simple harmonic oscillator in equation 3.4, through introducing the second variable velocity $v = \dot{x}$, the system then becomes:

$$\begin{aligned} \dot{x} &= v \\ \dot{v} &= -\frac{k}{m}x \end{aligned}$$

This is a second order dynamical system of the form of equation 3.8. With the same change of variables equation 3.6 would become:

$$\begin{aligned} \dot{x} &= v \\ \dot{v} &= -\frac{k}{m}x - \frac{b}{m} \end{aligned}$$

Further, considering the three systems:

$$\begin{aligned} \dot{x} &= rx \\ \dot{x} &= x^2 \\ \dot{x} &= sin(x) \end{aligned}$$

another observation can be made. In the first of the two equations, on the right-hand side $x$ appears to the first power,in the second equation it appears to the second power and in the third it is an argument for a function. The first system, therefore, is said to be *linear* while the other two are *nonlinear*. Another example of a non-linear system is that of the pendulum:

$$\ddot{x} + \frac{g}{L} sin(x) = 0 \tag{3.9}$$

Nonlinear problems are typically very difficult to solve (especially when the system is made up of more components). Linear problems, even if they have many components, can be broken apart more easily, with the possibility of analysing each part separately, and putting them all back together again without losing sight of the global system: what is known as the *principle of superposition*. This is much more difficult with nonlinear problems: in general, breaking apart components is near to impossible. Still, nonlinear systems are significantly more interesting: most of the processes and phenomena we are confronted in our everyday life are nonlinear. Nonlinearity is the realm of *chaotic* behaviour and complex systems theory. In the world of dynamical systems, these are the hardest, or more often impossible, to solve problems.

I'd also like to point to another important distinction among these systems. Take, by way of example, the forced oscillator system:

$$m\ddot{x} + b\dot{x} + kx = cos(t) \tag{3.10}$$

This is different from the previously noted systems, in that it is explicitly dependent on time $t$. This kind of system is called *nonautonomous*, as opposed to the *autonomous* systems we have looked at until now. Where the latter is self-contained, the former is explicitly dependent on some 'external' influence. This must be exerted from some kind of mechanism that is not internal to the system itself, but from the environment in which the system is located. In the case of the forced oscillator, for instance, this was an external force. It should be noted that these systems can be very difficult to study and share some commonalities with nonlinear systems, but in some cases a change of variables might help us to better look at the problem. For example, in the above example of the forced oscillator, we can manoeuvre $x_1 = x$, $x_2 = \dot{x}$ and $x_3 = t$ into the formulation of an *equivalent system*:

$$\begin{aligned} \dot{x_1} &= x_2 \\ \dot{x_2} &= \frac{1}{m}(-kx_1 - bx_1 cos(x_3)) \\ \dot{x_3} &= 1 \end{aligned} \tag{3.11}$$

We can also use some of these equations to shed light on some other ideas pertinent to this work, in particular on that of *phase space*. A simple *first-order* nonlinear dynamical system, for instance, might help clarify this concept.

$$\dot{x} = sin(x) \tag{3.12}$$

As I have previously mentioned, a dynamical system produces a *flow*, which in this case, as the system is one-dimensional, acts along a line. This equation says that in this system there is a *vector field* on the $x$ axis, with each vector acting as a force pushing and pulling each state towards another along that axis. Further, where $sin(x)$ is positive that push will result in a growth of $x$, as its derivative would therefore be positive and indicate a positive slope. If $sin(x)$ is negative then the slope would be negative and therefore $x$ would decrease. The arrows in figure 3.2 depict this vector field. At the points where $sin(x) = 0$, for $x = n\pi$, there is no flow: $x$ therefore remains constant. In general, points at which the flow is 0 are called *fixed points*, and the solid dots in the figure are *stable fixed points*, also called *attractors* or *sinks*, as the flow from both directions slopes towards them. The empty dots are *unstable fixed points* also known as *sources* or *repellers*. Regardless of name, however, these points are very important states for the system, as they are points of *equilibrium*.

Figure 3.2: Phase flow and fixed points of the a one dimensional dynamical system $\dot{x} = sin(x)$.

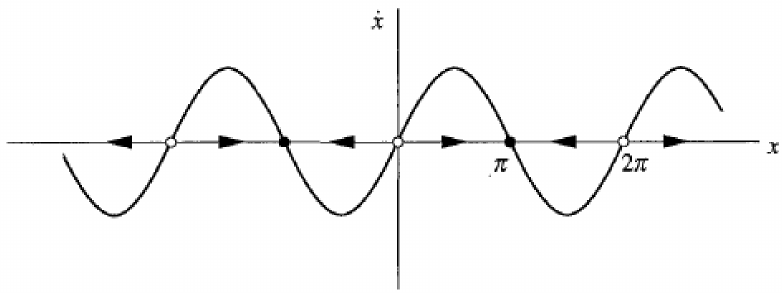


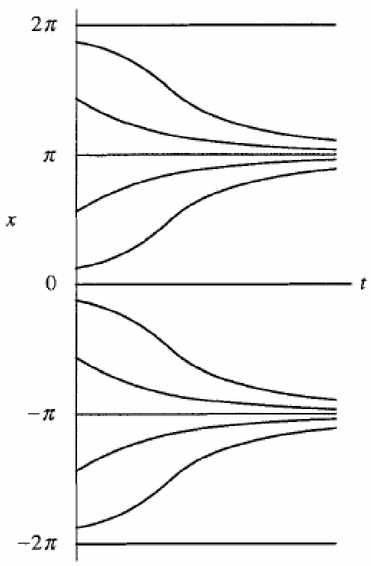


Figure 3.3: Some solutions for different initial conditions to the dynamical system $\dot{x} = sin(x)$.

Consider the trajectory of a point starting from a slightly positive state, say $x = 0.1\pi$. In that region $sin(x)$ is positive, but small, meaning that point willwill grow at a slow rate since the slope of its evolution is low. As it grows, its flow will increase too, exponentially until a maximum of the growth rate will be reached at $x = \frac{\pi}{2}$. Past this point, its growth will remain positive, but at a decreased speed, slowing until it reaches point $x = \pi$, where it will cease its evolution. Figure 3.3 shows the temporal evolution of the system with different starting conditions at $t = 0$.

The above example attempts to clarify a feature of the geometrical approach. Thinking in those terms means to imagine how a point moves, or changes its state, under the influence of a dynamical system's flow. It is to imagine and 'follow' its *trajectory* as time passes. The charm of this approach mostly lies in this physical, almost bodily and haptic depiction of temporal behaviour, even of very complex systems. That is the reason why I'm attracted to way of thinking about dynamical systems.

Another useful example of a dynamical system is the *logistic equation*. This is a simplified model for the growth of a population with a specified growth rate $r$ in an environment with a determined *carrying capacity* $K$:

$$\dot{N} = rN\left(1 - \frac{N}{K}\right) \tag{3.13}$$

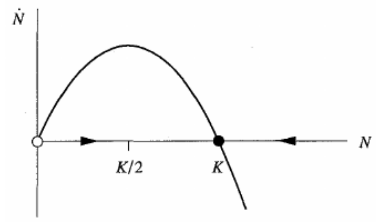


Figure 3.4: Phase flow and fixed points of the *logistic equation*.

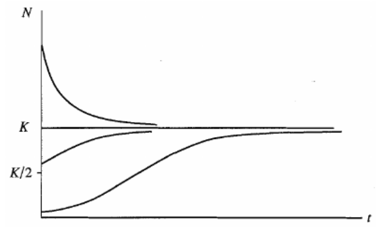


Figure 3.5: Some solutions for different initial conditions to the dynamical system base on the *logistic equation*.

As figure 3.4 shows, the flow of this system has two fixed points, $N = 0$ and $N = K$, the first being unstable and therefore repelling and the second stable and therefore attracting. That means that any state of the system will be pushed away from 0 and pulled towards $K$ and always tend to reach the carrying capacity of the environment. When $N > \frac{K}{2}$, the growth starts to decrease as it approaches point $K$. Figure 3.5 show the temporal evolution of $N$ under various starting conditions.

A general formulation for a second order *linear* dynamical system is:

$$\begin{aligned} \dot{x} &= ax + by \\ \dot{y} &= cx + dy \end{aligned} \tag{3.14}$$

or, in a more compact form using vector notation:

$$\dot{\vec{x}} = A\vec{x} \tag{3.15}$$

where

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \mathit{and}\, \vec{x} = \begin{pmatrix} x \\ y \end{pmatrix} \tag{3.16}$$

The solutions of equation 3.15 can be visualised as trajectories in the plane $(x, y)$, in this case called a *phase plane*. This system is linear, therefore, as noted earlier in this work, we know that a fundamental rule applies - if $\vec{x_1}$ and $\vec{x_2}$ are both solutions then also $\vec{x} = c_1\vec{x_1} + c_2\vec{x_2}$ is a solution for any $c_1, c_2$. Further, due to the linearity of the system there is only one fixed point $\vec{x^*}$. That point is $\vec{x^*} = 0$ since there it is always $\dot{\vec{x^*}} = 0$, meaning that at that position the flow is zero.

In mathematics, differential equations of this linear form are usually solved by setting:

$$\vec{x}(t) = e^{\lambda t}\vec{v} \tag{3.17}$$

with the $\vec{v}$ vector and growth rate $\lambda$ to be determined. Using the previousformulation for $\vec{x}(t)$ to find a solution

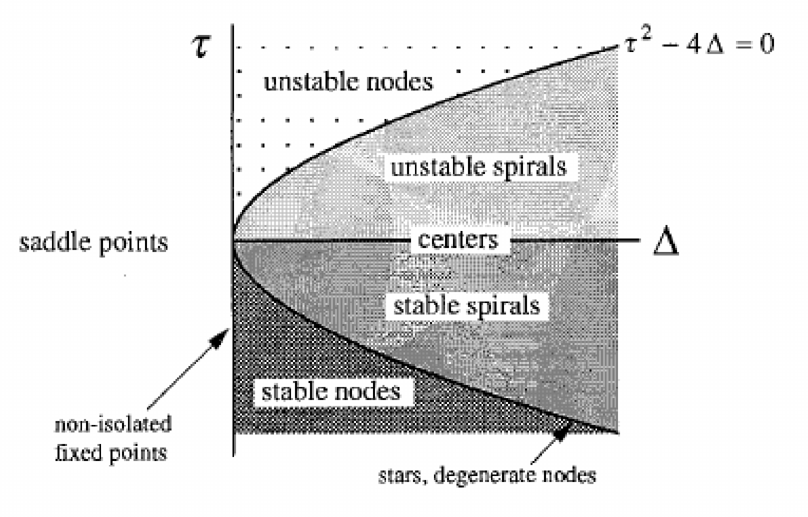


Figure 3.6: Classification of two dimensional fixed point types in dependence on the values of $\tau$ and $\Delta$ (see equation 3.20). Graphic taken from Steven H. Strogatz *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. Westview press, 2014

for equation 3.15, one obtains $\lambda e^{\lambda t}\vec{v} = e^{\lambda t}A\vec{v}$, while simplifying the non-zero factor $e^{\lambda t}$ yields:

$$\lambda\vec{v} = A\vec{v} \tag{3.18}$$

This consequently manifests as a classical *eigenvalue* and *eigenvector* problem. With this formulation we can reduce our problem to a significantly easier search for the (in this case, two) directions $\vec{v_1}$ and $\vec{v_2}$, which remain constant under the influence of the dynamical system. This kind of problem pertains to the field of *linear algebra* and is usually solved with the *characteristic equation* $det(A - \lambda I) = 0$, where $I$ is the identity matrix and *det* stands for the *determinant* function. This procedure is valid for general $n$ dimensional linear dynamic systems, even if in this text it is made explicit only for the case of a two-dimensional system.

The characteristic equation is thus:

$$det\begin{pmatrix} a-\lambda & b \\ c & d-\lambda \end{pmatrix} \tag{3.19}$$

The determinant gives:

$$\lambda^2 - \tau\lambda + \Delta = 0 \tag{3.20}$$

where:

$$\begin{aligned} \tau &= a + d \\ \Delta &= ad - bc \end{aligned} \tag{3.21}$$

The above quadratic equation gives two solutions for $\lambda$:

$$\begin{aligned} \lambda_1 &= \frac{\tau + \sqrt{\tau^2 - 4\Delta}}{2} \\ \lambda_2 &= \frac{\tau - \sqrt{\tau^2 - 4\Delta}}{2} \end{aligned}$$

Each of these two *eigenvalues* correspond to two *eigenvectors* - $\vec{v}_1$ and $\vec{v}_2$ - which can be found by substituting the two solutions back into the equation 3.18. One could think of these two vectors as the two main and independent axes along which the flow of a dynamical system unfolds. In particular, given the linear characteristic, any initial condition (starting state) $\vec{x}_0$ can be written as a linear combination of these two eigenvectors, $\vec{x}_0 = c_1\vec{v}_1 + c_2\vec{v}_2$. This allows us to write the general solution as:

$$\vec{x}(t) = c_1 e^{\lambda_1 t}\vec{v}_1 + c_2 e^{\lambda_2 t}\vec{v}_2$$

Dependent on the values of $A$ and, consequently, on $\tau$ and $\Delta$, the eigenvalues and eigenvectors bring forth very different phase flows that produce qualitatively different temporal behaviour. As can be seen from figure 3.6, these can be grouped into six classes of fixed points $\mathbf{x} = 0$.

- Stable Nodes: In the case of stable nodes, both eigenvalues are real and negative, e.g. $\lambda_{1,2} < 0$. This means that the time evolution for all points in the phase plane will be governed by an exponential decay towards the fixed point, since $e^{\lambda_{1,2} t}$ tends towards 0 as time advances.

  Figure 3.7 depicts the flow generated by the system:

$$A = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix} \tag{3.22}$$

  This is also known as a *symmetrical node* or *star*. igure 3.8, meanwhile, shows the flow of the system

$$A = \begin{pmatrix} -3 & 0 \\ 0 & -1 \end{pmatrix} \tag{3.23}$$

  In this case a stronger 'drag' is acting along the first dimension, and the two eigenvectors correspond with the two plane axes. For the system above the first axis, with the stronger pulling direction, is also called the *fast eigendirection* and the second *slow eigendirection*.

  This kind of fixed point is also an *attracting* node, and in this case in particular it is a *globally attracting* node for all points on the plane. It is also *asymptotically stable*, which means that all trajectories that start near to it will remain so at all times.

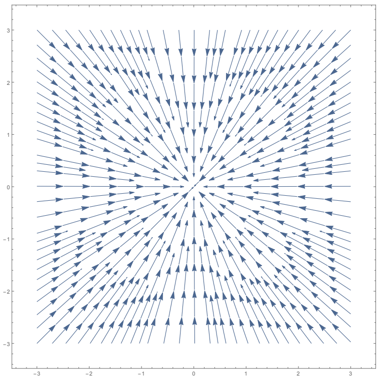

Figure 3.7: Phase flow of a *symmetrical node* or star fixed point.

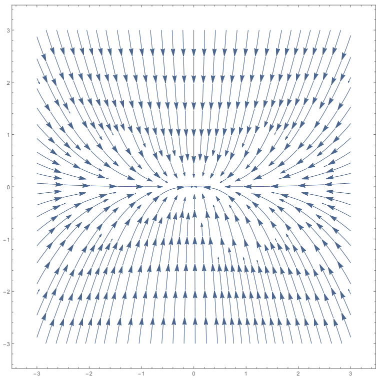

Figure 3.8: Phase flow corresponding to an *asymmetrical node* kind fixed point.

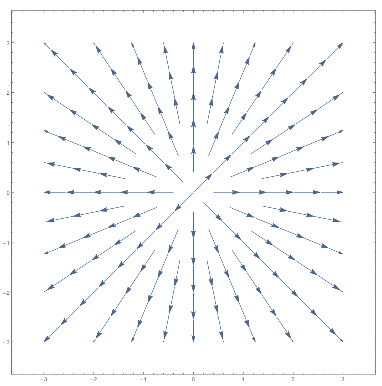

Figure 3.9: Phase flow of a *symmetrical unstable node* fixed point.

- UNSTABLE NODES: In the case of unstable nodes, both eigenvalues are real and positive $\lambda > 0$. Temporal behaviour corresponds here to an exponential growth starting at any point on the plane and, as time advances, leading away from the fixed point towards infinity.

  In figure 3.9, we see the system:

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \tag{3.24}$$

  is depicted. In figure 3.10 the system

$$A = \begin{pmatrix} 3 & 0 \\ 0 & 1 \end{pmatrix} \tag{3.25}$$

  Both are *unstable nodes*.

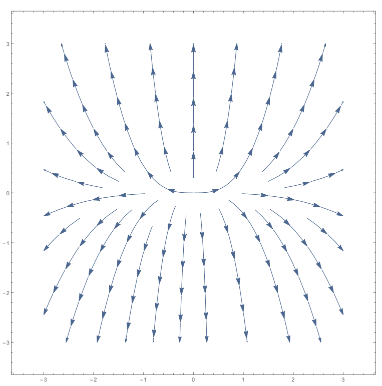

Figure 3.10: Phase flow corresponding to an *asymmetrical unstable node* fixed node.

- SADDLES: With saddles, the eigenvalues are both real, but one is positive and the other negative. Most trajectories grow to infinity away from $\vec{x}^*$. The growth occurs asymptotically along one direction, the eigendirection corresponding to the positive eigenvalue. Only when a trajectory starts exactly on the eigendirection corresponding to the negative eigenvalue will it move towards the fixed point.

  The system

$$A = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \tag{3.26}$$

  is depicted in figure 3.11. Here, the positive growth direction is the along the $x$ axis and the negative growth direction along the $y$ axis.

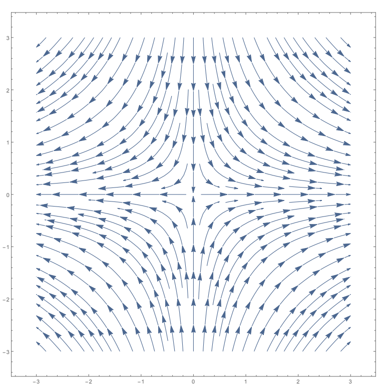

Figure 3.11: *Saddle* fixed point: symmetric flow with stable manifold along the $y$ axis.

$$A = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \tag{3.27}$$

  This system's (3.27) phase flow is depicted in figure 3.12. This is also a saddle fixed point, but in this case the eigendirections are rotated by 45 degrees. The negative growth axis is along the $(-1, 1)$ direction - also called the *stable manifold*. This is the set of initial conditions for which $\mathbf{x}(t) \to \mathbf{x}^*$. The direction of positive growth is $(1, 1)$, also called the *unstable manifold* of $\mathbf{x}^*$. A typical trajectory in the phase plane approaches the unstable manifold as $t \to \infty$

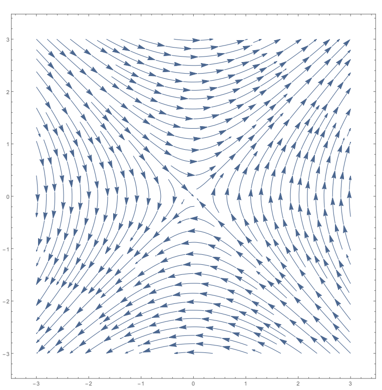

Figure 3.12: *Saddle* fixed point: symmetric flow with stable manifold along the $(1, 1)$ direction.

- Centres: The condition $\tau = 0$ results in complex eigenvalues $\lambda_{1,2} = \pm i\,\omega$ and eigenvectors. As the evolution of the system is governed by $e^{\pm i\,\omega\,t}$, the solutions are therefore oscillating, i.e. rotating on closed paths around the centre of the axes.

  Figure 3.13 depicts the vector field generated in the phase plane by the system

$$A = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \tag{3.28}$$

  Figure 3.14 shows the vector field of the system

$$A = \begin{pmatrix} 0 & 3 \\ -1 & 0 \end{pmatrix} \tag{3.29}$$

  As we can see, this vector field is slightly stretched stretched along the $x$ axis.

  Centres are regarded as *neutrally stable*, as trajectories near to the fixed point are neither attracted to nor repelled from it.

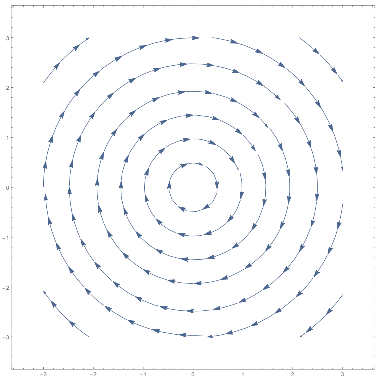

Figure 3.13: Fixed point of *centre* type: symmetric flow.

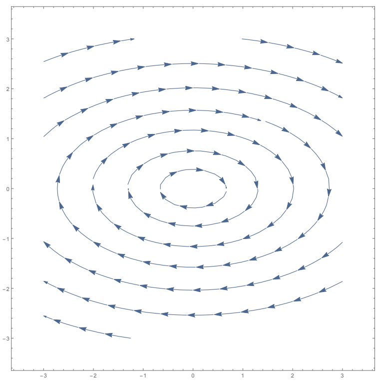

Figure 3.14: Fixed point of *centre* type: asymmetric flow.

- Unstable Spirals: Here the system's eigenvalues can have both a real and a complex part $\lambda = \alpha \pm i\,\omega$. If $\alpha > 0$, the temporal evolution $e^{(\alpha \pm i\,\omega)t}$ is a combination of the oscillatory behaviour of the centre type fixed point and of that of the unstable node: these are oscillations growing away from the fixed point. Figure 3.15 depicts the flow of this system:

$$A = \begin{pmatrix} 0.5 & 1 \\ -1 & 0.5 \end{pmatrix} \tag{3.30}$$

- Stable Spirals: If the complex matrix eigenvalues have a negative real part $\alpha < 0$, the temporal evolution of a system will exhibit *decaying oscillations*, slowly decreasing towards the fixed point. This is the fixed point type corresponding to a damped oscillator, noted earlier in this dissertation. Figure 3.16 shows the flow for the system with matrix:

$$A = \begin{pmatrix} -0.5 & 1 \\ -1 & -0.5 \end{pmatrix} \tag{3.31}$$

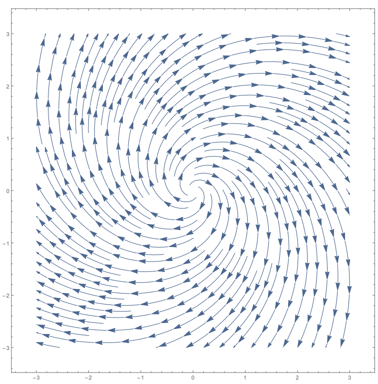

Figure 3.15: Phase flow corresponding to an *unstable spiral* fixed point.

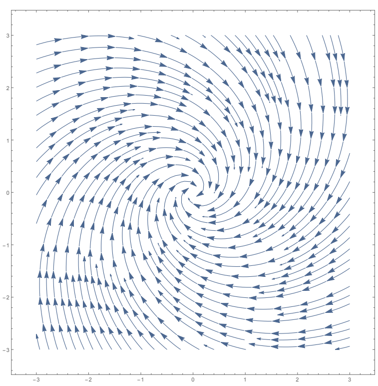

Figure 3.16: Phase flow corresponding to a *stable spiral* fixed point.

These six types of fixed point are fundamental in the study of all dynamical systems, particularly when it comes to the qualitative analysis of nonlinear systems. In fact, by virtue of the *linearisation* technique, the phase flow of a nonlinear problem can be well approximated with a corresponding linear flow near to a fixed point.

This approximation can be computed by performing a power expansion (as in a Taylor series expansion) of the flow function near a fixed point $\vec{x}^*$. That means that a general system

$$\dot{x} = f(x,y) \tag{3.32}$$

$$\dot{y} = g(x,y) \tag{3.33}$$

could be approximated near to the fixed point $(x^*, y^*)$ by:

$$\begin{aligned} \dot{x} &= f(x^*,y^*) + (x-x^*)\frac{\partial f}{\partial x} + (y-y^*)\frac{\partial f}{\partial y} + O(x^2,y^2,xy) \\ \dot{y} &= g(x^*,y^*) + (x-x^*)\frac{\partial g}{\partial x} + (y-y^*)\frac{\partial g}{\partial y} + O(x^2,y^2,xy) \end{aligned}$$

Here the partial derivatives are computed at the fixed point and $O(x^2, y^2, xy)$ is shorthand for second order terms in $x$ and $y$, which are very small and therefore negligible. Hence, in matrix form, the flow near the fixed point can be formulated in a linearised form:

$$\begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} = \begin{pmatrix} \frac{\partial f}{\partial x} & \frac{\partial f}{\partial y} \\ \frac{\partial g}{\partial x} & \frac{\partial g}{\partial y} \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \tag{3.34}$$

with the matrix

$$A = \begin{pmatrix} \frac{\partial f}{\partial x} & \frac{\partial f}{\partial y} \\ \frac{\partial g}{\partial x} & \frac{\partial g}{\partial y} \end{pmatrix} \tag{3.35}$$

This is called the *Jacobian matrix*, and is evaluated at the fixed point. The Jacobian matrix is the multivariable analogue of the one-dimensional derivative.

As an example, the following nonlinear system:

$$\begin{aligned} \dot{x} &= -x + x^3 \\ \dot{y} &= -2y \end{aligned}$$

has three fixed points at $y = 0$ and $x = 0$ or $x = \pm 1$. The Jacobian of the flow is:

$$A = \begin{pmatrix} -1+3x^2 & 0 \\ 0 & -2 \end{pmatrix}$$

which at fixed point $(0,0)$ becomes

$$A = \begin{pmatrix} -1 & 0 \\ 0 & -2 \end{pmatrix}$$

This is the fixed point of a stable node, with a fast eigendirection along $x$. At the fixed point $(\pm 1, 0)$ the Jacobian is:

$$A = \begin{pmatrix} 2 & 0 \\ 0 & -2 \end{pmatrix}$$

which is the fixed point of a saddle with stable manifold on the $y$ axis. That is, with this technique we can qualitatively look at a nonlinear system as a sort of 'combination' of multiple linear systems around its fixed points.

However, not all of the basic fixed points have the same *structural stability*, and an arbitrary small perturbation may change the stability of some fixed points. For instance, an arbitrary small amount of damping would transform a centre into a spiral, with completely different stability behaviour. Instead, fixed points with $Re(\lambda) \neq 0$ for both eigenvalues are called *hyperbolic* and are much more resistant to small perturbations as the centre type. Consequently, nonlinear systems with a Jacobian of a fixed point, in which all eigenvalues are with non-zero real part, will have a local phase flow near that point *topologically equivalent* to the phase portrait of linearisation. For two flows to be topologically equivalent there must exist a homeomorphism (continuous function with continuous inverse) that maps the one into the other.

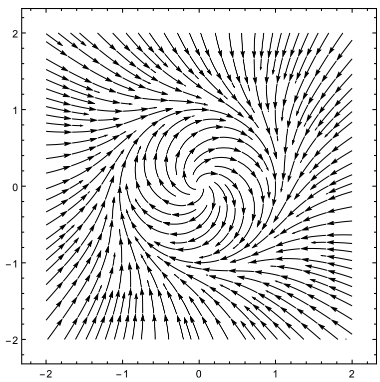

Figure 3.17: Phase flow of the *stable limit cycle* attractor.

A special and important variety of fixed points appearing in two dimensional systems is the *limit cycle*, which is a typical nonlinear phenomenon. A limit cycle is an *isolated* closed trajectory in phase space. While in the *centre* fixed point type all trajectories are closed on themselves, in this case there is only one closed curve cure in the plane. All other trajectories on the plane either spiral away or towards this closed curve. If all neighbouring trajectories spiral towards it, that would be a *stable limit cycle*, while if trajectories would spiral away from it we would observe an *unstable limit cycle*. The system:

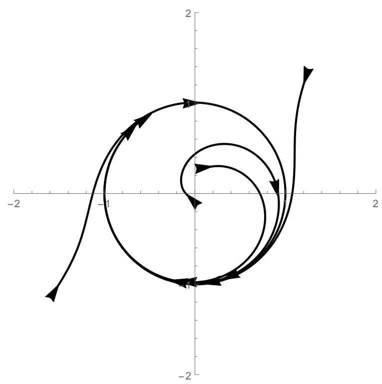

Figure 3.18: Some *phase trajectories* produced by *stable limit cycle* attractor.

$$\begin{aligned}\dot{x} &= y + x(1 - x^2 - y^2)\\ \dot{y} &= x + y(1 - x^2 - y^2)\end{aligned}$$

will produce the phase flow and phase trajectories as depicted in figures 3.17 and 3.18; the fixed point $\vec{x}^* = (0,0)$ is therefore a stable limit cycle. While the system:

$$\begin{aligned}\dot{x} &= y - x(1 - x^2 - y^2)\\ \dot{y} &= x - y(1 - x^2 - y^2)\end{aligned}$$

would instead result in a phase flow which pushes all trajectories (except the one lying exactly on the limit cycle's manifold) away from it, on one side spiralling down towards the centre, on the other towards infinity (see figure 3.19).

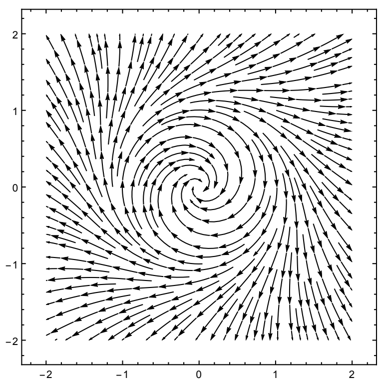

Figure 3.19: Phase flow of the *unstable limit cycle* attractor

Another important aspect to note is that in the limit cycle case the amplitude (and frequency, of course) of the oscillatory behaviour is determined by the system itself, from within it. In the case of centre fixed points, the amplitude of the oscillations is determined by the initial conditions and a slight perturbation will therefore have an effect which lasts forever. The limit cycle is therefore sturdier with respect to internal influences and is thus structurally more stable. In the case of the stable limit cycle, after a short time of adaptation to the perturbation, the system will return to its 'favourite' oscillation frequency and amplitude.

Limit cycles are of particular scientific interest as they model systems which exhibit self-sustained and activated oscillations even in the absence of external driving forces. This type of system is very common in nature and biology, e.g., the beating of a heart, the sleep/wake cycles, etc.

For systems of three dimensions, all the types of fixed points and limit cycles we have found in the two dimensional cases could appear. In many cases, the temporal behaviour of a dynamical system with three-dimensionality can be qualitatively analysed in terms of combinations of the fixed point types we have already encountered in the two dimensional cases. For instance, there can be flows in phase space which are the result of a centre fixed point on one plane and of a stable one-dimensional node on the remaining third dimension. Or of a saddle node on one plane and an unstable one-dimensional on the remaining axe. The technique of linearisation can also be applied, therefore complex and nonlinear systems may also be analysed in terms of the behaviour they exhibit near their fixed points.

But in three dimensional systems a very special behaviour type can appear, which is not possible in lower dimensions. It's the *chaotic behaviour*. Glimpses of the possibilities of a chaotic behaviour could already be found in Poincaré's work, but it is only in the 1970s that the groundbreaking work of Lorenz[5] began to be acknowledged as a central (if not *the* central) topic in dynamical systems theory. Studying meteorological phenomena and trying to model them mathematically, Lorenz discovered a dynamical system that was *inherently unpredictable*. This does not mean that the system evolves casually, without rules, or randomly. It means that its dependence on the initial conditions (i.e., the starting point in the phase space for a trajectory) is extremely strong. More precisely: consider two points in the three dimensional phase space of this system which are arbitrarily near to each other. Let them be the starting point of a trajectory in phase space, i.e. evolving according to the system's formulation. The two trajectories starting from the two points after some time will be very far from each other, eventually following completely different

[5] Edward N Lorenz. Deterministic nonperiodic flow. *Journal of the atmospheric sciences*, 20(2): 130-141, 1963

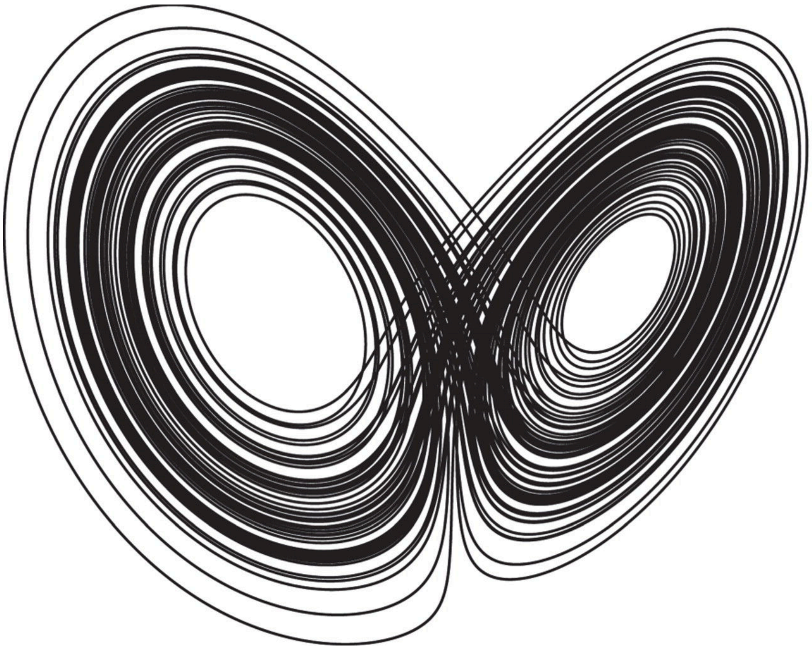

Figure 3.20: A phase space trajectory produced by the three-dimensional Lorenz system

paths. This is a qualitatively different kind of behaviour than all we have seen until now. In two dimensions, two points near to each other will follow trajectories that will *remain* near to each other. The consequence of chaotic behaviour is that, as it is not possible to determine the initial condition of any system with infinite precision, and only rough approximations are possible, these kind of systems are inherently *unpredictable*. In terms of differential equations, Lorenz's dynamical system is formulated as:

$$\begin{aligned} \dot{x} &= \sigma(y-x) \\ \dot{y} &= x(\rho - z) - y \\ \dot{z} &= xy - \beta z \end{aligned}$$

where $\sigma, \beta, \rho$ are fixed parameters.

But Lorenz discovered more. He has observed that, despite the inherent impossibility of knowing the exact state of the system at any time, there was some kind of *structure* in how different trajectories in phase space would evolve (see figure 3.20). The trajectories seemed to revolve and oscillate around a *strange attractor*, which he called an 'infinite complex of surfaces'. Today, we would say *fractal* for describing this special kind of spatial structures.

A full treatment of this topic is not possible in such sort introduction. Still, the important aspect to underline is that these kinds of systems approach the behaviour of everyday life - the structures of growing plants, the weather, etc. - phenomena which generate structures that are self-similar but at the same time always changing.

Even more essential is that the dynamical systems theory provides relatively compact descriptions and a powerful visual language through which *formulate* and *analyse* those phenomena.

As with most nonlinear dynamical systems, chaotic systems are not solvable. As I said before, this means it is not possible to reach a mathematical formulation, an equation which would directly describe the system state's evolution in time. That is, the behaviour of the system remains implicit in the differential equations that describe it and cannot be *unpacked* into a function $\mathbf{x}(t)$. But how, then, is it even possible to study those systems? How can a phase trajectory for the Lorenz system as in figure 3.20 be drawn? One possible method for the analysis of chaotic systems is *numerical simulation*. The idea is that, as these systems are not solvable but do provide definite rules of evolution in time, in order to observe how this evolution behaves (i.e. which kind of trajectories it produces), one must follow the evolution of one point of the phase space under the system's flow. Meaning to actually *sit* on this chosen point, compute the flow at that point and make a small step into that direction, then recompute the flow at that point and make the next step. Re-iterating this process means to *simulate* the system, to perform a *numerical integration* of it.[6] This results in a trace showing the system's temporal behaviour in dependence of its initial condition. The implementation and execution of simulations on computers has grown into one of the most-used methods in the study of dynamical systems, giving birth to a novel method for research in physics and mathematics, the so-called 'third way' between the empirical and the theoretical praxis.

[6] Refer to the section *rattle* integration algorithms in the Appendix for an introduction to the problem of numerical integration.

## *3.2 Dynamical Systems and Cognitive Science*

The enactive theory's understanding of cognition based on a sensory-motor coupling of the living agent with its environment, in an ongoing interaction, resonates well with a dynamical systems view. Enactive theory, in fact, directly refers to dynamical systems in describing its perspective on cognition.[7]

[7] Evan Thompson. *Mind in life: Biology, phenomenology, and the sciences of mind*. Harvard University Press, 2010; and Evan Thompson and Francisco J. Varela. Radical embodiment: neural dynamics and consciousness. *Trends in Cognitive Sciences*, 5(10):418-425, October 2001

In particular, enactive theories rely on a specific approach followed in cognitive science that describes cognitive systems as dynamical systems. According to this perspective, all aspects of action, perception and cognition should be tackled from a dynamical perspective. Aspects of the internal structure of agents, the coupling between sensory and motor systems, the interaction between perception and higher cognitive processes and, naturally, the mutual influence of agent and environment should all be formulated in terms of dynamical systems. That is, systems in which these elements are bound in a continuous interaction that is *defining*: they cannot be considered in isolation and together they form one system of intertwined parts.

This approach contrasts with (current) proponents of so-called cognitivist or computationalist hypotheses, which define artificial or natural cognitive agents in terms of computational *machines*, that is in terms of symbol-processing. This opposition is best explained in the words of philosopher Timothy van Gelder, one of the strongest proponents of this theory:[8]

> *The cognitive system is not a computer, it is a dynamical system. It is not the brain, inner and encapsulated; rather, it is the whole system comprised of nervous system, body, and environment. The cognitive system is not a discrete sequential manipulator of static representational structures; rather, it is a structure of mutually and simultaneously influencing change. Its processes do not take place in the arbitrary, discrete time of computer steps; rather, they unfold in the real time of ongoing change in the environment, the body, and the nervous system. The cognitive system does not interact with other aspects of the world by passing messages or commands; rather, it continuously co-evolves with them.*

An essential and characterising aspect of the dynamical approach is how it sees cognition fundamentally as a temporal phenomenon, as being *in* time. Time is at the heart of the dynamical perspective: its focus is on how the system evolves in time, on the temporal unfolding of its behaviour, rather than on some particular *state*. The beginning and end states of the cognitive process are secondary, or maybe not interesting at all: cognition is *ongoing*.[9]

Picking up on this perspective, Randall Beer attempts to formulate a simple model for agent-environment interaction in terms of a dynamical system.[10] He poses that both the agent $A$ and environment $E$ are dynamical systems: their states would then evolve by $\dot{\mathrm{x}}_A = A(\mathrm{x}_A;\mathrm{u}_A)$ and $\dot{\mathrm{x}}_E = A(\mathrm{x}_E;\mathrm{u}_E)$. $\mathrm{x}_{A,E}$ represents their internal state and $\mathrm{u}_{A,E}$ stands for parameters that are time-independent and internal to the systems. A further assumption is that both systems posses convergent dynamics: i.e., they *tend* to maintain their states in a bounded region and not diverge into infinity. In this instance, agent and environment system are in constant interaction, which means that some of their parameters are dependent on each other's state through a coupling function. This function will be $S$ for the sensory function

[8] Robert F Port and Timothy Van Gelder. *Mind as motion: Explorations in the dynamics of cognition*. MIT press, 1995

[9] Evan Thompson. *Mind in life: Biology, phenomenology, and the sciences of mind*. Harvard University Press, 2010; Tim Van Gelder. The dynamical hypothesis in cognitive science. *Behavioral and brain sciences*, 21(5):615-628, 1998; and Esther Thelen. Time-scale dynamics and the development of an embodied cognition. In Robert F. Port and Timothy van Gelder, editors, *Mind As Motion - Explorations in the Dynamics of Cognition*, chapter 3, pages 69-100. MIT Press, 1995

[10] Randall D Beer. A dynamical systems perspective on agent-environment interaction. *Artificial intelligence*, 72(1-2): 173-215, 1995

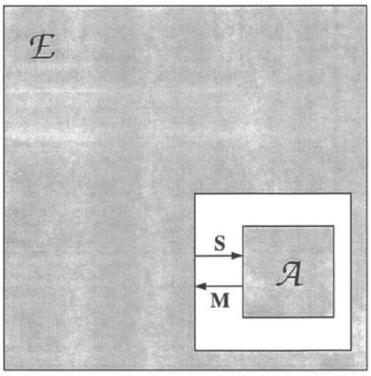


Figure 3.21: An agent and its environment as coupled dynamical systems. From: Randall D Beer. A dynamical systems perspective on agent-environment interaction. *Artificial intelligence*, 72(1-2): 173-215, 1995

coupling environment state and agent's parameters and $M$ for the motor function connecting the agent's state with the environment's. Thus the coupled system could be rewritten as (see figure 3.21):

$$
\begin{aligned}
\dot{\mathrm{x}}_A &= A(\mathrm{x}_A; S(\mathrm{x}_E); \mathrm{u}_A^{'}) \\
\dot{\mathrm{x}}_E &= E(\mathrm{x}_E; S(\mathrm{x}_A); \mathrm{u}_E^{'})
\end{aligned}
$$

Here the $\mathrm{u}_{A,E}^{'}$ stands for all parameters that are excluded from the coupling. This emphasises the role of *feedback* in the system. Every action $M$ of the agent modifies the environment's state, which in turn affects the agent through the sensory connection $S$. Both systems are thus continuously affecting each other's phase flow. Since not all parameters are under the influence of the other system, each element in this situation cannot specify the future trajectory of the other; rather it acts like a *perturbation* on the other's dynamics and trajectory. Beer underlines how agent and environment have to be considered as a whole system whose properties do not reside in either of the two interacting components. Further, the agent's behaviour is not located just in itself or the environment alone, but in the coupled system; the agent's behaviour is determined by its internal dynamics and its interactions with the environment, it *emerges* in the interaction process.

*Emergence* is a term that appears throughout the literature related to the dynamical approach. The term indicates a coherent and perceptible process arising from the interactions between the parts of the system. It is a higher-order organisation of the whole system into a recognisable structure which results from low-level interaction. That is, it is a behaviour of the system which is not pre-specified or formulated in the rules of interaction, but surfaces as a consequence of them. Most importantly, it is an *unpredictable* phenomenon which materialises spontaneously and surprisingly. Enaction theory in particular stresses that *emergence* has a two-way quality: it does not only indicate that a whole arises for the organisation of the parts, but also that the parts arise from the whole. The particular behaviour of each part of the whole is determined by the whole as much as the whole is determined by the interacting parts. This *dynamic co-emergence*, returning to Randall's system, means exactly that the properties of the joint system environment-agent are emergent from their interaction but at the same time co-determine their behaviour.

## 3.3 *Dynamical Systems in Electronic and Computer Music*

Dynamical systems enter the praxis of electronic and computer music in various ways, but mostly in the form of models implemented and simulated on a digital computer. The evolving state of a dynamical system, as it results from the numerical integration of its rules, is the material used musically. Various composers have been interested with the complex temporal paths generated by chaotic phenomena, as well as with the possibility of producing a wide range of diverse behaviour through the modification of only a few parameters in a model.[11]

[11] Strictly speaking physical modelling synthesis techniques are also part of this category as physical models are a subset of dynamical systems, but I will not refer to these methods here as I am more interested in approaches that receive the peculiarities of an approach based on dynamical systems

A first approach uses dynamical systems mostly in the form of iterated maps: a different formalism of dynamical systems, in which time is taken to be discrete rather than continuous, as in the theoretical introduction at the beginning of this chapter. In particular, these systems have been used in order to produce traditionally notated scores and can therefore be categorised under the praxis of algorithmic composition.[12] In these cases, dynamical systems, mostly chaotic, have been in a way *instrumentalised*, used as tools to generate temporal structures with both a high level of complexity and a coherence deriving from the simple rules they employ. The composers' fascination for dynamical systems leads back to the belief that a whole world of possible temporal processes lie at their disposal, through the slight turning of some system parameters.

[12] Michael Gogins. Iterated functions systems music. *Computer Music Journal*, 15(1):40-48, 1991; and Jeff Pressing. Nonlinear maps as generators of musical design. *Computer Music Journal*, 12(2):35-46, 1988

Chaotic dynamical systems are also used in sound synthesis: in this case, their evolution is more or less directly *audified* and translated into sound. Those systems can be realised through analogue circuitry or digital computation.[13] The encounter with the particular system's behaviour is here unmediated by the step of transposition into a traditional musical notation. There are some required steps in this process to transform the state of the system into sound, but the directness of the situation of listening in *real-time*, that is while the system is actually evolving, to the sound it produces, allows for an essentially different experience than in the previous algorithmic approach. The most salient qualities of this process being the wide range of timbral qualities and the space of diverse temporal and sonic behaviour that can be achieved and the low dimensionality of the parameter space.

[13] Dan Slater. Chaotic sound synthesis. *Computer Music Journal*, 22(2):12-19, 1998; and Agostino Di Scipio. Iterated nonlinear functions as a sound-generating engine. *Leonardo*, 34(3):249-254, 2001

Given the unpredictability of these systems' behaviour and their extreme sensitivity to initial conditions, this situation leads to a particularly explorative attitude. In order to build a perceptual *image* of the system's behaviour, a deeper engagement from the side of the musician/composer is needed. The parameters of the running model can be modified by employing interfaces, or eliciting the system's responses with some external perturbing input, so that it is also possible to act on the system during its evolution. A different type of contact could therefore be established extending into a more performative mode.[14]

[14] Tom Mudd, Simon Holland, Paul Mulholland, and Nick Dalton. Dynamical interactions with electronic instruments. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 126-129. Goldsmiths, University of London, 2014

The use of dynamical systems leads to a fundamentally different process than usually found in sound synthesis. Typically, the sonic result is pre-specified and the sound synthesis engine is implemented and adapted towards the achievement of that result. In the case of synthesis by chaotic dynamical systems, the model that produces the sound is specified by the composer formulating its rules of evolution, but the sonic result is unknown *a priori*. An exploration of the behaviour space generated by this process is necessary in order to construct the piece. I see here a connection to the concept of *non-standard synthesis* as depicted by Luc Döbereiner.[15] Sound synthesis through dynamical systems essentially consists of the formulation of rules of evolution or coupling. Sound is not specified by its perceptual appearance, but by the process that constructs it. This process becomes the object of composition. Of a composition that extends into sound synthesis, not in the sense of a composition *with* sound, but as composition *of* sound (in terms of processes). A process whose result is largely unknown at its onset and that has to be actually carried out, performed, explored and probed especially if it includes external disturbances.

[15] Luc Döbereiner. Models of constructed sound: Nonstandard synthesis as an aesthetic perspective. *Computer Music Journal*, 35(3):28-39, 2011

Further, temporal scales in which the system's evolution takes place are joined in a continuum: for example, oscillations can range from the audible domain to periods of seconds or hours depending on its parameters, while the formal system itself remains unaltered. Dynamical systems offer a formulation framework in which all temporal aspects of sound and its organisation could be integrated. Or, from the more deeply dynamical perspective as proposed by Di Scipio, to conceive of all temporal aspects of such a generative composition, microscopic to macroscopic as emergent from low-level nonlinear interactions.[16]

[16] Agostino Di Scipio. Iterated nonlinear functions as a sound-generating engine. *Leonardo*, 34(3):249-254, 2001

The works and thought of Agostino Di Scipio are paradigmatic for *thinking in terms of* dynamical systems in computer music. His work is central for this dissertation and a great inspiration for my artistic work. Di Scipio starts by observing that, in common computer music practice, an *interactive* musical system implies a linear relationship between performer and system: 'agent acts, computer re-acts'[17]. The problem he sees is that 'the sound-generating system is not itself able to directly

[17] Agostino Di Scipio. 'Sound is the interface': from interactive to ecosystemic signal processing. *Organised Sound*, 8(3):269-277, 2003

cause any change or adjustment in the external conditions set to its own process'. That is, it has no part in the determination of its own state and has to completely rely on the dynamics of the performer. His proposition is to conceive interaction as a 'by-product of lower-level interdependencies among system components'. Interaction is what happens when entities (or agents) are bound in an interdependent relationship. This perspective strongly refers to the enactive approach in Varela, Maturana and Thompson's work. In particular, Di Scipio sees the computer music system as a dynamical system that has both the faculty to sense external changes in the environment it is embedded in and to self-observe its own state, thus becoming a self-organising system - a perspective that resembles the definition of agent I have elaborated in the previous chapter (see 2.4 Enaction). In Di Scipio's work, performer, computer system and environment together form a system in which interaction is constituent. The interactions happening in the system would be the result of the interdependencies among the system's components: the planning and staging of these interdependencies is the region in which composition would then actually take place. In Di Scipio's words, referring to Chadabe's *interactive composing*:

> *This is a substantial move from interactive music composing to composing musical interactions, and perhaps more precisely it should be described as a shift from creating wanted sounds via interactive means, towards creating wanted interactions having audible traces. In the latter case, one designs, implements and maintains a network of connected components whose emergent behaviour in sound one calls music.*

Di Scipio understands his pieces and sound installations as components of an *ecosystem*, in which audience, performers, machines, and the room acoustics, have a structural, i.e. a constituent role. An ecological perspective on musical performance considers all elements, which traditionally may be considered as disturbances, sources of error, or unwanted deviation in performance as an essential component of the musical outcome.[18] Di Scipio directly addresses those elements in his pieces and makes them part of a complex network of composed interrelations. The tools he uses in this endeavour are *feedback systems*: computer systems which construct a closed loop between sonic input captured by a microphone and their sonic output projected by the loudspeakers. This is a computational mechanism which has only sound as its interface and is *immersed* in the physical

[18] Jonathan Impett. Interaction, simulation and invention: a model for interactive music. In *Proceedings of ALMMA 2001 Workshop on Artificial Models for Musical Applications*, pages 108-119, Cosenza, Italy, 2001; and Simon Waters. Performance ecosystems: Ecological approaches to musical interaction. *EMS: Electroacoustic Music Studies Network*, pages 1-20, 2007

world, 'constantly affecting the sonic ambience in that environment and constantly being affected by it'[19]. The aim is to realise a complete interdependency between computer music system and the ecosystem it is within:

> *There is no way to isolate the system input from its own output, as all output is an input. The very idea of an input/output system should be abandoned, in this context. The room space becomes the medium through which the process hears itself and acts upon itself.*

These pieces, and the computer systems they employ, do not exist without performance. They need the contingencies of a real performance in order to function: 'there is no form without performance'.

Emergence is a concept that is central in Di Scipio's work. Following the theory of enaction, he describes the phenomenon as upward and downwards-causation.[20] But apart from these definitions, what I believe is meant here by emergence is the appearance of the *unexpected* and *unpredictable*. That is, moments in which the mutual interaction of the entities in a sonic ecosystem appear to *self-organise* and bring forth a global behaviour otherwise impossible to produce, much as it is impossible to solve the nonlinear dynamical system they produce. In this sense, composing in terms of dynamical systems means to compose the interdependencies in an ecosystem for emergent phenomena to occur. It means to create the conditions for being surprised, and for experiencing the unexpected.

However, longing for the unexpected also exposes *failure*. Creating such a tight feedback coupling between computer system and its environment (or performer) aims at producing a completely circular situation. None of the components is in control of the situation, and nor can they be without disrupting interaction. Consequently, there is no way to predetermine whether or not errors will happen, and thus failure and its coping strategies should be made part of the composition. Di Scipio addresses this issue in his scores when he describes the *emergency measures* to take in case of failure in order to push the system towards more stable regions of behaviour.[21]

I believe that Di Scipio's work truly captures some of the most fundamental qualities of interaction I want to address. In particular that, in his understanding, interaction cannot be relegated to an ancillary function. Rather, it needs to be at the centre of compositional practice, becoming a form-giving or even a generative principle.

[19] Agostino Di Scipio. Listening to yourself through the otherself: on background noise study and other works. *Organised Sound*, 16(2):97-108, 2011

[20] Evan Thompson. *Mind in life: Biology, phenomenology, and the sciences of mind*. Harvard University Press, 2010

[21] Julia H Schröder. Emergence and emergency: Theoretical and practical considerations in agostino di scipio's works. *Contemporary Music Review*, 33(1): 31-45, 2014

In his words:[22]

*Understanding interaction as the object of composition means that the internal ecology of the musical process is captured in the mutual, causal interconnection of many component elements: changes in the ambience response [...] determine unpredictable but consistent reactions and adaptations in the machine's behavior [...], which in turn causes unpredictable but consistent reactions and adaptations in the ambience and the visitors' behaviour.*

[22] Agostino Di Scipio. Iterated nonlinear functions as a sound-generating engine. *Leonardo*, 34(3):249-254, 2001

I believe that dynamical systems provide an apt framework not only as metaphor in describing interaction in this sense, but also in realising actual computational artefacts (I avoid calling them computer music instruments) materialising such interactive systems.

## *3.4 A sense for change: behaviour*

For Alva Noë and O'Regan, 'to perceive is to exercise one's skilful mastery of the ways sensory stimulation varies as a result of bodily movement'[23]. Our perceptions, then, are based on differences detected by our sensory-motor system. That is, when our senses detect something, a difference is formed with respect to the motor system. Also, as Noë explains, our visual perception is strongly dependent on the continuous movements of our heads and eyes, which result in different visual images detected by our retina: the differences between those images are then reintegrated into the image we perceive. Our sensory-motor system even *produces* differences in order to be able to perceive. This is particularly evident in the case of visual perception, but there is evidence that similar mechanisms can be found also for other senses.[24]

[23] J Kevin O'Regan and Alva Noë. A sensorimotor account of vision and visual consciousness. *Behavioral and brain sciences*, 24 (5):939-973, 2001

I draw a connection here to the properties of the sensual receptors as described by neuroscientist Alain Berthoz:[25]

[24] Alva Noë. *Action in Perception*. The MIT Press, 2004

[25] Alain Berthoz. *The brain's sense of movement*. Harvard University Press, 2000

*Sensory receptor functions have a predictive quality. Receptors can detect the derivatives (that is, velocity, acceleration, changes in force and pressure) of the physical variable that stimulate them. Detecting changes in a variable allows the receptors to predict the value of that variable at a future time.*

Our whole perceptual system is very well trained to capture differences, especially time differences (derivatives). When we perceive something, we tend to 'solve' these embodied differential equations - not in a mathematical formulation, but in the sense of a predictive tension towards an anticipated temporal state.

I believe that Edmund Husserl, in his phenomenology of time, would call this tension *protention*. As Merleau-Ponty explains:[26,27]

> *Husserl uses the terms protentions and retentions for the intentionalities which anchor me to an environment. They do not run from a central I, but from my perceptual field itself, so to speak, which draws along in its wake its own horizon of retentions, and bites into the future with its protentions. I do not pass through a series of instances of now, the images of which I preserve and which, placed end to end, make a line. [...]. Time is not a line, but a network of intentionalities.*

That is, perception of time itself is based on a continuously performed *movement* that starts somewhere in the set of retentions, the just passed past moments, but still retained in our consciousness as a sort of lingering echo, passes through the present and protends into an anticipation of the future. Perceiving time means to continuously re-construct the coherency that unites past moments and projecting it before us, into the future.

In my words, this means that we are sensible to dynamical systems. In particular, we are *trained* to see and interpret the world in terms of dynamical systems, as a system or a collection of systems that produce temporal variations on the base of some rule, that of a differential equation. Perceiving or feeling this equation gives us the possibility to look in the future. This may seem obvious: the physical world we are immersed in is an aggregate of dynamical (physical) systems, and not being able to understand these would lead to evolutionary failure. Nonetheless, the interesting consequence here is that using dynamical systems for composing sonic processes would then resonate with the very fundamental mechanisms of perception in eventually making those more perceivable or *graspable*. This is one of the assumptions lying behind this work.

[26] Maurice Merleau-Ponty. *Phenomenology of Perception*. Routledge, 2002

[27] In the English version the french term *protention* is erroneously translated with *protection*. Here, I will use the term *protention* instead.

# 4

# *Case Studies*

In this chapter, a third path is pursued, this time through a series of concrete works. Namely, the implementation of a software framework for physical and dynamical systems modelling (with some experiments realised with it), the research project *Embodied Generative Music* (with some of its artistic case studies), and a proposal for a thought framework for interactive environments tested in a small experimental setting. I will draw a thread through some of the important works and research projects that influenced the development of this dissertation. For the sake of clarity, the narration I have constructed attempts to make this path appear smooth in terms of how each idea or realisation followed another. Of course, this is a construction and does not always coincide with the reality: temporal extensions might strongly overlap.

Still, this part is of particular importance. It aims to clarify how, in this dissertation, artistic practice and experimentation played a central methodological role. In the process of transferring ideas or utopias from theoretical abstraction to a material artefact (software, but also a composition or a sound installation), many decisions were be taken and new thoughts appeared that reflected on the starting conditions. Further, practical experimentation, observation, and experience of the aesthetic qualities of one's own ideas allowed for new formulations and a different kind of understanding. I would say that this self-exposure to my own thoughts in the form of aesthetic artefacts was (and is) one of the main methodological tools I have used in the course of this dissertation. This chapter tries to make this clear.

From another perspective, this chapter collects works of a very different materiality. These are drawn from different contexts, like mathematics, software development, artistic research, computer music research, philosophy, and cognition theory - fields that are far apart from each other. The risk therefore is that these materials might appear as too heterogeneous to construct a coherent discourse. This would probably be true if they would be left dangling without roots, but I hope that the previous chapters provide sufficient ground. Together with the artistic works collected in the appendix to this text, they form the framework of this chapter.

## *4.1 The* rattle *System*

*rattle* is the name I gave to the small software library I've been developing in order to formulate and realise experiments and the studies that follow. For many of those works, *rattle* has been a precondition. Still, it could be thought that the description of a software framework would not fit under in the chapter *Case Studies*; such a description would be better suited for an appendix or a decisively more technical section instrumental to some argument. This is not a mistake, but rather a statement and an acknowledgement. I consider the praxis of programming and of software development as inextricably intertwined with the evolution of thought and of the ideas that are pursued in a research project.

I claim that programming praxis and software development possess a sort of *excess*, as Hans-Jorg Rheinberger, citing Derrida, ascribes to the *means* by which experiments are conducted.[1] He indicates how those means inherently contain more and other possibilities than those to which they are actually held to be bound: they transgress the boundaries within which the research appears to be confined. In my view, this is especially true in the context of a research that is strongly driven by artistic aims, where the relation between realisation, formulation, and the artefact is extremely tight, in an almost inextricable compound. Coding, programming, and the continuous interaction with the formulations of an idea all play an active role beyond the purely instrumental. They sharpen and further ideas and, more importantly, raise new kinds of questions.

[1] Hans-Jörg Reinberger. Experimental systems: Historiality, narration, and deconstruction. *Science in Context*, 7(1):65-81, 1994

In its first incarnation, *rattle* is a physical modelling and simulation software framework. The software tries to offer a programming context in which physical systems – that is, systems resembling or exemplifying physical interactions between simple objects – can be formulated and simulated. It has been implemented in various programming languages: SuperCollider[2], C, Fortran, and a minimal implementation exists in JavaScript. Each reformulation of this same framework in a different programming language has contributed to streamlining the code as well as the underlying thoughts. Eventually the core of the whole framework can be expressed in very few functions and classes.

[2] http://supercollider.github.io/, accessed 23/05/2017

In principle, *rattle* implements a mass-force physical modelling method and, as such, it shares some commonalities with existing software. Two of the most notable examples of similar frameworks are *pmpd*[3], a library of objects integrated in the *pure data* open source visual programming language or the software *GENESIS*[4] and *CORDIS-ANIMA*[5] developed by the ACROE-ICA. However, *rattle* retains some differences in purpose as well as in implementation.

[3] Cyrille Henry. Physical modeling for puredata (pmpd) and real time interaction with an audio synthesis. In *Proc. of the Sound and Music Computing Conference*, October 2004

[4] Nicolas Castagné and Claude Cadoz. Genesis: a friendly musician-oriented environment for mass-interaction physical modeling. In *ICMC 2002-International Computer Music Conference*, pages 330-337. MPublishing, 2002

[5] Claude Cadoz, Annie Luciani, and Jean Loup Florens. Cordis-aniama: Modeling and simulation system for sound and image synthesis - the general formalism. *Computer Music Journal*, 17(1):19 - 29, Spring 1993

- Since the original idea was to develop a tool to explore different behaviours emerging from the evolution of physical systems, *rattle* is not limited to modelling and simulating systems governed by elastic forces or vibrating behaviour. Other physical modelling software used in computer music focuses on exactly this kind of system, a choice that is quite natural in a musical context. Oscillating phenomena play a central role if the aim is to synthesise sounds that imitate acoustic instruments, but may also generate control signals suitable for a high-level control of sound synthesis and formal structures in generative musical processes. For me, the interest in developing this tool was less the imitation of some specific behaviour (for instance, the oscillatory behaviour), and more to explore behaviour and its properties as a perceptual phenomenon. Therefore a wide range of behaviour formulations and implementations were sought, first limited to behaviour as it is specifically appearing in physical systems (systems that model interactions present in the physical world), then in a more broader understanding based on a dynamical systems perspective (refer here to Phase Space Thinking: an experiment).
- The very first implementation was in the *SuperCollider* platform and centred on the simulation of simple models of interacting objects, using their movements (e.g., the variation of speed and position) as control signal for spatialisation (see appendix A.2 cornerghostaxis#1) or for sound synthesis (see the spring scenario described in 4.1.2 An example and some considerations). Only the second reformulation of the framework into a C library allowed to run the simulations at audio rate thus to audify or sonify the movements of the simulated particles to directly synthesise sound (e.g. see appendices A.4 Interstices and A.5 Zwischenräume). Such a possibility distinguishes *rattle* from other frameworks that focus on generating control signals from the evolution of these kind of systems, i.e., low-frequency signals which are mapped to parameters of synthesis processes running in parallel.

- Interaction with simulated physical systems was one of the central points in developing *rattle*. Other software for physically modelling sound synthesis does not offer the possibility to enter the running simulation, even at audio-rate. Nor does other software allow one to control the position of single masses (e.g. spring scenario in 4.1.2 An example and some considerations), thereby changing their state, or change the parameters of the forces acting on each single particle in real-time. In general in such frameworks, after an initial phase in which a more or less complex object or a systems is constructed from simple building blocks (the particles), when the simulation starts the possibilities of interaction are strongly limited. It is not possible to change the structural qualities of the model as in *rattle*.

- The models that can be realised in *rattle* are three-dimensional, i.e., the positions and movements of the interacting objects take place in a three-dimensional Cartesian space. Other software makes distinctions between one, two, or three-dimensional masses and forces, which cannot be mixed, in order to enhance the overall performance of the software. In *rattle* there are no such distinctions, as all particles share the same space and can always interact with each other.

- *rattle* was from the beginning been an *open-source* project. In contrast to some of the tools mentioned above, I have always thought of it as a tool which could and might be used by other researchers and artists. It was devised with the intention of influencing others' artistic practice, as well as eliciting feedback from the artists themselves. In this way it would be an instrument that stimulates exchange. *rattle* is open, in the sense that its development and, in some way, its inner construction, corresponds with an attitude of sharing instead of possession, of questioning instead of instrumentalisation, of openness instead of control.

  It has been openly shared and used by colleagues I've been working with. But effectively, even if the software (at least in its latest implementation) is publicly hosted on an online open-source software development platform,[6] I didn't explicitly work on disseminating it widely and it has therefore not reached many other computer music practitioners. It still retains many idiosyncrasies due to the fact that, until now, I was the only developer. This act of dissemination is something I would like to pursue in the future.

[6] https://github.com/davidpirro/rattle, accessed 01/08/2017

The points above shed light on the ideas behind *rattle*. It is not conceived as a specialised tool, apt for solving selected problems in a fast way and providing potential users with a pleasant and easy-to-use graphical interface. Rather, it is a tool that aims at providing a platform for formulating and simulating generic physical systems, offering flexibility for unthought problems. This aspiration of generality leads, on the one hand, to a wider range of possibilities and, on the other hand (the downside), greater complexity and a certain degree of resistance in its use and possibly worse computational performance. As a result, a greater degree of involvement is required from the user who wishes to operate it: *rattle* is, in its current form, a tool that requires the users to get their hands dirty, mess with code, try, fail, and debug. In other words, *rattle* was not conceived as a complete instrument or a plugin that presents some ready solutions or effects; rather, I conceive it as an environment that offers possibilities to formulate a particular set of situations that may bring emergent and unpredictable phenomena to light. *rattle* affords the experiential exploration of these situations.

Due to the generic problems the software tries to address, in conjunction with the requirement that the systems it simulates should be at all times accessible to the user or performer to interact with in any way, no *a priori* analysis[7] can occur to reduce the computational steps needed in real-time. Anything could change during the simulation, and nothing can be predicted in advance, meaning that the simulated systems should remain open to a continuous adaptation to changes. Every step in the simulation numerically integrates the motion equations of the involved particles - the differential equations describing their behaviour in time. This computational step introduces errors that might accumulate very quickly and ultimately drive the simulation into unstable states. Strategies to reduce this effect involve reducing of the timestep used (i.e., more steps are needed to compute the change of state of a system after a time interval) and increasing the precision of the floating point numbers' representation used or the order of numerical integration method used (refer to the section *rattle* integration algorithms for details). All these known recipes increase the computation steps required to calculate each simulation's timeframe.

[7] For instance the *GENESIS* software offers the possibility to run a modal analysis of the constructed object thus pre-computing the spectral structure of its vibrational behaviour. This step greatly reduces the computational power needed in real-time.

These considerations eventually led to the decision to switch from the *SuperCollider* language implementation, to a realisation in the *C* language. This would allow more fine-grained control over these aspects, and would also produce faster code. In this implementation, simulations could be run at a much higher rate and thus be able to (as I have noted above) *directly* synthesise sound. *Directly* here means that positions or velocity values of the particles, i.e., the state of the modelled system, could be immediately used as audio signal, unmediated by control to audio rate conversions, upsampling algorithms or mapping functions which connect the states of the simulation with some synthesis process which runs separately or in parallel. In *rattle* there is no distinction between control and audio rates at all, and there are no implicit signal hierarchies: sound synthesis and modelling/simulation have a very tight connection.

Embarking on this programming endeavour also meant not relying on a series of handy functionalities already implemented in the *SuperCollider* language. These now had to be re-implemented. On the one hand, this meant a great loss of energy and time, but on the other hand, seen in retrospect, the process eventually led to a tightening and a clarification of formulations both in the code and my own thought. It also enabled a better understanding of its functionalities: the need to re-programme necessarily led to a better understanding of their workings and side-effects. I would say that this situation eventually gave me the possibility to develop a deeper understanding of common computer music algorithms and a better control over my own practice by opening up and rewriting ‘black boxes’, which I would otherwise have used without having a precise understanding of them.

Finally, I would like to point out how this situation of ‘re-starting from zero’, together with my limited programming skills and limited time, led to the condensation of a sort of method: keep things simple, focus on the essentials, reduce to the bare minimum, and eliminate every algorithm or function that is not fully under control and could be responsible for unwanted effects. This philosophy infected my way of thinking in many ways and is also reflected in various dimensions of this thesis: from a bird’s-eye perspective, looking how the whole work evolved, and trying to identify threads running from the beginning to its realisation, I found an enduring spirit of *reduction*, the removal of non-controlled transformations and a search of the essential qualities of a specific situation (as imagined or experienced). This was a method of elimination of unclear concepts or terms, attempting to peel away the layers of interpretation and praxis that might blur the sought-after core.

### *4.1.1 Modelling Paradigms*

*rattle* is based on a *fields* formulation of the forces acting upon particles. In classical mechanics (we will ignore quantum or relativistic mechanics effects), a field is a function that associates a scalar value (i.e., a number) or, in general, a tensor (e.g., a vector) to each point in space: it is a 'condition of space'[8]. Temperature, for example, is a scalar field $T(x,y,z)$ that associates the temperature value to each point in space. Instead, the electric field is a vector field originating from a charged particle and extending into the whole space. In this case, the field function $E$ indicates which force $F$ would be experienced by a particle with unit charge at each point in space. So, in general, a particle with charge $q$ in the field $E$ would experience the force:

[8] Richard P Feynman, Robert B Leighton, Matthew Sands, et al. *The Feynman lectures on physics, Vol. 2.* Addison-Wesley, 1964

$$\vec{F} = q\vec{E} \tag{4.1}$$

where the field $\vec{E}$ is is given by Coulomb's law:

$$\vec{E} = \frac{1}{4\pi\epsilon_0}\frac{Q}{r^2}\hat{r} \tag{4.2}$$

Here $Q$ is the charge of the particle emanating the field and $\hat{r}$ indicates the unit vector in direction of that particle, the centre of the field.

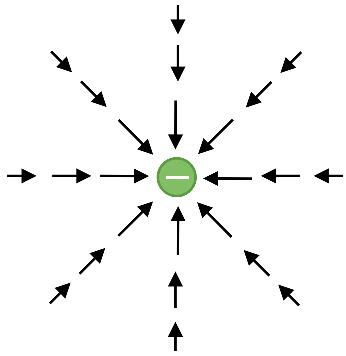

Figure 4.1: The electric field emanating from a negative charge

The electric field is an good example to illustrate this formalism, since historically *classical field theory* has been developed to formulate electromagnetic and gravitational fields. In the latter case, the field and the force acting on a mass $m$ would similarly be, by Newton's law of gravitation:

$$\vec{G} = -\frac{GM}{r^2}\hat{r} \tag{4.3}$$

$$\vec{F} = m\vec{G} \tag{4.4}$$

In *rattle*, fields are 'attached' to a particle. They are a property of particles, and a method in the particle object - in computer programming jargon, the *mass*. Most other particle-based physical modelling frameworks use the *links* metaphor for formulating interactions between particles. These links are in themselves objects, connecting two particles at a time. Using fields to formulate those interactions between particles allows (in many cases) for a more compact expression of complex interaction relationship networks in a model. Multiple particles might be under the influence of the same field (particle) and there is no need to specify a new *link* object for each of these interactions.

*Fields* in *rattle* are implemented as functions receiving one mass as an argument and returning the acceleration $a$

that the mass experiences as a result of being in that field.[9] Each *mass* is equipped with a default *field* function $f$ of the form:

$$\vec{a} = \vec{f}(p_0, p_1) = \frac{k_0}{m_1} m_0^{\beta} r^{\alpha} \hat{r} \tag{4.5}$$

where:

- $p_0$ is the particle at the origin of the field and $p_1$ is the particle it is applied to.
- $m_1$ is the inertial mass of the affected particle and $m_0$ is that of the field's origin particle. Note that $m_0$ enters the field equation as a parameter with an exponent $\beta$ in order to account for situations in which the origin's mass is part of the force's formulation (e.g., in the case of gravitational forces, see below). In the default case, $\beta = 0$, the origin's mass has no effect.
- $\hat{r}$ is the unit for the three dimensional vector that identifies the direction away from the origin particle $p_0$ and towards particle $p_1$.
- $k_0$ can be understood as a general coupling constant (or interaction constant), which controls the overall strength of the force particularly in relation to the other fields in the model, e.g., the spring constant in *Hooke's law* (see equation 3.3). This parameter also controls whether the force is attractive ($k < 0$) or repulsive ($k > 0$). The default value is $k = -1$.
- $\alpha$ is a parameter controlling the overall behaviour of the interaction force. For example, with $\alpha = 1$ (the default) the field would model the acceleration caused by a spring. With $\alpha = -2$, $\beta = 1$, and $k_0 = G$ the *gravitational constant*, gravitational forces could be modelled.

The idea of the above formulation for the field is to provide a possibility for modelling most of the fundamental interactions based on classical physics. Of course this formulation, even if very broad, would not be sufficient to cover all possible physical (and perhaps non-physical) possibilities for defining interaction forces. For instance, the field as defined in equation 4.5 could not model the magnetic force field; nor would it suffice for modelling the effects of non-linear springs or anisotropic force fields. Therefore, in the spirit of experimentation and openness, *rattle* gives the users the possibility to freely redefine the field function to fit most needs.

*rattle* also offers the possibility to assign more *fields* to the same *mass* so that it would affect other masses in different ways. This functionality, however, did not

[9] Technically, the *field* function receives two arguments as input, two masses; the first is always the *origin*, the mass that field belongs to (e.g., the centre of the electric field) as is automatically passed by the simulation callback.

prove to be essential in practice - the possibility to freely define the *field* function already provides enough flexibility for formulating the most diverse situations.

Of course, one particle might be under the effect of more fields of multiple *masses*. In this case, the 'superposition principle'[10] holds, i.e., the effects of the single force fields $f_i$ with $i = 1 \ldots n$ are added up in order to compute the acceleration $a$ for the mass under consideration. Eventually, a term for *damping* effects is added into the equation:

[10] Richard Phillips Feynman, Robert B Leighton, and Matthew Sands. *The Feynman lectures on physics. Vol. 1.* Addison-Wesley, 1963

$$\vec{a} = \sum_{i=0}^{n} \vec{f_i} - c\,\vec{v} \tag{4.6}$$

This velocity-proportional viscous damping force[11] ($c$, is known as the viscous damping coefficient) may be specified for each *mass* singularly and proves to be an essential variable.

[11] Richard Phillips Feynman, Robert B Leighton, and Matthew Sands. *The Feynman lectures on physics. Vol. 1.* Addison-Wesley, 1963

Many systems might exhibit instability in different ways. This is especially evident when they are interacted with, i.e., when a user acts upon a simulated model and therefore changes its state, injecting (or subtracting) energy from the system. The damping force is here a useful (or even necessary) tool to confine those instabilities that would drive the model into uncontrolled growth.

When all elements of a model are defined and in place, the simulation can begin. The simulation algorithm is a re-iterated process subdivided into two steps:

1. For each mass, the acceleration is computed using equation 4.6. A numerical integration using a symplectic-implicit Euler (or Verlet) algorithm (see appendix B *rattle* integration algorithms) computes the displacement and velocity variation vectors for the next frame (but these are not yet applied), as well as the effect of friction subtracted.

2. When the above step is finished for all masses in the system, the displacement and velocity variation vectors are applied to each mass.

These two separate steps are necessary to avoid the inconsistencies that would arise if the displacement of one mass would be applied before its effect on another mass could be computed. This would be especially dramatic in the case of two (or more) mutually interacting masses. Typically this would lead to chaotic behaviour, a result of the numerical integration routine's accumulating error, but not inherent to the simulated model.

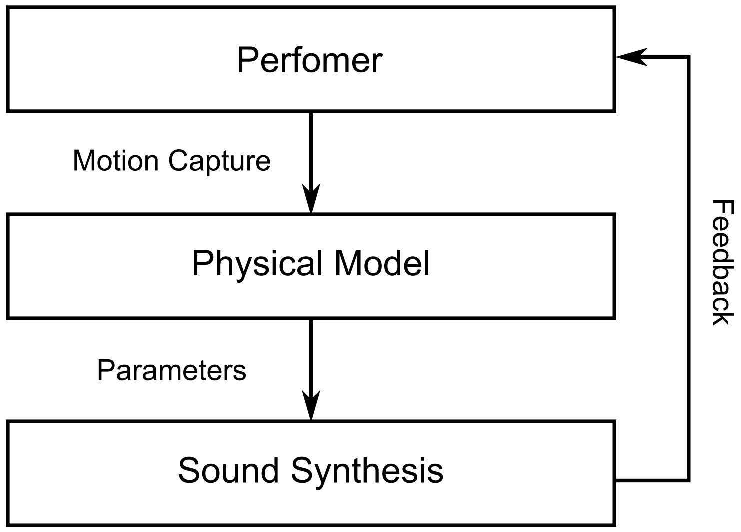


Figure 4.2: A simplified graphical depiction of the approach to the design of interaction used in the *simple spring mass scenario*

In the *rattle* framework there are three principal elements:

1. *masses*: these are the basic objects of the particle-based physical modelling. They contain the state of the particles, their three-dimensional position and velocity vectors. Further, each *mass* stores references - i.e., *pointers* - to the other *masses* in the system it exerts force on or it is affected. Each *mass* also stores other values controlling the form of the field it emanates (the coefficients $k$ and $\alpha$ and the mass $m$), as well as the friction force acting on it.

2. *fields*: these are properties of each mass. *Fields* are functions that implement the formalism of *classical field theory* followed here, i.e., functions that accept two *masses* as input, where one is the origin of the field and another is the mass that is affected by the field. Each mass has its own *field*.

3. *scenes*: a scene is a collection of *masses* interconnected by *fields*. Scenes are containers for subsystems in a model. Transformations (e.g., translation) applied to a scene are also applied to all elements in that scene. In other words, *rattle* scenes are a convenience tool to facilitate working with collections of masses. Furthermore, scenes can also contain specialised *fields* that are applied to all contained masses: for example, all masses in a scene might be subject to a gravitational force.

### 4.1.2 An Example and Some Considerations

To clarify how *rattle* has been employed to develop interactive scenarios, I will now introduce the *simple spring mass scenario*. This was one of the first experiments with physical models we performed in the *Embodied Generative Music* project (see section 4.2 The *Embodied Generative Music* Project).

In this scenario, as well as in the explorations that followed (e.g. A.2 cornerghostaxis#1 or A.3 Tball), the approach to interaction design can be simply formulated as in figure 4.2. The performers' state is captured via a motion capture system, and software then streams three-dimensional information about the position and orientation of an object held in a hand, a joint, or even the whole body, into a physical simulation running in *rattle*. A simple mapping function assigns this data to the state (i.e., position and orientation) of one of the elements in the running simulation. This *virtual* simulation space realises, therefore, a 'replica' of the real physical space being tracked: here the Cartesian axes of the two spaces correspond. The sound synthesis process parameters are updated according to the simulated system's state, producing a variation in its output. This variation is perceived by the performers, giving them feedback about the model's internal state evolution in response to their actions.

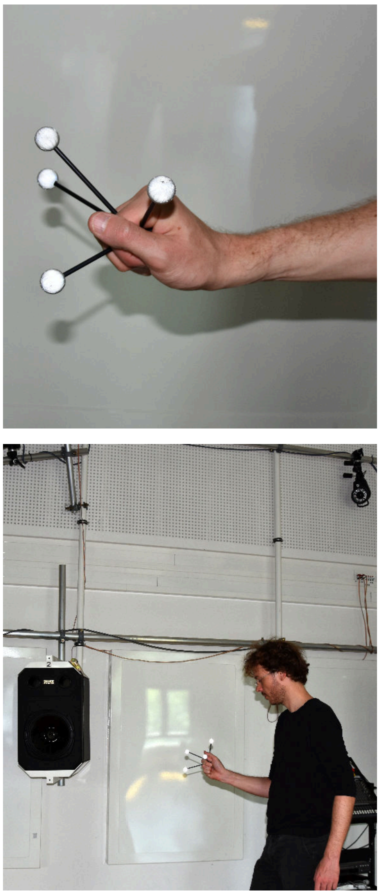

Figure 4.3: The *rigid body* tracking target (top) whose position is reconstructed by an infrared motion capture system by VICON in the CUBE Laboratory at the IEM (bottom, a tracking system camera in the top right corner).

In the *simple spring mass scenario*, a camera-based infrared motion tracking system captures the experimenter's movements into the physical model (see figure 4.3). More precisely, the tracked position of the hand holding a *rigid body* target is mapped to the coordinates of a particle in a simple model, in which two masses are connected with a spring (see figure 4.4). During the simulation, one of the masses (the filled black mass in the figure) follows the continuously changing position of the player's hand (via the tracked target he holds), while the other mass (the empty circle in the figure) is free to move, affected by the force of the spring.

Moving the hand, therefore, corresponds to a movement of the black mass and causes an elongation of the spring. The white mass would thenbe pulled and would start oscillating around the black one. As the movements and the model are in three dimensions, continuous and more complex movements by the hand would introduce more intricate paths. A degree of viscous damping is added to the system so that the oscillation would eventually fade out.

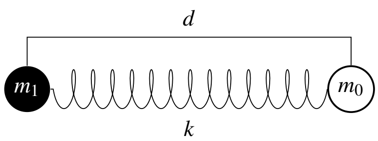


Figure 4.4: Graphical depiction of the *simple spring mass scenario*. The two masses $m_0$ and $m_1$ are connected by a spring (harmonic) force with Hooke constant $k$. The movement of the two masses is sonified by mapping the distance $d$ between them to the frequency of a sine oscillator.

The state of the system is translated into a simple (and crude) sonification by mapping the distance $d$ between the two masses to the frequency of a sine oscillator. It should be noted that sound is the only feedback offered to the user; there is no access to a graphical representation of the model's state while interacting with it. This scenario is of very limited musical interest, but is reported here as it offers some fundamental observations.

Even with such limited feedback, users could perceive the state of the simulated system and immediately establish a connection between their actions and the sound. This becomes especially evident when observing how quickly they could attune to the system. Anyone who tested this scenario could readily perceive that it was an oscillatory phenomenon they were confronted with and subsequently find and excite its resonant frequency. That is, after only a few hand movements, it is clear how to move the hand in order to keep the system constantly oscillating at the same pace performing only the smallest of movements. In its simplicity, this scenario shows how an interactive physical model may elicit immediate receptiveness, which then translates into a *bodily resonance*, where the connection between the perceived cause and effect is almost unmediated.

The role and effect of damping also needs some more attention. Damping was removed from the model from the very first test setting, which, rather interestingly, resulted in a scenario much more difficult to interact with. Each movement of the hand resulted in an energy injection into the system, which would then oscillate at its resonance frequency forever. In addition, the system would seem to continuously increase its energy and therefore the oscillations' amplitude, while the *felt causality connection* between one's own actions and their effects could not be established stably. Only after introducing the damping term could the behaviour of the system be fully grasped and performed with. The goal for those interacting with this simple scenario was to play with the resonant behaviour of the system, and it could even be said that the interaction consisted of exciting the model in such way to exactly counteract damping and remain in a stable energy regime. That is, damping not only seems to be a key factor in allowing one to grasp the system behaviour's at all, but also it could be interpreted as the parameter around which the users' movements revolve - a variable the body can immediately relate to and manipulate.

From a wider perspective, damping - as a (more or less) continuous energy loss - is a fundamental characteristic of any evolving physical system. Frictionless models are a common approximation in physics that are necessary in order to study and understand the basic behaviour of that system. Despite this, friction or energy loss is a constant force that plays a principal role in the world we are immersed in every day, a presence we feel in every action we perform in our reality. It is part of the *resistance* which we are accustomed to, and which we can experience and identify in any interaction we have with physical systems.

Practice and experience during the development of *interactive* physical models like the above 'simple spring mass scenario' has indicated that any of those systems, independently of the complexity of the model, need damping in order to be accessible to interaction or to resonate with the body. On the contrary, removing friction results in a simulated system that lacks any form of structural energy loss and therefore, the energy the interacting actor would inject into to system while engaging with it, would continuously accumulate. The resulting behaviour of unbounded growth is perceived as completely 'unrealistic' by the user or performer, even non-physical hindering a bodily understanding. These observations led to a few basic insights.

That the absence or the 'infringement' of basic physical laws is so clearly perceivable, leading to a sort of 'impossible' situation, might be taken as an indication of how much physical laws are ingrained in our perception - or, in other words, the extent to which our experience of the world and how it reacts to our own actions is embodied by physical laws.[12]

[12] That is, physical laws formulate what is already known by our bodies. At least what pertains to classical mechanics...

Also, the former observations might be taken as evidence for a perceptual propension to understand ourselves as part of a system, encompassing ourselves and whatever we may be interacting with. As opposed to being removed from or outside of a system into which we could inject energy at will, we are inclined to perceive ourselves as sharing and operating on the same energy balance our counterpart has access to. After all, our whole worldly experience is based on our being in a continuous exchange and with the physical world, the global system in which we are immersed.

## *4.2 The* Embodied Generative Music *Project*

Some of the paragraphs appearing in the following section and the next section 'Embodiment as inhabiting' are based on the paper 'On artistic research in the context of the project Embodied Generative Music' by Gerhard Eckel and David Pirrò, which appeared in the Proceedings of the International Computer Music Conference, ICMC 2009

The *Embodied Generative Music* Project was a research project hosted at the Institute of Electronic Music and Acoustics that ran between 2007 and 2010. The project was funded by the Austrian Science Fund (FWF) as part of its Translational Research Program. I was part of the core research team, also comprised of project leader Prof. Gerhard Eckel and researcher Deniz Peters. The project marked an important personal step as it was the first research project I was involved in at the *Institute of Electronic Music And Acoustics (IEM)*. Further, this dissertation, its initial thematic framing, and its conceptual and methodological foundation are strongly connected to the research issues addressed by the project.

The project's research questions and aims were based on early developments in electronic music and live-electronics in particular. More specifically, significant in forming the project's direction was the work of Joel Chadabe, who used an early real-time computer music system to compose and perform his pieces 'Solo' (1978) and 'Rhythm' (1980) with two 'proximity-sensitive antennas'[13]. He refers to the approach taken in these pieces as *interactive composing* (see also 2.2 Live-Electronics and interactive composing). This concept has grown out of his work since 1967 and is very close to the idea of an *Embodied Generative Music*, i.e., a type of generative music informed by the dancing body during its unfolding.

Other important approaches related to the objectives of the *Embodied Generative Music* project (now abbreviated to EGM) can be found in the work of Todd Winkler and Wayne Siegel. Winkler used Rokeby's VNS[14] system to create what he called 'motion-sensing music'[15]. Siegel explores 'rule-based composition'[16] in the context of the DIEM Digital Dance Project[17]. The development of SICIB[18] was also significant: a system 'capable of music composition, improvisation, and performance using body movements'. These technologies have become more accessible over the years, and a great number of works employing motion capture technology as in *EGM* (although not in a performance situation) can be found in various contexts, e.g., in gestural analysis and control[19]. Some of the questions raised by *EGM* touch upon movement sonification and are therefore also related to work in this field.[20]

The *EGM* project combined scientific and artistic research in order to further the understanding of the relationship between bodily and musical expression. In this endeavour, the research in *EGM* was driven both by a scientific and an artistic motivation. On the scientific side, questions concerning the role of the body in music creation, performance, and experience were approached from the perspective of music aesthetics. It is common sense that there exists a close relationship between the two forms of expression, one of which usually appeals more to the visual

[13] Joel Chadabe. Interactive composing: An overview. *Computer Music Journal*, 8(1):22-27, 1984

[14] David Rokeby / Very Nervous System: http://www.davidrokeby.com/vns.html (accessed 25/07/2017)

[15] Todd Winkler. Motion-sensing music: Artistic and technical challenges in two works for dance. In *Proceedings of the International Computer Music Conference*, pages 261-264, 1995

[16] Wayne Siegel and Jens Jacobsen. The challenges of interactive dance: An overview and case study. *Computer Music Journal*, 22 (4):29-43, 1998

[17] The Royal Academy of Music, Aarhus - DIEM:http://waynesiegel.dk/?page_id=214 (accessed 25/07/2017)

[18] Roberto Morales-Manzanares, Eduardo F Morales, Roger Dannenberg, and Jonathan Berger. Sicib: An interactive music composition system using body movements. *Computer Music Journal*, 25(2):25-36, 2001

[19] Frédéric Bevilacqua, Jeff Ridenour, and David J Cuccia. 3d motion capture data: motion analysis and mapping to music. In *Proceedings of the workshop/symposium on sensing and input for media-centric systems*, 2002; and Christopher Dobrian and Frédéric Bevilacqua. Gestural control of music: using the vicon 8 motion capture system. In *Proceedings of the 2003 conference on New interfaces for musical expression*, pages 161-163. National University of Singapore, 2003

[20] Alfred O Effenberg. Movement sonification: Effects on perception and action. *IEEE multimedia*, 12(2):53-59, 2005; Ajay Kapur, George Tzanetakis, Naznin Virji-Babul, Ge Wang, and Perry R Cook. A framework for sonification of vicon motion capture data. In *Conference on Digital Audio Effects*, pages 47-52, 2005; and Katharina Vogt, David Pirrò, Ingo Kobenz, Robert Höldrich, and Gerhard Eckel. Physiosonic-evaluated movement sonification as auditory feedback in physiotherapy. In *Auditory display*, pages 103-120. Springer, 2010

sense (literal body movement), whereas the other tends towards the auditory (metaphorical movement in music). As it turns out, it remains very difficult to characterise, understand, and explain the various forms in which the two are related in experiencing music, and how they are related in the creation of music. As a result, the main scientific objective of the project was to propose new elements of an aesthetic theory of the body/music relationship. This part of the project was mainly addressed by researcher Deniz Peters.

On the artistic side, and in the context of a *performance-oriented computer music*[21], the body/music relationship was approached from a *poietic*[22] perspective. We were interested in a kind of computer music that centres on a strong performative element based on the bodily presence and actions of a human agent. A main characteristic of this kind of computer music is the possibility it offers to dissociate the performer's movement from the sound production, making it subject to re-configuration and composition. The space for a poietic intervention in the relationship between bodily movement and sound has been already opened by early electronic audio technology in the late 19th century, but it was in the 20th century with the proliferation of digital computers and the tools they offered that this kind of re-composition could really become a central aspect of computer music. Although the possibilities for body/sound dissociation (e.g., transmission and storage or real-time synthesis and transformation of sound) have been used in music creation for a long time, the poietic questions associated with them are far from being clearly formulated, let alone them being systematically addressed or answered. The *EGM* project aimed to contribute to sharpening the questions associated with the *poietic* conditions of computer music production. In this sense, of central concern to the project was the following: through which means, and to what extent, are performers (especially dancers) able to shape the unfolding of a generative music composition through and with their living bodies?

[21] Guy E Garnett. The aesthetics of interactive computer music. *Computer Music Journal*, 25(1): 21-33, 2001

[22] The word poietic is used here to underline a fundamental quality of the perspective that has been taken in this strand of research in the *EGM* project, which is mostly based on productive tokens (from the Greek root of the term, *poieo*, 'to make') and creative processes, as an alternative to a theory-based approach.

In approaching this through various routes, both scientific (aesthetic) and artistic (poietic) questions are addressed, and there was a general acknowledgment that they could not be treated separately. Bringing a specific aesthetic aspect to light in an experimental setting means that poietic questions must be addressed when conceiving the setting. Poietic approaches to the body/sound relationship cannot be developed while ignoring their aesthetic implications.

The setting and research environment for the EGM project was the *aesthetic laboratory* (*ÆLab*). Physically, it was

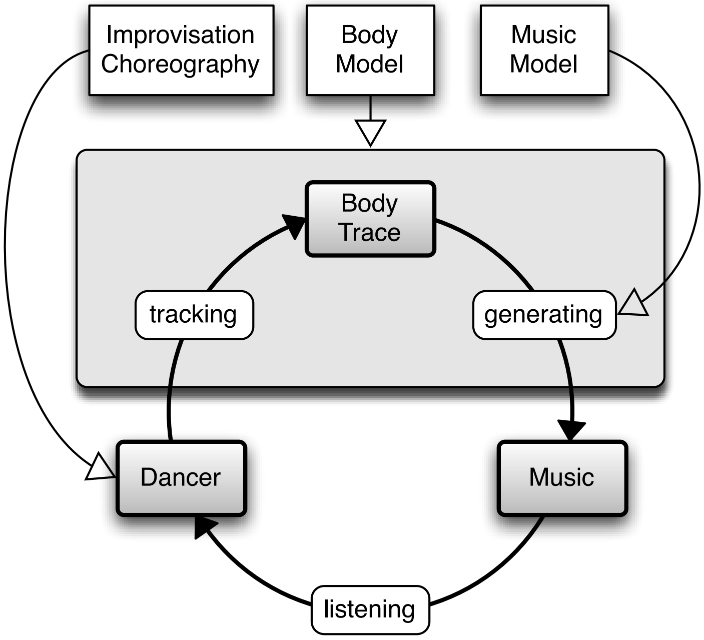


Figure 4.5: Schema of the conceptual and technical setup on the *Ælab*

installed in a 120$m^2$ studio space equipped with a 24-channel hemispherical Ambisonics-based sound projection system and complemented by an array of 48 ceiling-mounted speakers. Besides the sound projection and rendering infrastructure, a 60$m^2$ dancefloor and a *VICON*[23] motion-capture system with 15 infrared cameras was installed, allowing for high-quality full-body motion tracking. Working in the *ÆLab*, the dancer leaves a complex 'body trace' in time and space, which is then used to control sound and music synthesis (see figure 4.5).

[23] Motion Capture Systems from Vicon. Available: http://www.vicon.com (accessed 25/07/2017)

The 'body model' is inherent to the tracking technology used (see figure 4.6). The 'music model' represents the generative music composition. The resulting music naturally has a strong effect on the dance. In this tightly-closed loop, the dance is as much subject to the structure of choreography and/or the dancers' improvisational skills as it is driven by the music unfolding as a consequence of the dancers' movements - i.e., music that the dancers perform themselves.

As such an endeavour would probably change - or at least shift - the established understanding of choreography, improvisation, and composition, we approached our goal step-by-step in order to tackle its complexity. We reformulated our problem in terms of building a new instrument that could be played by the dancer, knowing well that the terms 'instrument' and 'to play' serve only as auxiliary constructs: we meant an instrument for playing on a structural level. An underlying assumption of this

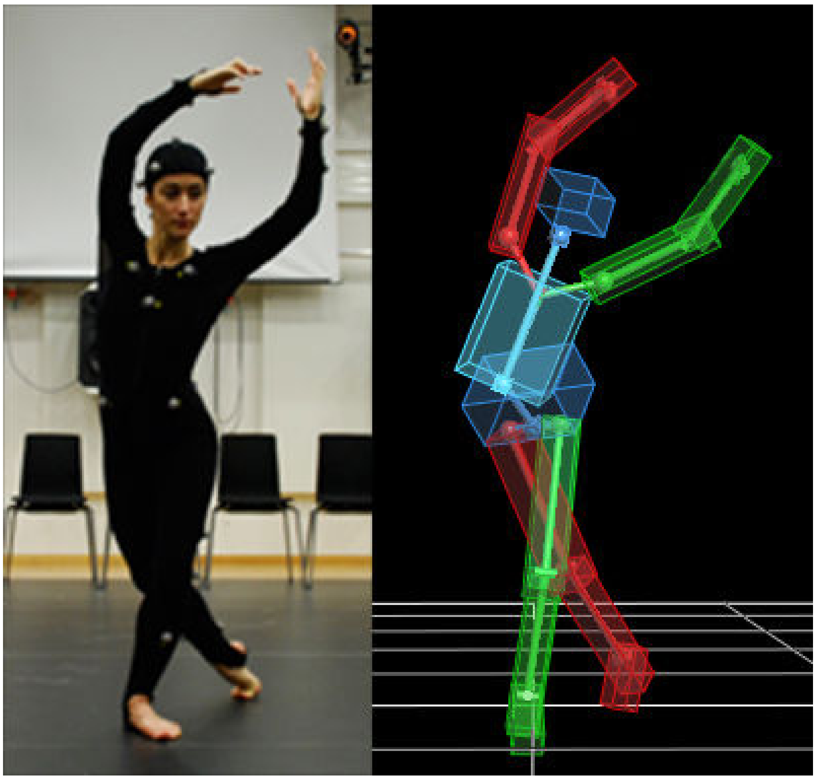

Figure 4.6: Dancer Valentina Moar in the full-body motion tracking suit (left) and the body model reconstructed by the Vicon motion tracking software

approach was that the expressive means and the bodily memory of the dancer's body would be best suited to fulfill our desires of an embodied generative music.

The first step in approaching our overall objective was to take our instrument metaphor literally and have the body produce sound. This was achieved by directly mapping the tracking data to sound synthesis parameters, thus achieving a kind of sonification of the dancers movements. Aside from solving the underlying technical and practical problems of full-body tracking and interpreting the enormous amounts of tracking data, this step also generated many new ideas and insights for possible approaches to the overall goal.

Methodologically, this exploratory part was organised as smaller units of experimentation that we called *scenarios* (or *case studies*). In each scenario we tried to address one specific means of connecting bodily movement and sound, and each was characterised by one particular mapping establishing such a connection. The fundamental idea of the *scenarios* was to further subdivide the complexity of the problem we wanted to tackle into smaller units, each realising simple, distinguishable and observable aspects. An analytical (i.e. a systematic) approach underlaid this methodological structure, operating on the assumption that each aspect could appear in isolation.

Mappings were simple, but offered parameters to be adjusted during the experimentation phase with the dancers. Alterations of parameters in such essential mappings would cause sensible deviations of their aesthetic qualities. Implicitly following a sort of *variational principle*[24], such differential relations would allow us to explore the aesthetic space of body/sound relationships. Eventually, this method would enable us to establish a stable connection between aesthetic experience and the mappings' formulation, thus making it available for composition.

[24] I use the concept of *variational principle* here in a metaphorical sense: in *mathematical analysis*, the original context of the concept, it stands for a general method for finding the functions which satisfy certain (extreme) conditions. That is, in general the principle is used to determine the underlying function given its observed variations relative to an independent variable.

As may now be clear, our methodology was strongly shaped by *aesthetic means*. Aesthetic criteria were used to determine whether a mapping was appropriate in eliciting a specific experience. Dancers can judge with confidence if a sound model and its motion mapping fit the movement or not (i.e., if the change in the sound feels right for a particular movement with respect to realising a particular idea). Motion mappings were developed through several iterations of an empirical process, in which dancers and composer informally assessed and discussed the quality of the mapping using their own embodied perception (aisthesis). The main measurement instrument in the *ÆLab* was therefore the aesthetic experience of the artistic researchers - hence the name of the lab. This experience, which was discussed among the researchers, forms the basis for the aesthetic judgement that determines the process.

After an intense period of exploring various kinds of motion-to-sound mappings with different dancers, we felt the need to summarise our findings in a short dance solo piece, which became *Bodyscapes* (see section A.1 Bodyscapes in the appendix). We found a big difference between this and an experiment in a laboratory situation, in which we may abstract from many aspects which are part of the problem we are treating in order to concentrate on a few central one. One could say that in trying to address 'researchable' questions by producing and showing artworks complicates the situation with respect to a more scientific approach, in which disturbances are reduced in order to measure and analyse results. But, when producing a piece that follows an artistic idea, we are forced to acknowledge all aspects of the production and performance - and their complex network of relationships - which raises different kinds of questions that otherwise would never have been asked (let alone answered). At least from an artistic perspective, these questions and the unrevealed perspectives they bear often are more valuable than quantitative results.

*Bodyscapes* has a special place in this dissertation, despite there being no physical models or dynamical systems involved in its composition. I mention this work because the decision to produce and perform this piece and was methodologically motivated at that time and had a great impact on the way this dissertation research was conducted afterwards. Most of the research work I present here is based on experiences gained during the production of artistic works. I could say that the production, exposition, and performance of artistic artefacts became, after this initial experience in *Bodyscapes*, a stable tool in the set of methods I have been working with. This method provided a balancing counterweight to the *scenarios* approach and the more analytical perspective it embodies.

### *4.2.1 Embodiment as Inhabiting*

The overall objective of the artistic research in *EGM* was to use new technology to develop new intermedial means of artistic expression combining dance and generative music, choreography, and composition. There are a number of research questions arising from this overall goal. But the core of the whole project, its driving force, was the utopian concept of an embodied generative music. Therefore, a central question was: in which manner may the dancer's movement influence the unfolding of a generative composition in an intuitive, or *embodied*, way.

The project is grounded on the concept of an *embodied interaction* (see 2.3 Embodiment), which framed how we observed and developed relationships between the dancing body and sound. Interpretations of how *embodiment* might be defined are numerous and can vary enormously depending on the research context in which they appear. In the context of the *EGM* project, however, we would understand *embodiment* as:

> *the extension of the dancer's body into the music – both on the level of the sound production as well as on the level of the unfolding of the compositional structure.*

We used the word *inhabiting* to describe this essential quality of the relationship between the dancing body and a musical composition. We imagine the dancers to be able to inhabit the music (as well as their dance). By this, we meant that they would know the music well, feel 'at home' in it, would feel at ease navigating it, and would be able to achieve a symbiosis of movement and sound, of dance and music, of choreography, improvisation, and composition.

The understanding we reached with the dancers is that a scenario can be thought of as a kind of *sound costume*. In this sense, a successfully composed scenario has to be *wearable* by the dancer. Wearing the sound costume will –

much like a real costume - highlight certain features of the movement and suggest a certain form of movement. It would suggest ways to use the sound-extended body. It may also constrain movement, which may or may not suit the artistic and aesthetic idea.

With respect to the sound production, *EGM* offered the dancer a kind of virtual instrument. For the dancer to be able to *inhabit* this instrument, a number of requirements have to be met - some of which were assumed essential at the outset of the project, while others were identified during the course of the project.

- Real-Time Requirements: Richard F. Moore's term 'control intimacy'[25] denotes a useful concept for illustrating many of the requirements that have to be met in order for an embodied sound generation to become accessible to a dancer. In his paper, Moore focuses on the temporal aspects of the problem - the time lag between performer action and audible result, and the jitter this causes. Both are very important in the case of *EGM* - the time lag had to be as short as possible and the jitter as small as possible. In the *EGM* setup we worked with a time lag of less than 20*ms* from movement to sound, and a jitter of no more than 5%. These values were measured with a *VICON* system comprising 15 M2 cameras covering a tracking volume of about $100m^3$ and running the iQ2.5 software. In most cases, a tracking rate of 120*fps* was used, at which the position and orientation data were provided by the system. Higher rates would have been possible at the cost of a reduction of the spatial resolution of the system, and this was soon found to be essential for embodiment to occur. At 120*fps* the system resolved positions in three-dimensional space with a precision of about 1*mm*.

[25] F Richard Moore. The dysfunctions of midi. *Computer music journal*, 12(1):19-28, 1988

- Real-Space Interface: As much as we had to provide the dancer with a real-time interface, the interface was also required to qualify as a real-space interface. Only the aforementioned spatial resolution and its consistent availability throughout the whole tracking volume could guarantee that even the most subtle movements of the dancers would be captured and translated into sound. The noise introduced by the system described here is of the same order of magnitude than the noise inherent to a dancer's body - this being a minimum requirement for embodiment to occur with most types of mappings, especially, of course, with space-based mappings. This aspect has been described very well by David Wessel, when he writes: 'Musical control intimacy and virtuosity require both spatial and temporal precision in the sensing of gestures'.[26] To keep the dancers' movement intact, key was to ensure that position and orientation information was available at least at the resolution of about 1*mm* for all body segments. This would allow for a sufficiently detailed and fully three-dimensional representation of the dancers' instantaneous posture. The quality of the posture's representation must be independent of the dancer's position and orientation in the tracking volume. The overall quality of a real-space interface is determined by its spatial precision, the size of the tracking volume covered, and its reliability (meaning that the system has to be able to track any posture a dancer may take).

One of the most successful scenarios in the *EGM* project was, departing from the above instrumental qualities, the so-called *Springer-Tempophone* scenario. In many explorations it seemed to offer dancers opportunities for being inhabited, and was the basis for many other scenarios developed during the project and for some scenarios in *Bodyscapes*. In its most simple incarnation, the mapping this scenario employs uses the tracked value of the position along one of the Cartesian axes that the tracking system draws into the space, e.g., the $x$ value of the three-dimensional positions of the dancer's right hand. The mapping function then appropriately scales this value and transforms it into an index in a pre-defined sound file. In other words, a space coordinate is transformed into a time coordinate. A granular synthesis algorithm extracts a small window or *sound grain* 100*ms* to 22*ms* long from a sound file at this time index: this sound grain is then reproduced in a loop. Changing the hand's position would cause an update to the window's position in the file, meaning a different sound grain would be produced.[27] The sound synthesis realises a sort of *Springer-Tempophone*[28] mechanism, an analog tape recorder/player which allowed to independently control playback speed and transposition.

At its most fundamental level, this scenario realises a simple connection between position in space and sound: moving in space means also to traverse and hear the sound contained in the file. Further, the *real-time* and *real-space* qualities of this instrument ensured that this connection was as *direct* as possible. Matching the body's own spatial and temporal perceptual sensitivity, in this scenario embodiment in terms of inhabiting could emerge. As an example, dancers repeatedly reported how strongly the sound in the recording structured their spatial awareness and even described experiencing virtual haptic illusions.[29]

[26] David Wessel. An enactive approach to computer music performance. *Le Feedback dans la Creation Musical,Lyon: Studio Gramme, France*, pages 93-98, 2006

[27] In its condensed form, this description attempts to transport the essential idea of the scenario. Still, mapping and the sound synthesis process were more fine-tuned in each new version. For instance, the length of the window was related to the height of the tracked joint in order to provide broader sound colouring; further, the central position of the grain was very slightly jittered so as to avoid the sonic artefacts that would appear when reproducing exactly the same sound bit in a very short loop.

[28] Peter Manning. *Electronic and computer music*. Oxford University Press, 2013

[29] Jana Parviainen. Seeing sound, hearing movement. In Deniz Peters, Gerhard Eckel, and Andreas Dorschel, editors, *Bodily expression in Electronic Music*, chapter 5, pages 71 - 81. Routledge, 2012

### *4.2.2 From Embodiment to Enaction*

Our experiences in the *Ælab* have shown that conditions can indeed be created for dancers to extend their bodies into the instrument. They can *embody* the sound synthesis processes, as per my description in the previous section. This could be realised by establishing direct or mostly unmediated sound to movement connections. However, those scenarios were limited to a rather *instrumental* understanding of the body-sound relationship. Looking back at the project now, I might say that it did not manage to realise an utopian *embodied generative music*, a form of interactive generative music in which dancers could embody the process – the algorithmic model by which music is generated – and not just the details of sound synthesis. Still, some of those experiences, especially the ones driven by an artistic approach as during the production of *Bodyscapes*, forced a re-thinking of the project's assumptions and revealed different perspectives that otherwise would not have been foreseeable.

The most important of these moments, and one that triggered many thoughts and questions that only now, after years, are becoming clear to me, is connected with what we called the *Delay Scenario*. This scenario appears in *Bodyscapes* as *The Partner* (see Bodyscapes), and departs from the aforementioned *Springer-Tempophone* scenario by introducing two critical modifications.

The first consists of a delay of a few seconds (e.g., four seconds) introduced between the dancers' movements and the sounds generated as a consequence of their actions. Dancers would move without instantly hearing that movement's sound, instead hearing it after several seconds. As they start to move in response to this material they would in turn generate a new sound, which would then appear only in the future.

This modification breaks the immediacy of the relationship between sound and movement. A new mediation, the delay, is introduced between action and sound, invoking a radically different situation, in which cause-effect linkages are tranformed. Such a transformation only appears to be possible on the basis that the dancers still be able to hear their own bodily agency inscribed into the sound, albeit delayed. This (delayed) agency gives consistency to the sonic output, allowing the construction of coherent perception. Even if the mediation impedes this perception, the reaction of the computer music system does not appear as random or unaffected by their actions. The heightened proprioception and bodily memory allows dancers to reconstruct a pathway into the past and, consequently, they can also learn to project sounds into the future through the actions they are performing in the present, sounds to which they will dance[30]. A play between expectation and correspondence can be experienced not only by the performers, but also by the audience. The play was taking place on a temporal level, meaning that even if only simple articulations of action/reaction and repetition/difference seemed possible, the introduction of a delay showed the potential for an interactive environment acting on the sound organisation in time.

[30] The dependence of the dancers' ability to cope with this situation with the length of the delay has not been studied systematically. Nevertheless, it is to expect that delay times approaching short-term memory duration (~ 18 seconds) would cause a sensible degradation of their capacity to reconstruct temporal relations.

Now, a further parameter is added to the scenario. The delay time is made variable and dependent on the speed of the performers' tracked joint. The dependence is inversely proportional, meaning that if the speed is minimal the delay is maximal (e.g., 4 seconds), while at maximum speed[31] the delay is minimal (0 seconds), and the sounds would be produced immediately. The changing delay value is smoothed with a simple integrator (a low-pass filter). The synthesis process does not compensate for the delay variation: as a consequence, the dancers' movement articulations produce clearly audible pitch shifting effects. Accelerating actions mean diminishing delay times, causing the delay line's read head to move towards the write head, therefore speeding up the reproduction of the sound in the buffer and eventually producing an upward transposition. A deceleration, however, would drag the read head towards sounds produced in past, far from the current sound, meaning a deceleration of sound play-back and thus a transposition to lower frequencies.

[31] The value of maximum speed has been determined experimentally and adapted to each performer.

In contrast to the effect of the delay, these artefacts appear instantly as the movements' articulations change. This version of the the scenario unites both mediation and immediacy, and these two aspects are non-trivially interwoven in such way that affecting independently one or the other is not possible.

The scenario presented performers with an environment where responses and temporal behaviour were far from the qualities of interaction we had searched for and experienced before. It was a situation that was difficult to grasp rationally or analytically; there was no certainty as to which output an input would correspond with; without the addition of the variable delay this was still possible, but at this point there was no clearly reconstructable cause-effect relationship. Rather, the sensibility of the system to both present and past events would induce variations in the output depending on the whole history of events preceding that moment, thus making every movement unique and almost unrepeatable. In a sense, the scenario had all the characteristics that we didn't want: it was neither simple nor it did present a clear and direct correlation between movement and sound.

This scenario clearly was more complex[32] than any other we had tested before. It was striking that a relatively small change in the algorithm induced such a profound effect on this scenario. It might seem an obvious observation: even a small change in a simple set of rules of an algorithm has the potential to produce unexpected and unforeseeable results. Still, experiencing this so dramatically evidences how fragile a controlled situation can be, how delicate and unstable. And moreover, how interesting this instability is.

[32] The term *complex* here is used to indicate a situation consisting of multiple interwoven and mutually interdependent variables.

Even more unexpected was that dancers could cope with this complexity. They could enter the scenario and establish a relationship; they could dance with it. Even if it was, at best, confusing when looked at with rationalising attitude, they could apparently *read* it with their bodily perception. They could '*grasp the dynamics of the system with their bodies*'[33] as Gerhard Eckel put it when describing this scenario.

[33] Gerhard Eckel. Embodied generative music. In Deniz Peters, Gerhard Eckel, and Andreas Dorschel, editors, *Bodily expression in Electronic Music*, chapter 10, pages 143 - 151. Routledge, 2012

Dancers engaged in a performance that continuously oscillated between action and reaction, togetherness and opposition. The sonic reactions were unpredictable, but they still exhibited a coherence, a felt *agency*, or a *behaviour* that could be grasped by both the dancers and the audience. This contributed to a crucial aesthetic change of perspective: the computer music system appeared as an actor in the environment, a source 'external' to the dancing body. These two actors - dancer and computer - exchanged, at times struggled, exhibiting different and changing modes of interaction that fluctuated between synchronicity and conflict. As such, the performance itself emerged as a process generative of different aesthetic experience of the connection between body and sound. It was the nearest to the imagination of an *embodied generative music* I could say to have seen during the project and I would say artistically one of the most rewarding experiences of the *EGM* project, both for the dancers as for the audience. In the attempt to understand how this scenario would fit into the project's frame, new questions demanded our consideration. How could this situation be interpreted in terms of our understanding of *embodiment*?

Juxtaposing this scenario with others previously developed reveals some implicit limits of our approach. Aiming at essential and traceable aspects of the connection between bodily movement and sound, we employed mappings establishing a simple and direct relation between the two. In retrospect and at a closer look, I would say that the paradigm we were still implicitly relying on was that of direct and unmediated *control*. We thought of the entire situation in terms of a traditional musician-instrument relationship, where virtuosity is directly proportional to the detail of control musicians would have over the instrument. Our assumption was that this kind of control was the substrate enabling the body to grow into the sound, to embody it. The *delay scenario* showed us a different mode of interaction, transcending the limits of this assumption. It was evident that performers were *not in control* of the details of the situation; neither they were completely disconnected or non-interacting. In fact, issues of control (who is in control? who is acting upon whom?) did not play a role for them or the audience. A central concern therefore materialises: how much does the (implicit or explicit) *control* paradigm inhibit possible alternative perspectives on interaction. And what concept could replace it?

Our idea of *embodiment* as inhabiting the sound synthesis, which would become a 'costume' for a moving body, involves the dissolution of specific qualities of the 'other' element, those which would make it appear as an artefact *present-at-hand* (to use Heidegger's language). Ideally, the computer music system would be a transparent interface completely permeable to the (total) control of the performer. In the *delay scenario* however, both actors retained their own identity: it is 'as if the music were an other creature dancing with her'[34], as Susan Kozel described dancer Valentina Moar in this scenario.[35] They *resisted* each other, but did not dissolve: this is central to this scenario. Thus, our idea of *embodiment* is revealed as insufficient for describing and framing the mode of interaction in this scenario.

An *enactive* perspective offers an alternative. The *enactive* approach (see 2.4 Enaction) holds that cognition is the result of a mutual, ongoing interaction between two entities. Performer and environment, both provided with *agency*, engage in a circular relationship. This is a temporally evolving connection in which the action and perception functions of each are interlocked. From this perspective, the *delay scenario* can be seen as a representation of an *enactive* process. Dancers are continuously challenged and re-adaptat their bodily understanding based on the environment's responses. Experiencing this scenario from the audience's point of view means to observe the (continuous) unfolding of the dancers' (and environment's) cognition.

The scenario seems to offer the possibility of being such an environment, of possessing the kind of agency or behaviour that enables this particular *enactive interaction* mode to emerge. It is, on the one hand, an agency not too

[34] Susan Kozel. Embodying the sonic invisible. In Deniz Peters, Gerhard Eckel, and Andreas Dorschel, editors, *Bodily expression in Electronic Music*, chapter 10, pages 61 - 70. Routledge, 2012

[35] This specific performance took place during the symposium *Bodily Expression in Electronic Music* (*BEEM*) which was the final event the the *EGM* Project, held in November 2009.

dependent on the performers, and therefore not possessing a clear identity: if it were too dependent on performers, the scenario would fall into a 'single actor' mode, where the computer music system is reduced to a merely reactive machine or instrument. On the other hand, it is not entirely free from external influence: the effect of the dancers' actions should not be too small in order to allow a perceptual correspondence to be established between movement and variations in the system's output. If this were true, the scenario would risk falling apart into parallel performances of two independent actors.

With the *delay scenario*, we have found, through an empirical process of calibration and testing, a good middle ground between these two possibilities. We can assume that a crucial factor in structuring the system's output appropriately is the fact that its temporal articulation contains the dancer's own bodily agency; transposed and distorted, perhaps, but still recognisable enough to provide agency. I suspect that the most important quality of this trace is that such temporal articulation moves in a range of variability still ascribable to a body.

Crucially, this scenario and the switch to an *enactive* perspective opens new ways to conceive and develop interaction in computer music praxis and performance. These are modes that allow generative processes to become tangible and permeable to performers' actions on a structural level, rather than as merely instrumental. Therefore, pursuing this direction, the question at this point is: to which extent, and with which conceptual and practical tools, could such agency be *composed*, not just found? And what are the conceptual implications of this *enactive* approach?

## *4.3 Dynamical Systems Thinking*

Departing from the experiences made in the *EGM* project (and especially with the *delay scenario*), we began to move towards an *enactive* perspective on interaction. This seemed to open up an integrative view on both the bodily aspects of performance and the intrinsic generative computer music processes. This move is motivated by the observation that *agency*, as a perceived quality of the computer music system, would play a critical role in constructing and defining interactive scenarios: agency is an essential ingredient to the theory of enactive cognition.

It is also important to recall and clarify here that, at the core of this inquiry, is the question of how computer music systems can be interactive, and what are the conceptual and performative consequence of the development and composition of these systems. The understanding of computer music systems I employ here is not limited to an instrumental one, but includes and centres on the *generative potential* of algorithmic processes such systems may spawn. Further, *interactive* here means that those systems, formulated for musical performance and composition, should involve the performing musician or composer they are put in connection with, transcending an attitude of *control*. Rather, they facilitate and require a deeper cognitive and physical effort by tapping into the very constitutive building blocks of perception. In this sense, I regard this research as a continuation on a slightly different premise from that of the *Embodied Generative Music* project from which the core research themes are inherited.

In the previous parts of this chapter, the term *behaviour* has been used for referring to a distinguishing temporal evolution, e.g., of simulated physical models. The intuition behind this section is that this idea of behaviour - as in general exhibited by dynamical systems (of which physical models are a subclass) - correlates strongly with the idea of an agency as observed and imagined in the *EGM delay scenario*. Following this path, a collision is staged between the two concepts on the grounds of the definition and characterisation of agency as it can be found in enactive cognition theory (see 2.4 Enaction). I expect that this convergence would bring forth a sharpened definition of agency: as the quality of a computer system which sustains interaction on the level of processes rather than on the level of states. The intent here is to refine a formulation to stabilise a conceptual thinking framework, which would also offer a concrete tool set for actual realisations.

To perform this collision, at this point a more-or-less precise definition of *behaviour* is needed - or at least a clarification of the meaning of the term in the context of this dissertation. I borrow the term behaviour and its meaning from physics or mathematics, disciplines which also lack a clear definition of the term. In these fields, behaviour is used to indicate ‘how’ a function or a system evolves from one point or state to another. For example, how the function $1/x$ reaches 0 when $x$ tends to infinity is a behaviour proper to that function and to that function only. Or how the velocity of a mass $m$ attached to a spring changes periodically in time - this is the behaviour specific to that system. With behaviour, I denote the way the state of a system changes from one moment to the next, or from one coordinate to the other. It indicates the unfolding of changes, the time-ordered variations of a system when it proceeds from one state to the other. Behaviour is made up of the differences produced by the system and is observed given a starting condition. Most importantly, behaviour is an identifying characteristic of a particular system: that is, all oscillators (e.g., mass-spring systems) exhibit similar and recognisable behaviour and every oscillatory behaviour may be ascribed to the evolution of a dynamical system of the mass-spring class.

What I am trying to establish is a bridge between dynamical systems and aesthetic and perceptual qualities. Behaviour, in the definition above, is the path of evolution of a system, where its particular temporal unfolding is as perceptible a quality as its colour. But there is also the idea of behaviour as found in physics or mathematics, which is inscribed into a precise mathematical formulation. Of course, these two meanings do not have to coincide and it is not my intention to equate them: only that in both cases the word behaviour refers to the temporal structure (or dimension) of things. In some way exploiting the somewhat unclear meaning which the term covers in both contexts, I attempt, therefore, to establish a correlation between the two fields of dynamical systems: mathematical modelling and the enactive cognition theory. The final (and maybe utopian) aim is to develop a language for formulating behaviours which could then be translated into programmes or algorithmic entities: mathematical formulations offer this possibility.

To be clear: I do not assert that any perceived behaviour can be readily transposed into a mathematical formulation of a dynamical system. The assumption here is that *the temporal behaviour of a mathematically formulated dynamical system has a perceptible correlate*. Most of the works collected in the Appendices (see Appendix A, A catalogue of works) might be seen as experimental (and experiential) studies that put the former statement to the test via different gradations of intervention by an external performer. These range from interactive performer-system settings, to reactive installations, to acousmatic pieces evolving without (or with little) performer influence The qualities of perceived behaviour are the central aspect explored by these works and their connection to the underlying formulations in terms of physical models or, in general, dynamical systems.

Now, as anticipated, I will attempt to bring together the two terms of behaviour and agency, based on the characterisation of the latter given in enactive cognition theory. Three qualities are fundamental: *individuality*, *activity*, and *adaptability* (see section 2.4 Enaction).

1. Individuality: The system exhibits a clear and perceivable identity. As seen in multiple artistic case studies, this characterisation strongly resonates with a sensible quality of the behaviour of dynamical

systems I have often described using the term *coherence*. What I would like to delineate with this word is the felt consistency of the system's evolution: that is, that the perception of the temporal path drawn by the sequence of states between two chosen points $A$ and $B$ has something in common with, or is similar to, the evolution from $B$ to a later point $C$. It corresponds to an intuition of a constant evolution rule which lies 'behind' that path, that which is driving the system. In some way, it is an affordance the system's evolution offers for being *integrated* into a coherent perceptual *image* (see also 3.4 A sense for change: behaviour).

In some way, stating that a coherent system might be sensed or perceptually reconstructed from its temporal evolution, is similar to affirming that the differential equations governing a dynamical system can be perceived from its behaviour. Not so much formulated or reconstructed in a mathematical form, but rather sensed in their presence and in their specific characteristics. The specific way dynamical systems structure time seems to be perceptible, but this also contributes to an identity that can be differentiated from other systems. This form of perceived identity might also come into relation with the specific form of geometrical flow inscribed by a dynamical system into *phase space*.

This observation seems to be valid for both isolated systems and systems interacting with an external agent or performer. In the latter case, I would describe *coherence*, as the consistency the performer may feel between effects and their causes. That is, even if the effects exhibited by a dynamical system as a consequence of action are not always precisely predictable, they can be nonetheless brought into a sound relation: the surplus of non-predictability becomes part of the system's agency or individuality.

Another observation is that dynamical systems yield a space of potential behaviour that is not infinite. Even if there is a wide range of possibilities, viable paths lie in a bounded space. Not every temporal evolution, therefore, can be followed or produced by the system. This is a fundamental quality of such a system and is essential in understanding how their identity may be perceptually constructed. If every behaviour would be possible, no underlying coherence could be reconstructed: limitations and constraints are an indispensable ingredient.

Concluding: temporal behaviour, as it is generated by a dynamical system, can be brought into relation with a perceived quality of the agent's identity.

2. ACTIVITY: An agent is a source of energy for the coupled system, meaning that it acts in absence of an external excitation. An agent does something, and is a source of excitation for the environment it is in.

   This can be regarded as a fundamental difference to the systems I have described previously in this chapter, namely the *simple spring mass scenario* (see section 4.1.2 An example and some considerations) and the *delay scenario* (see 4.2.2 From embodiment to enaction). In both of those cases, the performers were confronted with an interactive environment that reacted to input, showing a response which then decayed over time. In the first example, due to the system's attrition setting, the spring's oscillations would fade away. In the second case, if the dancers were still, the computer music system would not produce any sound by itself.

   However, it seems obvious that a process should have a certain degree of independence and action in order to be perceived as an agent. Further, *activity* - as the synthesis of an acoustic output independent of an input - might be seen as a precondition of generative processes as they are understood in computer music. That is, processes that unfold their own temporal structure according to certain rules.

   Experience with artistic works has shown me that dynamical systems might well be modelled and simulated in such way that they could be a source of activity, or generative in a computer music sense. In terms of the mathematical formulation, from this activity it would follow that such a dynamical system would not have an asymptotically stable fixed point at its origin (see section 3.1 Theory). From the presence in the system of such a type of critical point it would follow that the system would, sooner or later, fall into it, remaining in that state indefinitely - or at least until re-excited again. In other words, the system would be built around a more instrumental conception. A simple attractor that would fit this description would be the centre attractor, which is paradigmatic for oscillatory phenomena (e.g., the undamped oscillator or the *limit cycle* attractor).[36] The system should, at least, have a lower boundary for energy, which would prevent it falling into a fixed point and no longer able to move.

   This activity quality, therefore, necessitates narrowing down its possible behaviour types. It suggests a dynamical systems without asymptotically stable critical points.

3. ADAPTABILITY: This quality refers to the connection the agent has with its environment or the other actors in it,

[36] At present it is unclear if from the request of activity in the sense described here, together with the condition of boundedness (i.e., the state of the system cannot become infinite) would strictly mean that the attractor types and behaviour addressed here can be only be oscillatory. This aspect should be addressed in future research.

and how these relations may influence it. Adaptability posits a *coupling* of the agent's system with the environment, therefore addressing the role of action and reaction, and of *interaction* in the agent system. It is therefore of central interest here.

From the perspective of dynamical system modelling, adaptability requires the behaviour of the system to be influenced in some way by the state changes of the external environment or by the actions of any other agents it might interact with. How can this influence be better formulated? It is clear that an external input affects the system's behaviour, but how can this happen? With reference to the diagram 2.3, how does the bottom coupling arrow enters the system's constituting process? A good example is, again, the *simple spring mass scenario* (see 4.1.2).

As a dynamical system, we are looking at a simple linear system, which might be written as:

$$\dot{\vec{u}} = A\,\vec{u} \tag{4.7}$$

where $\vec{u} = (x, v)$ is the state vector of the system, $x$ and $v$ the position and the velocity of the mass respectively, and $A$ is the *Jacobi* matrix:

$$A = \begin{pmatrix} 0 & 1 \\ -k & 0 \end{pmatrix} \tag{4.8}$$

With this definition this formulation would therefore reduce to equation 3.9. Extending the dimension of $\vec{u}$ and modifying $A$ according to the problem at hand, equation 4.7 is a valid formulation for any general *linear* and *autonomous* dynamical system (see section 3.1 Theory). Including the external influence on the system (i.e., the moving and tracked hand position in the example of the *simple mass scenario* (see 4.1.2)) means to include in this formulation a *time-dependent* external component.

$$\dot{\vec{u}} = A\,\vec{u} + G(t) \tag{4.9}$$

where $G(t)$ is a function that, for this simple example, could be re-written as:

$$G(t) = \begin{pmatrix} 0 \\ kx_h(t) \end{pmatrix} \tag{4.10}$$

where $x_h(t)$ is the (time-dependent) function of the hand's position.

The inclusion of the external influence $G(t)$ on the system, from a mathematical perspective, has a qualitatively dramatic effect on the type of problem we are now concerned with. Mathematically, any such system would now become a *non-autonomous* dynamical system. This category of problem has completely different qualities than the the autonomous systems we have considered so far. Even if the form of the function $G(t)$ would be known in advance, formulating a solution to these problems is much more difficult . In general, for most of these problems, one cannot find mathematical solutions at all: strange attractors and chaos lurk in every corner. The only possibility is to simulate such systems: that is, to computationally evaluate the system's state time-step-by-time- step in order to reconstruct an image of its behaviour. And this is what we do here, especially as, in our case, the form of the external input's function $G(t)$ cannot be known and is dependent on the performer's actions.

Now, having this formulation, we are left with another important consideration. From the *simple spring mass scenario* and the formulation we have found above, one can see that the external input enters the system by modifying its *flow*: that is, it influences how the system will evolve in time. Typical coupling models, or more commonly mappings, between external action and computer music system, would consist of a functional relation of some kind between input and internal *state*. Returning to our previous example, this would mean that moving the tracked hand would produce a change in the state of the simulated mass, for instance by directly translating it, and placing it instantaneously onto another phase state path. In the present case, such an approach would mean to momentarily (for the duration of the input) suspend the system's own behaviour in order to move it into the desired state.

In light of the previously discussed qualities of identity and activity, a direct manipulation of the system's state should therefore not be allowed. Any kind of direct manipulation would mean suspend, even if partially, the system's identity and activity in order to apply instrumental control. It would mean to apply an ideally infinite force to the system from an external god-like perspective over the system. No equal partnership between a computer music agent and a performer, like that seen in the *delay scenario* (see 4.2.2), could be realised on these grounds.

Instead, influencing the system on the level of its evolution would allow its output coupling (again referring to the diagram 2.3) - being a function of its state output - to inform both of its internal dynamics and of the effects of the input. Perceiving or hearing both of these aspects in the system's output is, as experience has shown, fundamental for understanding the agent's behaviour. Again, taking the *simple spring mass scenario* as model, the performers moving their hand would induce an output response that is determined both by the system's own dynamic and by their input. That is, the performers can perform a sort of *perceptual impulse response* testing of the system, by listening to the effects of their actions *through* the system's output. This would then construct an *image*, of sorts, of the system's internal dynamics. Of course, when the systems composing the agent become more complex, non-linear, and *active* in the sense described above, it is not possible to speak of impulse response testing in a traditional sense, i.e., as in classical signal processing. Still, I believe that as a metaphor the idea holds for the purpose of explanation.

Returning to our starting point, we see that the question of adaptability refers to much more than simply the nature of the agent's coupling with the environment. Adaptability also indicates the agent's ability to *modulate this coupling*. According to the enactive cognition theory, this modulation is performed by the agent aiming at maintaining its *norm*, which in general corresponds to the self-maintenance of the processes which constitute it (see 2.4 Enaction). Further, this modulation has to be a function of the system's state. In terms of a dynamical system and extending the previous formulation at equation 4.9, this would mean:

$$\dot{\vec{u}} = A\vec{u} + H(G(t), \vec{u}) \tag{4.11}$$

where $H$ is the function modulating the effect of the external input $G$ into the system. This is not only dependent on time, but also on the state vector $\vec{u}$. In general, we can assume that the function $H$, the time-dependent modulation of the system's coupling with the exterior, should once again be a dynamical system. With this change, the above system becomes not only non-autonomous, but also *non-linear*: even if we suppose a very simple linear system at the core of the agent's behaviour (e.g., in the example above) as a consequence of the request for adaptability, this system is drawn again into a qualitatively different class of dynamics. The behaviour it could exhibit, especially in concert with the external environment or agent, would be qualitatively different: complex, emergent, and chaotic.

### *4.3.1 Phase Space Thinking: An Experiment*

In the preceding sections a better understanding of the concepts of agency and behaviour has been gained, and their reciprocal connections have been delineated. At this point a fundamental question can be addressed: *How can agency and behaviour be composed?* Specifically, in detail, and grounded on our previous insights about the specific definition of enactive agency, how could this could be reformulated in terms of dynamical systems behaviour? Or, in even greater detail: *How can dynamical systems be composed such that their behaviour generates agency? How can these dynamical systems be employed in the composition of interactive computer music environments?*

These questions aim at the development of tools or a framework that allows us to realise the ideas above, putting those to the test in practice, exploring the space of possibilities they could provide. At this point, we lean on an established and existing tool used to observe and analyse dynamical systems qualitatively: the *phase space representation* (see 3.1 Theory).

This representation reformulates the systems of differential equations, defining dynamical systems into geometrical structures, or so-called *attractors*, spatial movements, and *vector flows*. These representations have the indubitable advantage of being able to transport abstract mathematical formulations into a more directly sensuous experience. They provide alternative access to the qualities of dynamical systems based on the visual sense: I would claim that those figures and geometrical diagrams have an immediate bodily correlate. By looking at these, one gains a sense of what a dynamical system is 'up to'. These qualities of the phase space representation can be readily experienced when looking at a well-known book by mathematician Ralph H. Abraham and visual artist Christopher D. Shaw: *Dynamics, the Geometry of Behavior.*[37]

[37] Ralph Abraham and Christopher D. Shaw. *Dynamics-the geometry of behavior*. Addison-Wesley, Advanced Book Program, Redwood City, Calif., 1992

It is important to note here that the representation of a system's behaviour that its phase space provides is not just a more or less faithful 'image' in geometrical terms. It is well known that phase space and the mathematical formulations using systems of differential equations are *isomorphic*, meaning that they contain exactly the same information: they can be used as equivalents. Phase space and in particular the 'construction' of the *Lorentz attractor* as it is depicted in the book cited above might serve here as a metaphor for illuminating which utopia I am pursuing.

With reference to the following figures, we may at first look at the computer music system and the performer as

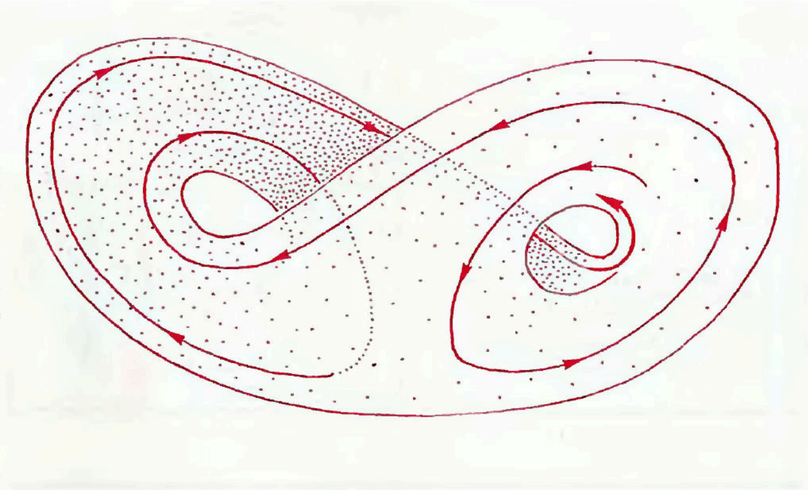

Figure 4.7: One of the Lorentz attractor's phase space representation: Abraham & Shaw *Dynamics - The Geometry of Behavior*, p. 286

two disjointed dynamical systems (see figure 4.8). We may regard these first two dynamical systems (performer and computer music system) as being of some type *A* (three dimensional attractors of saddle type with spiral outset in the figure). A second attractor *Y* (a saddle attractor with nodal dynamics on its inset in the figure) could then be the phase space representation of the dynamical system of the coupling between performer and system.

Now, setting two *A* type attractors (a performer and a computer music system) in the same phase space, we see that their mutual interaction through coupling *Y* will provoke a transformation and deformation of the phase space flows exerted by the individual attractors into a global flow (figure 4.9). Eventually a new attractor emerges affecting the whole phase space.

This combination produces a new unitary and complex structure, or a new dynamical system: the single attractors from which we departed are still there, but are now reciprocally modulated by the other systems occupying the same space. Their intertwining produces something new, which cannot be decomposed into a sum of the effects: the single attractors are instrumental in generating this new system, but at the same time they disappear as separable elements. In other words, the idea is to create the conditions in which performer and computer music system are enabled to mutually interact such that their joint evolution might result in a coherent and synchronous dynamic evolution.

This idea, based on an understanding of interaction as a coupling between dynamical systems in phase space, has been put to the test in the context of a small case study, the *phase space experiment*. In this case study, a relatively reduced experimental setup was realised, in which a performer was asked to interact with a computer

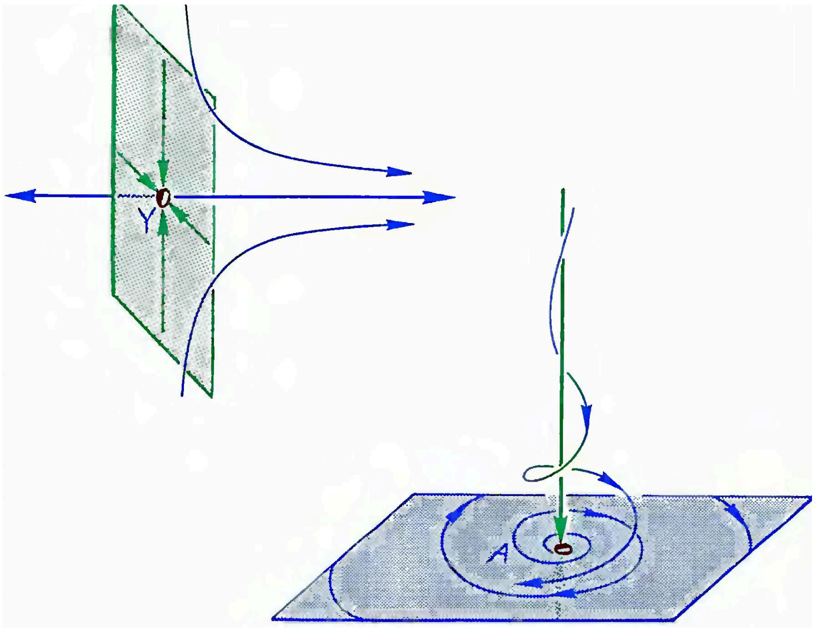


Figure 4.8: Two three dimensional attractors: *Y* saddle attractor with two dimensional *inset* (node), *A* saddle attractor with spiral *outset*. Abraham & Shaw *Dynamics - The Geometry of Behavior*, p. 383

music system (in the following CMS) whose sound output is modulated by the evolution of a simple two-dimensional dynamical system (*centre* attractor, the prototype of all harmonic oscillator systems), which is in turn perturbed and influenced by their playing. The aim of this experiment is to observe:

- If and how the performer's reactions to the sound produced by the CMS are informed by the specific attractor type used in designing the system.
- How salient the behaviour induced by the attractor is for the performer.
- If phase space structures give rise to perceptually distinguishable musical gestures.

Again, it is important to note the relevance of global behaviour of the combined system. Does the behaviour of the whole coupled system (composed of the CMS and performer as experienced by an audience) present significant and identifying characteristics (evident when the dynamical system implemented in the CMS is varied)? If it does, this could point towards an effective employment of the phase space thinking model in the composition of interactive live-electronic environments.

The case study was too small and too reduced to be considered a fully-fledged experiment providing clear scientific insights. Still, going through its implementation and witnessing the reactions of two professional musicians (Saxophonist Joel Diegert and

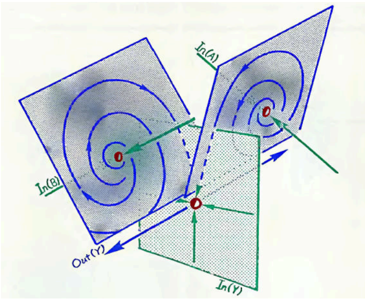


(a)

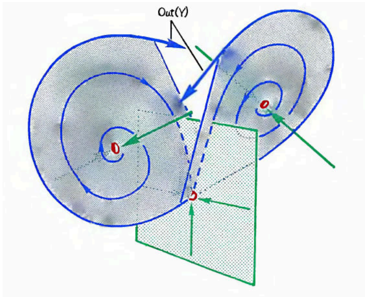


(b)

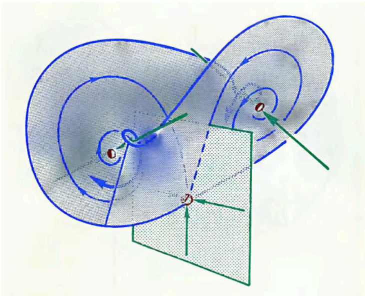

(c)

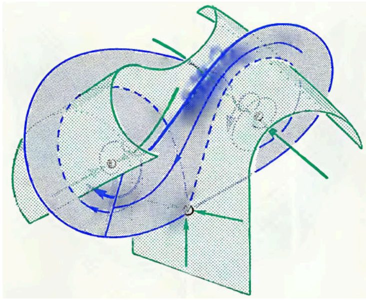

(d)

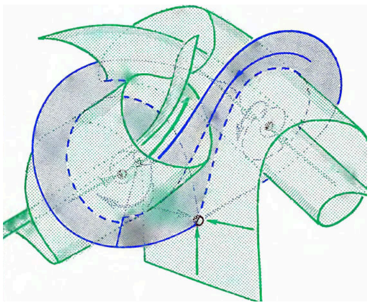

(e)

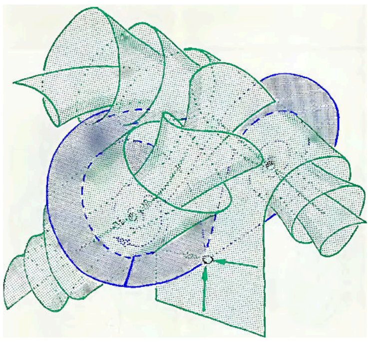

(f)

Figure 4.9: The ‘construction’ of the Lorentz attractor by the interaction of 3 different attractors. Abraham & Shaw *Dynamics - The Geometry of Behavior*, pp. 384-389

Violinist Lorenzo Derinni, see figure C.2) exposed to such kind of interactive environment yielded precious experiences and observations that could be the basis of further investigations in the future.

Both musicians had experience with works including live-electronics to different degrees. They were asked to play, hear, and react to the CMS's sound and find their own way to interact with it: they were not given a prior explanation of the system's functioning. Eventually, in informal interviews after the testing sessions both stated that the kind of interaction they had felt had definitely different qualities to what they were accustomed to in previous situations.

- Clearly for both, what they heard coming from the CMS was not only an effect. Rather, the interaction was felt not only as one-way causal relationship, where the sound they produced was the only source of activity. The system was perceived as having some kind of own activity.

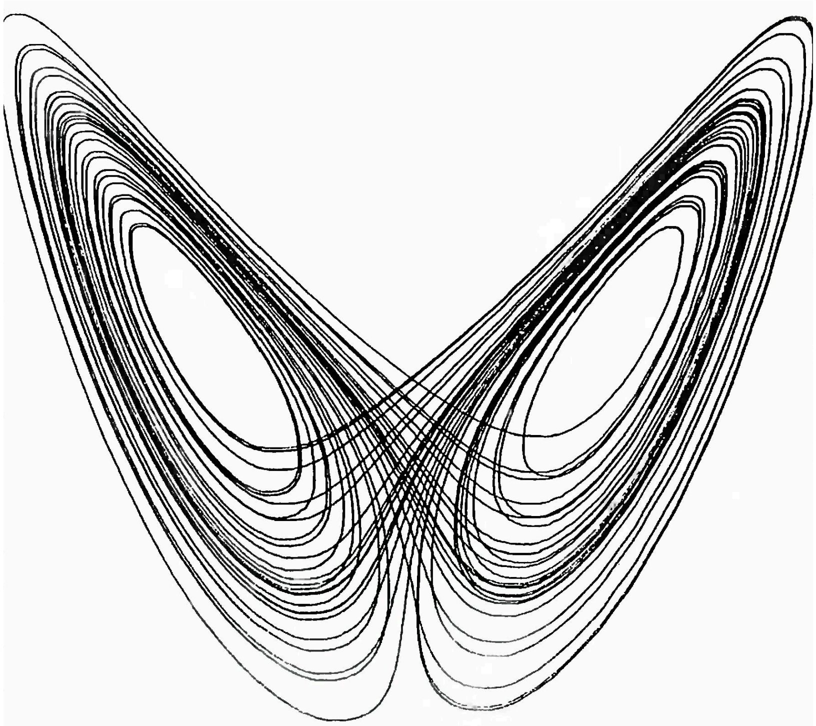

Figure 4.10: The phase portrait of the Lorentz attractor. Abraham & Shaw *Dynamics – The Geometry of Behavior*, p. 387

- Even if the dynamical system used in the model was very simple, they could not exactly formulate in words this behaviour. But, they clearly felt a sort of own *will* in the CMS, which they could influence at times more or less effectively.
- After some initial adaptation time, they showed some synchronisation to the base system's evolution in their playing: a sign that its behaviour resonated with their perception.
- Both musicians underlined an aspect that was not really clear prior to the test: It is not possible for the performer *not to interact*. There is continuous *contact* with the CMS; they could not choose to simply 'back off'.
- They defined the system as very *sensitive*; clearly any kind of movement or sound they produced was reflected as some modification in the CMS's reaction.
- Despite the immediate perception of activity, behaviour, and sensitivity, finding a working interaction mode is difficult (and maybe frustrating at times), but intriguing.

Especially interesting was the experience while working with violinist Lorenzo Derinni. After some hours of testing, a different mindset made its way into his play. At that point, the circularity of the interaction between system and performer became evident – as though the performance became simultaneous playing *and* listening. After this experiment, he also reported that at that moment he experienced a *heightened sense* of hearing, of his instrument, and of time.

The professional musicians involved in this study were highly sensitive to sound and very detailed in producing it. We did not consider, however, that even if asked to concentrate only on the sound and ignore more musical approaches for the test, for them these two aspects (sound and musical structures) are inseparable. Therefore, for both musicians, the scenario was too simple, meaning that they would expect more variation in the system's behaviour.

Further, both asked for more sensitivity from the CMS, especially towards small impulsive or rhythmical structures: the implementation of the system used smoothing (low-pass filtering) at various levels of signal conditioning in order to keep it in a more stable region, bounding the system's evolution. This choice, however, limited the sensitivity of the system. This meant that the musicians both asked for faster reactions and actions by the system.

Yet, on this premise, further development and a more extended and systematic study seems to be promising. Suggestions and observations should be taken into account in a future experiment. Also, additional questions and features could be addressed – for instance, if and how the affect of qualitatively different attractor types could be observed. The focus should remain also on the qualities and consequences of the coupling: this aspect should be the object of a more in-depth study. Questions pertaining to how this coupling and, in particular, its modulation as part of the dynamical system's agency are surely central and possibly, in view of these collected experiences, even essential.

# 5

# *Conclusions and Outlook*

This thesis develops some ideas about interaction in the context of computer music. The interplay of three elements forms the basis for this work:

- A scientific and theoretical analysis of the concept of interaction and of how these are understood in relation to computer music, particularly with regards to theories of *perception* and *cognition*.
- The mathematical theory of *dynamical systems* and its applications for both computer music and cognitive sciences.
- A direct, personal *aesthetic engagement* in the development of interactive computer music environments and in concrete *artistic experimentation*, putting issues of interactivity at its centre.

Each of these aspects is represented by one of the preceding chapters.

I am aware that these elements, and therefore their chapters, might appear in part thematically disparate and in some way isolated from each other: the path leading from one chapter to the other might appear somewhat broken. To maintain this impression over smoothing it out was a conscious decision: such contrasts might help in bringing out an overall image of interplay. Also, connections pointing across the chapters are provided throughout the text, albeit not always fully developed. I will try here to give a concise résumé of the main themes and their connections, before providing a condensed account of the core claims of the thesis.

## *5.1 Résumé and Central Claims*

Chapter 2 (Interaction) introduces an understanding of computer music in terms of *generative computer music*. The means provided by a computational medium for formulating and performing processes are here the essential characteristics of this form of composition. As praxis and technological development in the context of computer music evolve, the inherent separation between sound source and interface centres questions of performance and *interaction*. A key assumption of this thesis is that a precondition for this understanding is a clearer insight into human *perception and cognition*, since interaction is central in how humans think and experience relations with other entities, machines or organisms populating our world. The *enactive approach* to cognition provides such an insight: this theory establishes tight links of dynamical exchange between the processes of cognition, action, and perception. From this perspective, interaction becomes the main mode through which *knowledge* about the world is constructed, via a *mutually influencing* exchange between perceiving agents and their environment. This is the understanding of interaction this work relies on. The theory of *agency*, which is part of the enactive approach, describes which qualities an entity should exhibit in order to enter an interactive relationship with a counterpart: *individuality, activity* and *adaptivity*. A further thought here is that, providing a computer music generative process with those qualities would allow a performer to engage in a mutually interactive relationship.

Chapter 3 introduces *dynamical systems* as a mathematical theory that is concerned with the temporal evolution of entities or ensembles of entities under the rules of their mutual interactions. The theory arises from the observation of complex physical phenomena and provides the tools for a qualitative analysis of their temporal behaviour in terms of geometrical structures in *phase space*. The language of dynamical systems can be applied to an extremely wide range of phenomena and is abstract enough to transcend the boundaries of the purely physical world. Dynamical systems offer a *process-based* way of thinking and an *ecological* perspective that looks at the connections and interactions between all of the elements involved in a system, rather than isolating them. Hence, this language finds its way into diverse fields in which those aspects are central, like cognitive sciences and the theory of enaction, as well as theories of perception and a specific praxis of computer music. While, in most approaches, this language is used at a metaphorical level, this thesis attempts to establish an approach which concretely employs the mathematics of dynamical systems in the realisation of interactive computer music environments. That is, a language that formulates and constructs the interdependencies between entities in a system in the form of differential equations. Through the flow of time, a composition of mutual interactions will emerge.

The practice-based artistic engagement with interaction in computer music takes a tangential path and is presented in Chapter 4 Case Studies. Interaction is, at first, understood as bodily involvement and therefore physical models (a subset of dynamical systems) are used for tapping into the bodily sensory-motor knowledge of performers. The hypothesis in this chapter is that, by modelling processes with physically inspired simulations, a behaviour could be composed that would exhibit a *temporality* resonating with that of our previously gained sensory-motor experiences and with the mechanisms of human temporal perception in general. The software framework *rattle* is the basis for the development of these environments. Through personal exploration and collaborations with performers and dancers in the *Embodied Generative Music* project, I could continuously explore and aesthetically experience interactive environments. This was an essential part of the process that made my implicit assumptions visible and allowed me to make more precise formulations. Through this experimentation a paradigm of *mutuality* appears in opposition to a *control* or instrument-based approach, in which the computer music system is understood as an extension of the performer's body. Interaction with a generative process requires reciprocity between performer and system, where the computer music system *appears as an agent* with which the performer interacts; it does not disappear as fully embodied instrument. The final case study explores the idea that, in using the language and formalism of dynamical systems, the computer music system can be provided with the *affordances* of an agent. This provision allows a truly interactive relationship between performer and generative process.

I would highlight the three following, mutually dependent points as central claims:

- *Interaction is the process of continuous mutual influence of two coupled agents.* Interaction is thus a *process* rather than a state. It is an ongoing continuous exchange of influences between two agents, who are coupled as a dynamical system. Each agent affects and is at the same time affected by the other in its temporal evolution. Interaction is also *situated*, as it is the result of a process that has to be performed, and is not a condition that can be set *a priori*. To compose interactions means therefore to formulate interdependencies between agents so that a process of interaction might emerge.

- *The language of dynamical systems allows us to formulate and analyse processes of change and interaction*. Dynamical systems theory is the mathematical language of *change* and *behaviour*. It offers a perspective of the world in which temporal processes emerge according to rules of change connecting the entities populating it. These rules of change are *couplings*. They are bi-directional and cyclic connections: change in one entity provokes change in another coupled entity, which in turn influences the former. As our perceptual apparatus is especially sensitive to change (i.e., to the perception of change in form of *derivatives*) dynamical systems provide means for formulating processes whose evolution, temporal structure, or behaviour resonates with our perceptual structure.
- *Agency is a perceptual quality modelled with dynamical systems*. Enaction theory describes an agent as an entity having *individuality*, *activity*, and *adaptivity*. An agent, therefore, is a recognisable source of activity in the environment with which it is coupled, and its adaptability means that it is capable of *self-regulating* this coupling. An agent defines itself through the qualities of coupling it exhibits; it appears and can be recognised as an agent only in the process of interaction. Agency is therefore a temporal perceptual phenomenon. The internal structure of an agent, as well as the form of coupling it exhibits, can be formulated in terms of dynamical systems.

*Agency* develops in the course of the dissertation to a central theme. It is this perceptual quality in a generative computer music system that allows for a mutual interaction with a performer. The generative sound process itself is at the core of the agent; it is its active and identifying character making its behaviour perceptible and sensible. This process is expressed in terms of a dynamical system. At the same time, this dynamical system's structure connects to its environment and exposes itself to external influences. Agency becomes the *affordance* the generative process offers for being interacted with: it is the quality that enables it to be touched, grasped, and interacted with. The haptic metaphor is here used following Alva Noë's description of vision as an active process of perception.[1] It is the surface presented by the process towards the composers/musicians/performers' influence, both *resisting* them as part of its agency and offering opportunities for being acted upon.

[1] Alva Noë. *Action in Perception*. The MIT Press, 2004

The composition of mutual relationships between the computer music agent's state space and external conditions is the basis for a composition of agency. This is a composition that will emerge in the process of interaction as a mutual shaping and forming of the space of possible actions.

## 5.2 Open Questions

What are the consequences of this approach in general and on computer music practice in particular? This question opens up many new directions that have not been addressed directly in this work, beyond suggestion. Following are a few of those directions, which should be object of future research.

What does this perspective imply for the composition of interactive computer music? It seems clear that this kind of understanding puts interaction itself at the very core of every piece employing it. Such an understanding requires the piece to *emerge* from the unfolding of the interaction. It is a situated process which develops in that moment, on that stage, with those performers, etc. Anything else would be in opposition with its enactive roots.

So, where is the piece? It cannot be a score in the sense of a traditional notation. The performers cannot follow a predetermined path: they are instead placed in certain conditions to act and react and co-determine the shape of the piece. The piece *is* the unfolding of interaction, it is the evolution of the joint dynamical systems of performer and computer music system. A different kind of notation seems to be required, one that could possibly reflect this situation.

What is the piece? If everything happens in the moment and is dependent on every aspect of the ecology of the performance, can we speak of a defined 'piece' at all? Again, probably not in a traditional sense: a piece consists of the formulation of a situation in which a specific set of processes must emerge. Composition here, though, means creating the conditions for the emergence of an interaction process.

A dynamical system perspective requires all involved entities to be on the same level: all contribute to the system's evolution and are indispensable to its path. That is, composer, musician, computer music system, audience, and venue all participate in the same system and have a crucial influence on the performance and, therefore, on the piece. Traditional hierarchical relationships are pulled into question from this ecological (and political) perspective. Who is the author?

Control paradigms are contrary to this perspective. Couplings between the agents in the systems are always mutual, and there is no unidirectional action of one agent over the other. Every entity senses and acts in accordance with its inner structure. The individuality of the agents is always respected: it is the motor of the mutuality. The computer music system has therefore to be carefully liberated from all tendencies of 'instrumentalisation' that would transform it into a tool in the hands of the performer. Further, the performer can be liberated from its role of 'controller' or 'interpreter' and be put in a more active and determining role in the composition. Control is a circular phenomenon: as cyberneticians have understood, control is dependent on the viewpoint from which the relationship between controlled and controller is seen.[2] From the perspective of a heating mechanism, for instance, the thermostat is the controlled entity, the controller, the thermostat, has to adapt to the control mechanism as much as the controlled. As cybernetician Heinz von Foester had observed, paradigms of control limit the possibilities of both entities equally and circularly, controller and controlled. These thoughts were responsible for the birth of second-order cybernetics, which sought a meta-perspective resolving this circularity of relationships between observed entities. Does the perspective developed in this work imply a kind of 'second-order composition'?

[2] Heinz Von Foerster. *Understanding understanding: Essays on cybernetics and cognition.* Springer Science & Business Media, 2007

Questions regarding the consequences of a *circular* concept of relations should also be addressed. Thinking in circular terms, implicitly assumes that the *temporal dimension* of things would be essential to their definition and their existence. An interesting direction to pursue, in this respect, would be the philosophy of Alfred N. Whitehead, who poses temporality as an essetial ontological cathegory: his work is currently experiencing a growing interest - a consequence of the work of Luciana Parisi on the status of the algorithm in generative art.[3]

[3] Alfred North Whitehead, David Ray Griffin, and Donald W Sherburne. *Process and reality: An essay in cosmology.* University Press Cambridge, 1929; and Luciana Parisi. *Contagious architecture: computation, aesthetics, and space.* MIT Press, Cambridge, MA, 2013

Most importantly, I think that this dissertation shows the need for an inquiry in the *specificity of computer music*. It seems to me that paradigms of traditional musical praxis have been more or less 'blindly' applied to the field of computer music. This is, of course, in itself reasonable: those paradigms have worked well until now, so why should they not work for computer music? Still, I do believe that computer music affords a qualitatively different approach, a unique way of thinking and praxis, due to its essentially generative character. These aspects are still not fully acknowledged. Concepts of instrument or control, composition as static, a solution to a 'problem', are - in my opinion - in opposition with the process-based way of thinking at the core of generative music. Instead, I call for a *radicalisation* of the concept and definition of computer music: a clear formulation of its core qualities with the intent of marking a departure from traditional musical praxis. This is the definition of a boundary, not as an insurmountable trench, but rather as a rhetorical tool for eliciting new modes of thinking.

## *5.3 Future Directions*

In this last section, I would like to report on the planned or already ongoing research projects or case studies that push this work's research questions further.

### *5.3.1 Phase Space Thinking: Experimental Explorations*

The case study Phase Space Thinking: an experiment reported in 4.3.1 was valuable for this dissertation. Conceiving the experiment and carrying it out, even in such a small scale, contributed greatly to the sharpening of questions and concepts. The necessity to conceive and realise an experimental setup that exposes the test persons to the correct questions drives a process of reduction of those questions. The simplest and most essential formulation of the problem is the key to a successful experiment and is simultaneously already of great scientific value. Furthermore, some aspects of a phenomenon can only be seen through a systematic exploration. Therefore, a continuation of these explorations would be an important factor in further research.

In particular, the next phase of experiments should focus on sharpening the description of the qualities of the coupling function between the agent's internal dynamical systems and external input. That is, how input energy is 'digested' by the system without a continuous accumulation, which might lead to instability and over suppression. This oversuppression might result in a suppression of the agent's own activity and identity. There is therefore a trade-off between these possibilities, which should be carefully evaluated and described as they have critical consequences.

A further theme which needs attention is the *adaptivity* of the agent. A clarification of this aspect and a formulation in terms of dynamical systems might lead towards a very useful approach in realising interactive environments in general. As Agostino Di Scipio has already noted, the key to true mutual interaction is the *self-observing* character of systems, which is tightly related to their adaptivity.[4]

[4] Agostino Di Scipio. 'Sound is the interface': from interactive to ecosystemic signal processing. *Organised Sound*, 8(3):269-277, 2003

### 5.3.2 *Agency and the* Algorithms That Matter *Project*

*Algorithms That Matter* is an artistic research project funded by the Austrian Science Fund (FWF, PEEK AR 403-GBL) and led by my colleague Hanns Holger Rutz and myself. The project takes on the questions of agency in the context of computer music praxis and is tightly related to this dissertation.

The basic assumption is that computational processes, or algorithms, possess an inherent agency as an irreducible and defining quality. This agency is perceivable both in the actual execution of the algorithm, the unfolding of the computational process, and the traces it produces in the very process of constructing algorithms, a process fuelled by the interaction between the programmers or composers and the computational medium.

The project ask questions about the medium-specificity of computation, in contrast to an approach employing the computer as merely a tool to solve well-defined problems through the execution of programmes. Algorithms, in fact, have traditionally been understood – in computer science, music and art – as a formalisation of thought; much like ideas, they were seen as immaterial and timeless. For instance, early algorithmic composition practices fall into this characterisation: in the words of Gottfried Michael Koenig, the computer is concerned with finding the solution to the problem 'given the rules, find the music'[5].

[5] Gottfried Michael Koenig. Kompositionsprozesse. In *Ästhetische Praxis*, volume 3 of *Texte zur Musik*, pages 191-210. PFAU Verlag, Saarbrücken, 1993

In contrast, the *Algorithms that Matter* project picks up impulses coming from current cultural studies and philosophy that suggest such praxis is characterised by two *entanglements*, first between the human and the apparatuses of creation (e.g., computers, software, algorithms, experimental arrangements, materials) and second between apparatuses and the objects produced (the arrangements and processes vs. the pieces of music or artistic knowledge). The concept of entanglement is borrowed from Karen Barad's work and means that two sides do not exist prior to their interactions, their separation happens only analytically.[6] These entanglements form the starting point of the project *Algorithms that Matter*.

[6] Karen Barad. *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning*. Duke University Press, Durham & London, 2007

Algorithms are taken as the crystallisation point of an inseparable human-machine agency in computer-based composition. Thus algorithms are studied as *performing entities* that emerge from specific artistic practices. And vice versa, the project is interested especially in how these practices are transformed by the agency of algorithms. While existing research often focuses on the refinement of algorithms, machine learning systems, etc., *Algorithms that Matter* looks at the process through which the algorithms and codes have come into existence.

The central question of the project is, therefore:

*How do algorithmic processes in experimental computer music structure artistic praxis and the understanding of composition and performance?*

In other words, the hypothesis is that these processes can unfold a specific *agency* that retroacts and changes the compositional praxis, becoming a new organising principle.

In addressing this question, the *Algorithms that Matter* project builds upon a new understanding of algorithms as entities bearing a material *performative* aspect that exceeds their design. An excess that, for instance, becomes material when algorithms have unintended consequences, such as crashing machines, etc.

Luciana Parisi's work serves as a basis for this perspective.[7] In Parisi's theory, an algorithmic object not only possess a finite material form, its particular implementation, and set of instructions. It is complemented by an abstract reality that makes it possible to produce and transform novel data. This surplus value is unwritten and non-implemented, *in-compressible* in the sense that it cannot be formulated. Through a *material engagement* oscillating between these two perspectives - the experimenting with and the observing of algorithms - the project aims at constructing an experimental system in which compositional practices serve as an epistemic tool in exploring the algorithms' performative essence.

[7] Luciana Parisi. *Contagious architecture: computation, aesthetics, and space.* MIT Press, Cambridge, MA, 2013

In the project, research questions are concretely addressed through principles of iterative experimentation. The approach to observing processes is inspired by Karen Barad's concept of 'diffractive reading', which describes 'an iterative (re)configuring of patterns of differentiating-entangling'.[8] That is, through a series of connected, but diverse, *re-configurations*, we attempt to observe the boundaries drawn by the agency of algorithms which may lie transversal to presumed boundaries such as a specific piece, performance, composer, format, etc.

[8] Karen Barad. *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning.* Duke University Press, Durham & London, 2007

The project is thus divided into four subsequent configurations. Each configuration brings together a group of artists and researchers who, over a period of two months, develop a series of algorithmically related sound pieces. The process is observed and transcribed into multiple forms of presentation and discourse and a continuous online exposition is complemented by gatherings and symposia. Each group consists of three persons: the two principal investigators, Hanns Holger Rutz and David Pirrò, and one additional artist/researcher. This will be a person invited to ensure a greater level of effectiveness and validity, reaching beyond the individual experience of the main investigators. In pursuing the investigation, one host environment in which algorithms are implemented and run will be used by all researchers in each configuration: this framework forms part of the laboratory apparatus. There are two software systems that will be alternatively used to this end: *Sound Processes*, maintained by Hanns Hogler Rutz and *rattle* (David Pirrò).

The research project is hosted at the Institute of Electronic Music and Acoustics (IEM) which is part of the University of Music and Performing Arts Graz (KUG) for a total running time of three years from 2017 to 2020.

# A

# A Catalogue of Works

This chapter contains a *catalogue* of selected artistic works and studies tied to this dissertation, either as tools or aesthetic experiences. T hese works served as a test and use case for the technical and conceptual framework. The process of developing and staging those works also enabled observations that otherwise would have not been possible. In this sense, I regard these works as *experimental*, in the original meaning of the term. They were trials, or tentative procedures; acts of testing a principle or a supposition; operations staged for the purpose of revealing something unknown.

These works have been developed, staged, performed, or embedded in some of the artistic research projects I've been part of. Also, my artistic practice plays an important role here, as it is intertwined with those research activities.

## A.1 Bodyscapes

Some of the paragraphs in the following section are based on parts of the paper 'On artistic research in the context of the project embodied generative music' by Gerhard Eckel and David Pirrò, which appeared in the Proceedings of the ICMC 2009

*Bodyscapes* is an interdisciplinary piece at the intersection of dance and computer music. It was realised in a collaborative artistic research process by Valentina Moar (dance, improvisation and choreography), Gerhard Eckel, and myself (composition, live electronics, interaction design and software development) over a total of seven days during two working periods in December 2008 and January 2009. The collective research and creation were carried out in the context if the *Embodied Generative Music* project at the IEM in the *aesthetic laboratory* (see The *Embodied Generative Music* Project), where the piece was also premiered on January 20th. A documentation video of the premiere is available online[1].

[1] https://vimeo.com/4949316: accessed 19/07/2017

After a prolonged period of development and experimentation carried out within the *EGM* project, exploring different motion-to-sound mappings with various dancers, we felt the need to condense our findings and observations into a short piece. This became *Bodyscapes*. The piece consists of different scenarios, (i.e. *bodyscapes*) each enacting a particular relationship between bodily movement and sound, bearing a recognisable characteristic and aesthetic identity. Each scenario revolves around a specific artistic idea or a metaphor, which serves as the basis for the development of the sound model and its mapping to the dancers' movement. The behaviour manifested by the resulting sounds would then, in turn, induce particular dynamics in their movements.

The following section describes the four bodyscapes appearing in the piece: they are named after the main metaphor driving their conception.

- *The Persona*: Starting our inquiry, we decided to concentrate our investigation on the dynamics of bodily movement. We searched for the most basic metaphors and sonic images connected to body dynamics. We were seeking ideas connecting body dynamics and sound, while making these relations clearly readable for the audience and 'wearable' by the performer.

  We identified the following characteristics this bodyscape should incorporate:

  1. directness of the link between movement and sound;
  2. simplicity of the relation;
  3. clearly readable causality of sound dynamics.

  As the body is always moving within air, the sound created by such movement inspired the sound model we used in this bodyscape. This was a low-pass-filtered noise that mapped the cutoff frequency of the filter to the speed of the movement. Our idea also implies that sound should only be produced if there is movement at all – as one of the clearest and most readable mappings of body dynamics to sound – so we extended the mapping such that the speed of movement also controls the volume.

  In this bodyscape, we took into account the spatial position of every joint in the body of the dancer, computing the speed of each, and using the fastest at any moment in the mapping. This is motivated by the assumption that the attention of the audience follows the fastest body part, meaning the sound dynamics follow the visual focus.

  The details of the mapping were determined empirically in experimental sessions preceding the production phase.

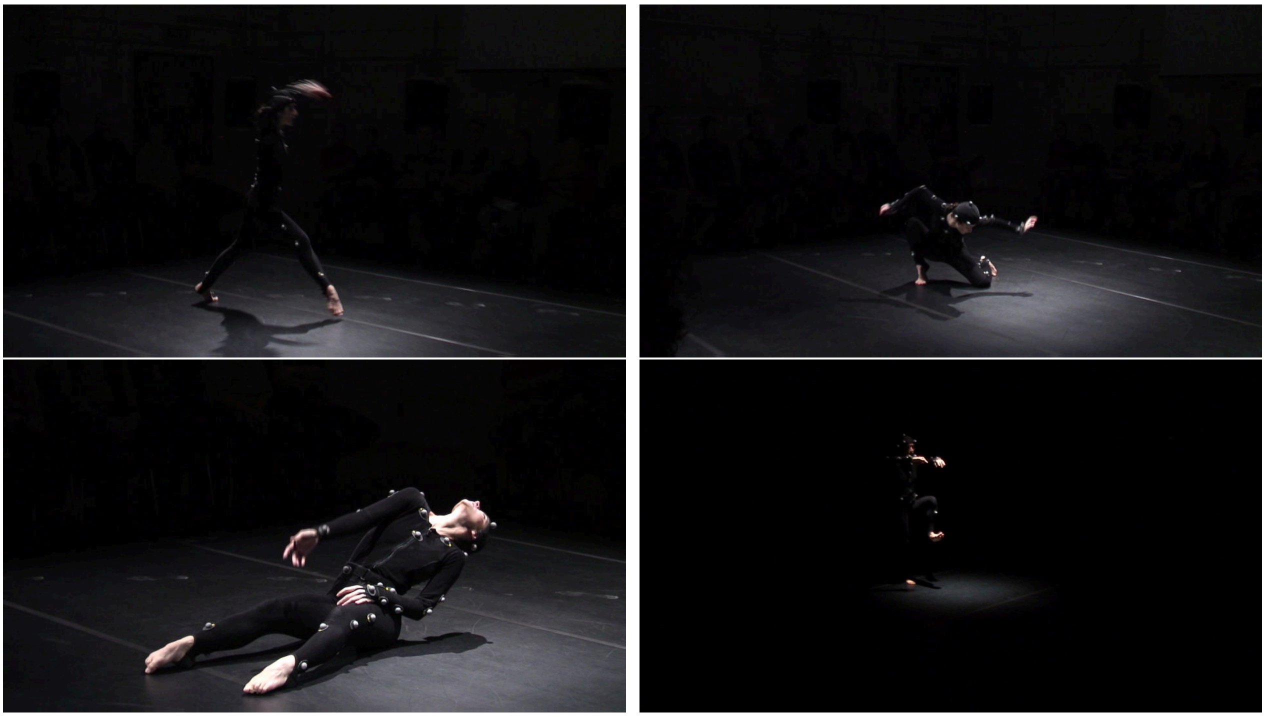

Figure A.1: Four moments of the *Bodyscapes* performance in the CUBE at the IEM. Referring to the explanations in the text, from top left the *persona*, the *partner*, the *frame* and the *object* on the bottom right.

The precise aim of these sessions was to establish a mapping faithfully portraying the *effort* involved in carrying out the movement. A dancer's judgment on the aptness of the sonic feedback was used as criterion in the process. The resulting mapping related the square root of the speed to the logarithm of the filter's cutoff frequency.

In order to keep the focus on the dancer, we had to avoid diverting away focus, which was achieved using a localised sound source (e.g., a single loudspeaker). We drove the hemispherical array of the 24 loudspeakers in our performance venue with 24 de-correlated filtered noise sources sharing the same cutoff frequency and amplitude mapping. This madesure that the dancing body remained the centre of attention in this bodyscape.

- *The Partner*: Moving away from the directness we achieved with the 'Persona' bodyscape, we tried to imagine a situation where sound and body were not so closely linked, and could therefore engage in a dialogue, albeit one where the sound was entirely caused by the performer. We imagined the following basic qualities of this bodyscape:

  1. indirectness of the relation between movement and sound;
  2. loss of complete control by the performer;

3. possibility for the dancer to establish a dialog with the sound produced.

The sound used in this bodyscape was produced through granulation of recorded sound material arranged in two sequences, in which voiced and unvoiced vocal sound fragments were separated by silent passages. Vocal sounds were chosen since they hint at the presence of a dialogue partner. The material was selected and arranged together with the dancer. The time axes of the two sequences were identified with the two axes of the horizontal plane of the tracking volume (the stage), mapping the time in each sound sequence to the position along one of the axes. The position of each hand of the performer on these axes functioned as an index in the sound sequence and determines which part is reproduced through periodic granulation (a kind of 'sound scrubbing'). The resulting two signals were then dynamically delayed, with the delay time mapped to the square root of the speed of the corresponding hand. The maximum delay of $2.5s$ was reached when the hand no longer moved and no delay occured at maximum speed. The variable delay produced a increase in pitch of the reproduced sound material whenever the hand accelerated, and a decreased when it slows down. The details of the mapping (maximum delay time, smoothing of the speed) were defined together with the dancer in a process in which she improvised and tested different settings. The two signals were discretely projected from four loudspeakers placed at the corners of the performing space, thus being clearly localisable and giving both the dancer and audience the opportunity to relate to the resulting voices in a 'theatrical' way.

- *The Frame*: The dancer's performance unfolds in a space that in itself is neutral, although it constitutes the frame in which bodily movement can take place. It is not the geometrical space we want to address in this bodyscape, but rather the environment – the fixed and unmodifiable or uncontrollable context through which the dancer is moving. We formulated some basic characteristics of this bodyscape:

1. the relation between movement and sound should be felt rather than readable;
2. unpredictability for the performer;
3. neutrality of the sound produced.

In this bodyscape we adopted a similar sound model as in the 'Partner' bodyscape, using granular resynthesis of a previously prepared sound file. The material used here is constituted from a selection of recorded impulsive and explosive sounds produced by heating a pot with a wet bottom on a boiling plate. We used eight sound generators corresponding to eight chosen joints in the body of the performer: left and right hands, elbows, feet and knees. The positions of each of these points along one of the two axes in the horizontal plane were used to find the position in the sound file reproduced through periodic granulation.

In this bodyscape, we wanted to design an environment that surrounds the dancer like a diffuse atmosphere in which she is moving. Therefore, we arranged the short impulsive sounds in a very dense distribution covering all of the tracking space, so that the dancer could not control the production of single sounds. The performer, moving in the space, generated a sort of cloud, as if she were lifting small dust particles into the air which left traces for us to hear. The speed of the movement determined the density of the projected sound events, so that the relation between movement and sound was more evident when moving slowly.

The sound in this bodyscape was very neutral, filling the whole space indifferently and homogeneously, projected into the performance space from an array of 48 small loudspeakers hung from the ceiling.

- *The Object*: Whereas the 'Partner' bodyscape allowed for a dialogue with an animated counterpart, the 'Object' bodyscape created a situation for the dancer to interact with an inanimate object positioned at a particular location in space (the centre of the stage). In this bodyscape, we wanted to explore the possibility of interacting with sound only through the change of position of the dancers body. The basic ideas for this bodyscape could be summarised as:

  1. clear spatial structure;
  2. complete control of sound production for the performer;
  3. directness and simplicity of position/sound relation.

  The sound model of this bodyscape distinguishes three zones: inside the object, outside of the object, and on the surface of the object. We found that especially the latter plays an important role in the clear identification of the zones. The perception of the surface is strongly linked to the sense of touch, and we therefore paid much attention to the surface's sound design.

  The joints taken into account in this bodyscape were, as previously, the left and right hands, elbows, feet, and knees. This time we used noise passed through a comb filter with long feedback, thus generating clearly pitched tones. Whenever one of these joints entered a cylindrical region of one meter radius placed in the center of the stage, an ADSR envelope was triggered, remaining open until the joint left this region. Out of this region, no sound was produced. In this way we represented an even clearer subdivision of the space through the presence or absence of sound. The envelope had a sharp attack in order to augment the feeling of touching/passing through a surface. When the dancer was in the region previously defined, the pitch of the generated tones varied slightly in a range from 3.2 to $4.5kHz$. according to the distance of each of the considered joints to the pelvis, which in this bodyscape represents the body centre.

Such a setup clearly defines an *object* external to the performer: something she cannot move or modify, but she can interact with, by moving through and being in.

The development of these movement-to-sound mappings and the calibration of the details of their parametrisations, in practice followed an empirical process which can be described as classical trial-and-error, shaped and guided by the aesthetic experiences of the participating dancers and composers. One of the main accomplishments of the *EGM* project was the development and the formulation of this *method* by which we tried to gain access to dance performers' implicit bodily knowledge about the aptness of movement/sound relationships. By composing virtual instruments in the framework of the EGM ÆLab, we could realise particular body/sound relationships, and rendered certain aspects of this knowledge explicit through sound. The collaborative composition of such a mapping is a tedious empirical process paved with failures and frustrations. This process met its objective only when the result *felt right* for the dancer (and the observer). Once the sound generation was felt to be embodied, and the dancers were able to fully engage with the sound, did they report a heightened awareness of the details of their movement. This opened new possibilities for the choreographic work, as structural aspects suddenly became audible.

After the first working period in the *Bodyscapes* production, dancer and choreographer Valentina Moar explained her experiences in an email in the following way: '*after a while it seems to me there is no more difference between the sounds and my skin.*' We take that as a clear indication for having reached a high degree of embodiment in the sound generation with the virtual instruments built

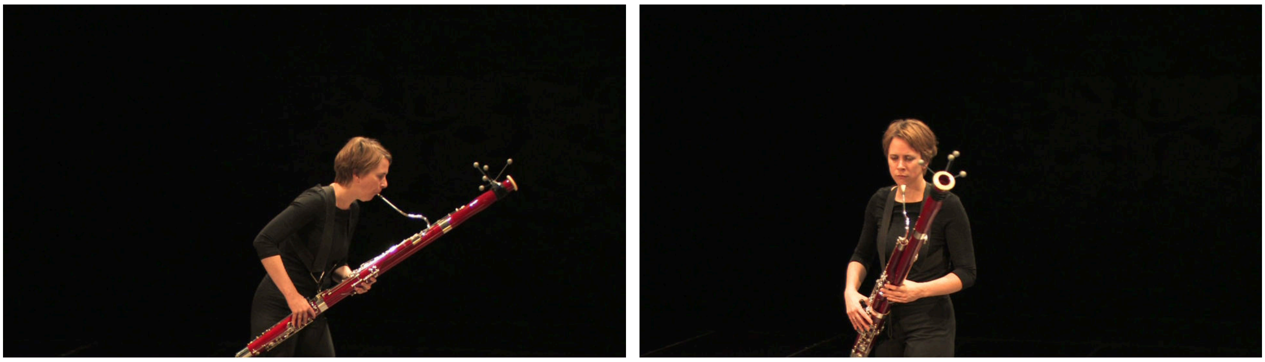

Figure A.2: Two moments of the performance of *cornerghostaxis#1* by Stephanie Hupperich.

in EGM. The symbiosis of movement and sound experienced by the dancers is the basis for the choreographic and compositional work. It is also a prerequisite for the special body/sound relationship to be performed and made accessible to an audience.

## A.2 *cornerghostaxis#1*

Some of the paragraphs in the following section are based on parts of the paper 'Physical modelling enabling enaction: an example', by David Pirrò and Gerhard Eckel, which appeared in the Proceedings of the International Conference on New Interfaces for Musical Expression 2011

*cornerghostaxis#1* is an artistic work, a composition for solo bassoon and live-electronics employing physical models in the design of the interaction between the performer and the electronics. The piece was premiered during IMPULS Academy 2009 in the context of the Motion-Enabled Live-Electronics workshop at the CUBE of IEM Graz, (bassoonist: Dana Jessen). It is the result of the collaborative effort of a team of three people: Stephanie Hupperich (bassoon), Gerriet K. Sharma (composition) and David Pirrò (physical modelling/interaction design).

The idea of the composition is to put the performer in a dialectical relationship with four electronic sound sources. These sources are dynamically spatialised on an array of loudspeakers. This is accomplished by a physical model establishing a gestural and bodily connection between the four channels of an electro-acoustic composition, the sounds the performer plays on her instrument, and her movements in space.

In the composition, the position and orientation of the tracked instrument is used as input for the physical model. The virtual space in which the physical simulation is taking place is a representation of the real space where the performance took place, including the positions of

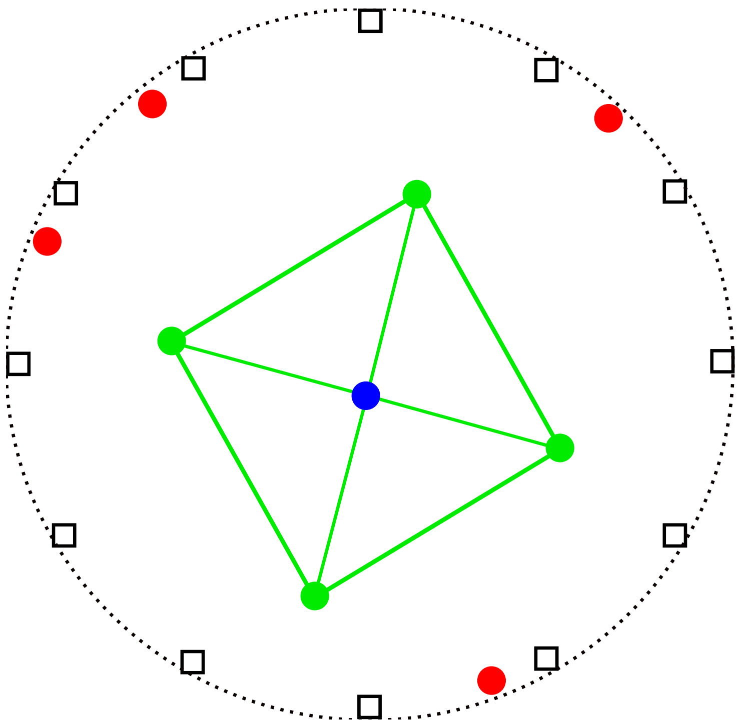

Figure A.3: Graphical depiction of the *cornerghostaxis#1* physical modelling environment. The red masses are free to move but are bound inside the disc, whose border is the dashed line. These masses interact with each other with an electric-type repulsive force, as if they were particles with the same electrical charge. They represent the spatialised position of the four channels of the electroacoustic composition on the loudspeaker array (the loudspeakers are the empty boxes at the boundary). The green square is centred on the blue mass whose position and orientation is controlled by the tracked bassoon. The four masses fixed at its corners also exert electric-like repulsive forces on the red masses.

the loudspeakers and the instrument. The physical objects that move and interact in this space are constrained on the surface of a virtual hemisphere, onto which also the loudspeakers are placed according to their actual positions.

The objects in the physical simulation have a very clear relationship: one can imagine them as electrically charged masses with the same charge. That is, the forces acting between the objects are repulsive[2]. The tracking data is used to control the position and orientation of a square with four 'charged' masses placed at its corners. The other masses are free to move on the hemisphere spanned by the loudspeakers: they are also 'charged' and repelled by the ones on the square as well as from one another (refer to figure A.3). The distances between these masses and the virtual loudspeakers are used to control a simplified DBAP algorithm (for a description of the algorithm see DBAP and ADBAP) that determines how the four channels of the tape composition by Gerriet K. Sharma are spatialised onto the arrangement of physical loudspeakers. Furthermore, the amplitude of the four sources is slightly modulated according to the movement speed of these masses and the distance to the performer. If the performer is close to one of them (i.e., she 'captured' one, see below) that source gets louder[3].

The piece has been conceived as a whole. None of its different aspects (e.g. interaction design, electroacoustic composition or the bassonist performance) overpowered the

[2] A short video of the model's simulation is available at http://pirro.mur.at/nime11/CGA-Model.mov (accessed 22/07/2017)

[3] A documentation video of the performance at Mumuth Graz is available at http://pirro.mur.at/nime11/CGA.mp4 (accessed 22/07/2017)

others and the development of each part advanced in parallel to the others. The physical model is not just an effect used to spatialise the tape composition: it is part of the piece, and part of the environment in which the composition unfolds.

In the next section, I summarise how the approach described above reshaped the working routine in the explorations and rehearsals of the piece, with respect to our aims. I therefore collect the most important observations made by the performers and by us. But we also attempt to condense our reflections based on our own aesthetic experiences gathered throughout. We understand the whole realisation of the piece, beginning with the design of the physical model, passing onto the preliminary explorations with the performer, to the rehearsals and the final performance, as part of an experiment aimed at putting this strategy into practice and observing what happens and how. An interpretation or evaluation of these observations is not explicitly given, but will be the object of future research.

The most important feedback came from the performers who played the piece. The musicians underlined that they felt having achieved a clear understanding of the dynamics of the sound spatialisation and how they could influence it. They could quickly establish an intimate control of the interface/model and they could rapidly learn how to play it.

This understanding also changed the communication between musician, composer and programmer. Relying on the physical metaphor, on which the programming and the whole realisation of the piece are based, the performers could more easily communicate with the composer and programmer. In this sense the physical modeling layer appears as a platform for the exchange and refinement of ideas which are shared among all the participants, regardless of their technical knowledge. For example asking ‘Could you make the masses heavier?’ is straightforward for the performer. At the same time it is easy for the programmer to understand and, knowing the model, to accomplish. This is one of the main reasons the performers were actively involved in the development of the piece.

Basically, in performing the piece the musician and the masses play a game of hide-and-seek. The sources try to escape the performer, always placing themselves at the points most distant to her. This dynamic became quickly clear to the performer in the first experimental session and her instinctive reaction was to try to find ways of stopping their continuous slippage, blocking one of them by pinning it down, or ‘capturing’ it. During the performance, the aim for the performer was to ‘catch’ one precise mass out of the four at a specific moment during the score. But the sound sources, which represent the mass positions in the model, seem to have their own will and try to hinder the musician to achieve her goal, in order to ‘win’ the game.

It is important to note here that understanding the rules of play means to understand the laws on which the physical model is based, which are coherently and continuously followed by the simulation and are inscribed in the sounds’ positions and movements. In our experience, this gaming quality greatly contributes in making the interaction more clear, interesting, and engaging. The reactions of the model are complex but retain a certain predictability. Thus the performer does not perceive erratic reactions in the model, which would destroy the illusion of a coherent environment - although the model and the sources are very difficult to control. It is tough to achieve exactly what the composer or the performer wants. The model resists the performer’s actions, at the same time offering great detail in interaction, as every little position or rotation change has audible consequences.

In our observations, the resistance of the model coupled with the refinement of control greatly enhances the felt embodiment. As a matter of fact, the musicians, after a short time of experimenting with the model, feeling challenged, asked for a more difficult setup, which was initially kept simple. That meant more resistance of the environment to their actions, but also more detail for their control. Resistance and detail of control create a continuous tension between performer and model that can be seen and felt clearly. This tension captures attention and engages the musician, as well as for the audience assisting at the performance.

The performer could thus fully engage in play with the environment and with the piece itself. The consistency of the interaction qualities and the resulting sonic feedback caused a ‘suspension of disbelief’ for the performer, who could truly and bodily trust the coherence of the model’s responses, of the connection between her movements, and the reactions of the sources. This link was so clear to one bassoonist that she started giving them a ‘body’, regarding them (in her own words) as ‘colleagues’, like she would do with other human players in an ensemble. Furthermore, she reported an enhanced sensibility, not only in the perception of the spatial location of sound, but also of her own movements, her position in space, and an increase of her proprioception.

I underline at this point that the model was neither visible to the audience nor to the performer, neither during the rehearsals nor the concerts. It was not clear to the viewer how the model worked or exactly which forces were acting in the simulation, as this was not explained before the concerts. It was not my aim to make this aspect evident. In our approach, the physical modeling layer is not intended to be clearly perceivable as such, but its purpose is to enhance the enactivity of the interaction.

Nonetheless, during the informal discussions that took place after the performances, it appeared that the relationship between movement and sound, between action and spatialisation, between the player's sounds and the electronic sounds, was clear also to the audience attending the performances. The player's efforts, inscribed in the qualities of her playing as well as in her body could be seen and conveyed to the spectator.

## A.3 *Tball*

---

Some of the paragraphs in the following section are based on parts of the paper 'Motion-Enabled Live Electronics' by Gerhard Eckel and David Pirrò, which appeared in the Proceedings of the Sound and Music Computing International Conference, SMC 2009

---

Similarly to *cornerghostaxis#1* (see previous section cornerghostaxis#1), the piece *Tball* for trumpet and live-electronics is a composed environment in which the musician and performer participate in a real-time physical simulation. The simulation establishes a relationship of interaction between the tracked performer, his movements and gestures in space (as well as the sounds he produces), and the spatialised movements of a sound source moving in the space. *Tball* was developed in collaboration with musician Paul Hübner and performed during the *Motion Enabled Live-Electronics* (*MELE*) Workshop that took place at IEM in the context of the Impuls 2009 Festival and Academy in Graz[4].

[4] A documentation video of the performance can be found here: http://pirro.mur.at/Tball/MELE_tball.mp4 (accessed 22/07/2017)

This composition's environment was inhabited by two agents: the trumpet player and performer and a moving invisible ('virtual') sound object, the *Tball*. As in *cornerghostaxis#1*, the virtual space in which the simulation takes place is a representation of the performance space, including the loudspeaker positions, the floor, and the tracked position of the trumpet's bell (see figure A.4). The spatial movement of the *Tball* has been modelled according to a simple spring-mass physical model. It is a point-like object attached by a spring to the a point at the centre of the stage, around 1.8$m$ above the floor. Once set into

Figure A.4: Two moments of the performance of *Tball* by Paul Hübner.

motion, it will oscillate with a given frequency according to its mass around its anchor point, possibly hitting the floor, where it will then bounce off. The position on stage and the orientation of the tracked trumpet are linked to the position of the second object in the simulation (refer to figure A.5)[5]. This object (marked as the blue filled ellipse in the previous figure A.5) can be imagined as a 'prolongation' of the trumpet from the bell's position in direction of its orientation: when the performer produces a sound with his instrument, this object exerts a strong attractive force on the *Tball*. In effect, whenever a sound is produced with the trumpet, this object 'grabs' the *Tball*. Additionally, this force is scaled by the relative angle of the *Tball* and the trumpet's direction: that means that this 'grabbing' force is greatest when the *Tball* faces the exact direction of the playing trumpet, whereas it is minimal (or even zero) when the *Tball* is to the side of the trumpet (e.g. at 90 degrees to the trumpet's direction). The grabbing force continues as long a sound is produced from the trumpet and is switched off when there is no sound.

The *Tball's* sound is spatialised on the loudspeaker array of the IEM CUBE according to its position relative to the loudspeakers in the model using a the *ADBAP* algorithm introduced in DBAP and ADBAP. That is, by listening to the sound from the loudspeakers and its dynamics, or the behaviour of its changes, the performer can hear where the *Tball* is. He then can engage in a game of catch, grabbing the ball and launching it away or against the floor.

During the whole performance the sounds played by the musician were continuously recorded into a ring buffer of fixed length: this recording is the basis of the sound the *Tball* produces. As new input are added into the buffer,

[5] A short video of the model's simulation can be found here: http://pirro.mur.at/Tball/Tball_Model.mov (accessed 22/07/2017). The black dots represent the loudspeaker positions, the moving green point the trumpet's bell, and the red point the *Tball*.

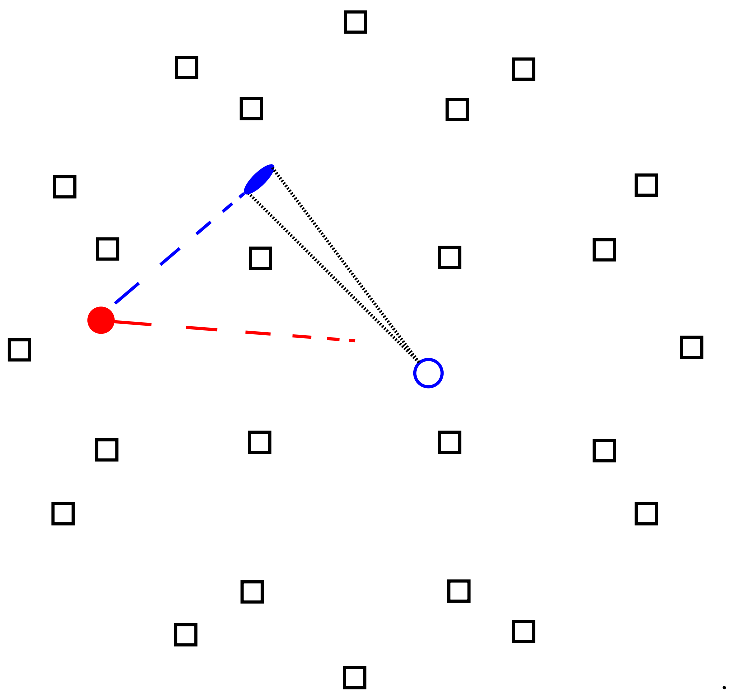

Figure A.5: Top view of the *Tball* environment. The red point represents the *Tball* object, which is attached with spring to the center of the stage (the red dashed line). The blue squashed disc represents the prolongation point of the trumpet bell, starting at the tracked bell's position (empty blue circle). This objects also exerts a force in the simulation (the dashed blue line) on the *Tball* which, whenever the trumpet produces a sound, 'grabbing' the *Tball*. The empty black boxes represent the positions of the 24 loudspeakers in the IEM CUBE (organised in 3 rings) on on which the *Tball* sound source is spatialised

this sound becomes more and more dense as the performance goes on. Accordingly, as the *Tball* accumulates all of these sounds, its becomes bigger and heavier, making it more and more difficult for the performer to grab and control it. The performance ends when eventually the ball 'explodes' like a balloon that has been inflated too much.

Composing the piece, input from and work with the instrumentalist were of great importance in implementing a well balanced environment allowing for a high degree of embodiment. Further, even if only small time slots where at our disposal for rehearsals, the performer could quickly construct a detailed image of the environment's dynamics, resulting in a high degree of control. Surprisingly, in spite of never having seen the graphical representation of the running simulation, the performer even asked for a more difficult parametrisation of the 'grabbing' force. Initially, the angular range at which the force could act effectively was left quite broad, as the performer could rely exclusively on his hearing in order to locate the moving sound source. It would appear though, that the dynamics and the behaviour inscribed in the sound source's movement in space was a clearer cue than expected, allowing him to predict quite precisely the object's location, extrapolating from its previous path - which, after a short period of time, even became 'too easy' in the musician's own words. We ended up by narrowing the angle range, thus requiring a better alignment of the trumpet with the object for the 'grabbing' force to be effective.

Eventually, the trumpet player and performer could engage in play with the *Tball*. By listening to the sound resulting from the interaction and watching the behaviour of the instrumentalist, the object could also appear in the audience's imagination.

## *A.4 Interstices*

Some of the paragraphs appearing in the following section are based on parts of the paper 'Exploring sound and spatialization design on speaker arrays using physical modelling' by Georgios Marentakis and David Pirrò, which appeared in the Proceedings of the 9th Sound and Music Computing Conference, SMC 2012

*Interstices* is a multi-channel sound installation, which I realised in cooperation with colleague Georgios Marentakis and was exhibited at *esc medien kunst labor* in January 2012.

The installation is an investigation into the spatial appearance of sound projected by a non-standard[6] speaker distributions. More precisely, it is an *artistic exploration* of how the composed temporal organisation on different time scales of synthesised sound, (i.e. the *behaviour* it exposes on multiple levels) affects its perceptual spatial appearance.

[6] *Non-standard* here refers to loudspeaker distributions which are not reducible to the spatial arrangements required by standard spatialisation techniques.

At origin, this work hypothesised that the shape of the composed temporal evolution of a sound is a strong cue – possibly even stronger than expected by psychoacousticians – that contributes the construction of a coherent perceptual sonic image. Most 'standard' algorithms that relate or transpose sound sources into space depart from a static conception of sound sources, ignoring movement or dynamical qualities. This is in some way understandable, as the effects of sound source movement on localisation have not yet been sufficiently studied and formulated, mostly because of the lack of sufficiently advanced speaker array systems that would allow a systematic investigation. However, it is known that, for instance, head movements clearly contribute to localisation. Even if the relevance of this cue cannot be clearly stated at the present time, it seems plausible to presuppose that dynamic qualities of sound including, but not limited to, its location could play and important role in perception.

The work naturally relates to the general theme of sound spatialisation in electronic music but, grounding on the perspective we introduced above, abandons the tendency to conceive sources as perfect points. Instead it attempts to generate sounding 'geometries' emerging through a consistent behaviour inscribed in their dynamic evolution.

*rattle* served here as the framework to formulate and compose such behaviour on different scales within time and space.

Here, the particle-based physical modelling and simulating version of *rattle* was used. I used a general formulation in which systems of particles connected together in a network, linked by variable forces acting between them. The behaviour of each of those objects and eventually of the whole compound is determined by the form of the interactions, i.e., the form of the forces acting between them. To define and possibly even change those interactions, would mean to re-compose and alter the dynamics on both the level of single objects and the whole network compound: one actor affects and is affected by all the others. Eventually, running the simulation, elements show a coherent behaviour according to the model's composed interactions. The primary aim was to experience how this behaviour affects sound and its spatial appearance. Depending on the relations governing its internal mechanics, these systems exhibit dynamics that lie within a continuum ranging from single organic entities, to extended subspaces, or to a collection of disjointed particles. Exploring this range of possibilities and making it subject to composition was a central aim of this exploration.

This approach is applied in parallel to different timescales of sound generation and spatialisation: three distinct systems or layers were used to compose the choreography of sound in space. We refer to these as *microscopic*, *mesoscopic* and *macroscopic* respectively. Each work on different time steps and rates, ranging from sampling rate (*micro*) in terms of sound synthesis to much slower transformations and bigger time steps when it comes to sound spatial distribution (*macro*).

- The *microscopic* layer is realised with a simulation that acts on the smallest time steps, i.e., at audio-rate. Displacements of the particles are directly audified: this modelling layer is therefore the sound generator, responsible for its morphology at small time steps - its microstructure.

  The model for this layer (see figure A.6) is constructed from a network of four mutually interacting masses, whose movements are confined within a sphere with an elastic boundary. The values of the speeds of these interacting masses at every audio-rate time step are translated into sample values in a four-channel audio stream. The weights of these particles' masses and the magnitude of the forces connecting them to each other have been chosen so that their movements exhibited changes with

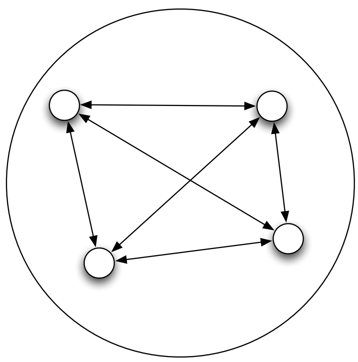

Figure A.6: Graphical depiction of the *microscopic* layer: the four empty dots represent masses interacting with each other and bound in a spherical region (the solid line boundary) frequencies within the audible range. The morphology of the sound output is thus a function of its weight, of the forces connecting the masses to each other, and of the attrition acting upon them. For instance, spring-like forces lead to simple, relatively static harmonic spectra. Gravitational-like forces, however, produce more complex and inharmonic sounds with unstable and changing time behaviour. With attractive forces, the sound exhibits clear pitches, while repulsive forces cause more impulsive, noisy sounds or bursts. During the preceding preparation phase, the different spatial perceptions were examined independently as a function of the different sound microstructures in dependence of the parameters of the model. Varying those, a 'behavioural' space can be identified, encompassing a range of distinguishable timbres ranging from harmonic to quasi-harmonic to transient. While the installation was in operation, these parameters were gradually changed to explore this space of possible behavioural states. Two of these models, oscillating between different states, ran in parallel in the installation, yielding substantial timbral variation, juxtaposing different sonic textures, and enhancing a differential perception of their specificity.

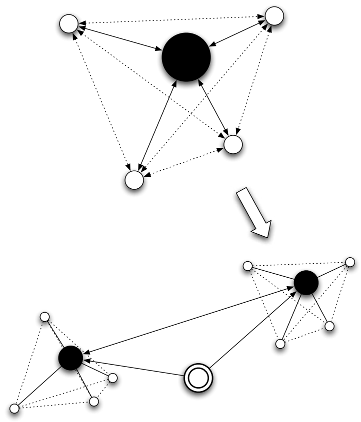

Figure A.7: Top: *mesoscopic* model. Five masses interact with each other through gravitational forces. The black object is bigger and heavier than the other four. Bottom: *macroscopic* model. The central fixed object acts on the black masses of the previous model attracting them; at the same time these objects repulse each other

- The *mesoscopic* modelling layer (and the *macroscopic*, see below) are tightly connected to how the sound produced by the *microscopic* layer is projected through the loudspeaker array.

  This layer (figure A.7 top), was implemented using five particles interacting through gravitational-like forces. The model was designed so that one of the objects acted as main attractor, keeping the other four orbiting around it. Its mass was substantially larger than the masses of the other four particles and the forces connecting the lighter objects to the attractor were substantially stronger than the ones connecting them to each other. Masses and forces were chosen so that the movements were significantly slower than in the previous *microscopic* model: the time needed for one of the smaller objects for a complete revolution was ca. $1-3$ seconds. This model was also updated dynamically, changing the magnitude of the attractive forces and yielding variable spatial configurations with objects moving in a loose or tighter way relative to each other, ending up very concentrated or more dispersed. The location of each object in this *mesoscopic* model defines where the sound (i.e., the displacement) each mass in the microscopic model produces, would appear on the loudspeaker array: that is, the movement of the single particle in the *microscopic* model are spatialised, not their sum or a mix. The qualities of their movement and their relative positions relate to how localized or extended the sounds projected by the loudspeaker array would be perceived.

- The *macroscopic* modelling layer connects to the movement (and rotation) of the *mesoscopic* model layer in itself and in relation to the whole space defined by the loudspeaker array.

  A similar approach was used here as in the previous mesoscopic layer. Again, a bigger mass, a fixed 'sun', was placed at the origin of the coordinate system. The two bigger attractors of the *mesoscopic* level revolve around this object, as they are attracted to it via gravitational forces. They are also mutually repelling each other through similar gravitational forces, so that the *mesoscopic* systems slowly revolve around this central sun, remaining mostly separated from each other and only occasionally mixing (Figure A.7 bottom).

In effect, there are two disjointed models working in parallel: the *microscopic* on the one hand and the *macroscopic* and *mesoscopic* on the other. The latter systems share the same simulation space, a rectangular box with reflecting walls which constrains the movement of their elements.

As has been already mentioned, this installation uses a loudspeaker array: the *The IEM modular speaker array system*. This is a 48-channel system that uses affordable Class-D amplifiers and small, easy-to-mount speakers, which provide the opportunity to rapidly prototype and quickly test diverse speaker array configurations. For *Interstices*, the 48-speaker array was divided into four speaker clusters, each containing twelve speakers (see A.9 and A.10).

For the spatialisation, an approach loosely leaning towards the Virtual Microphone Control (ViMiC) approach[7] mixed with a modified DBAP algorithm (see DBAP and ADBAP) was used. The simulation space was put in correspondence with the physical exhibition space by establishing a simple correspondence which maps specific positions in this space with the actual loudspeakers positions in the array. Each speaker was represented in the *macro* and *mesoscopic* system space using a single point. The sound of each of the *microscopic* model masses was rendered to the loudspeakers with an intensity that was determined based on the distance of its corresponding *mesoscopic* mass to each of those loudspeaker points. To avoid an excessive *blurring* between the single sources, the algorithm's parameters are chosen

[7] Jonas Braasch. A loudspeaker-based 3d sound projection using virtual microphone control (vimic). In *Audio Engineering Society Convention 118*. Audio Engineering Society, 2005; and Nils Peters, Tristan Matthews, Jonas Braasch, and Stephan McAdams. Spatial sound rendering in max/msp with vimic. In *Proceedings of the 2008 International Computer Music Conference*, 2008

such that each sound could appear on a maximum of three loudspeaker at the same time (see figure A.8).

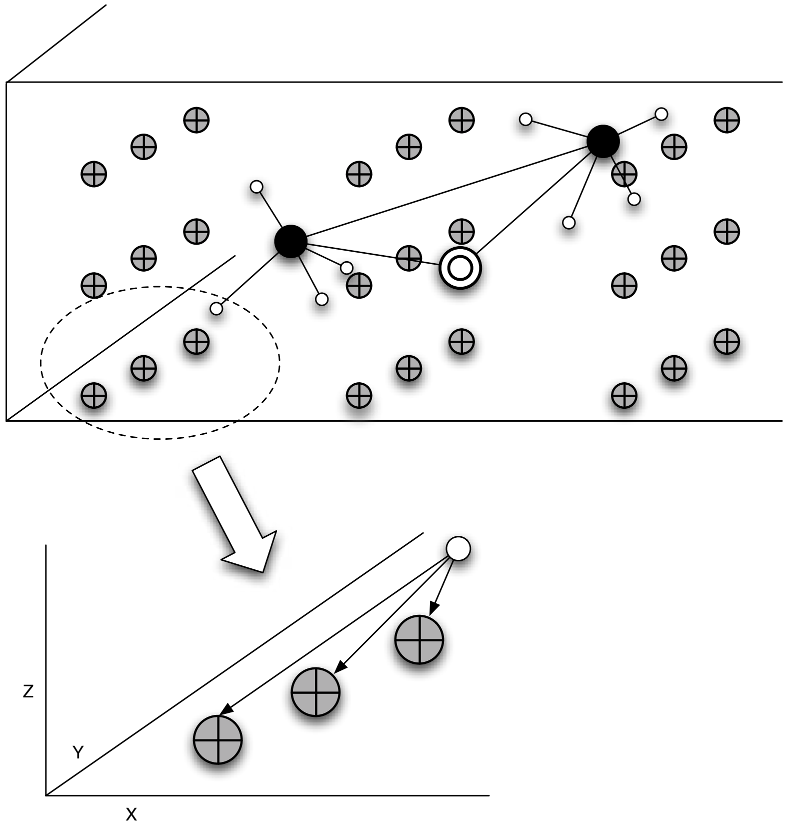


Figure A.8: Graphical representation of the spatialisation algorithm used in *interstices*

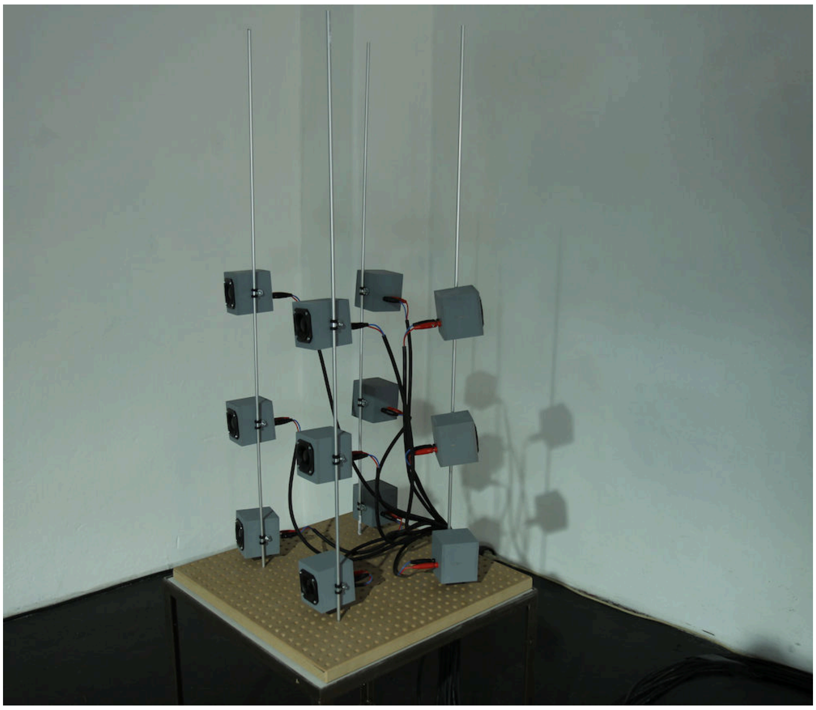

Figure A.9: One of the loudspeaker clusters used in the sound installation *interstices*. Foto: Martin Rumori

Figure A.10: Final distribution of the loudspeaker clusters in the esc medien kunst labor space. Foto: Martin Rumori

## A.5 *Zwischenräume*

Some of the paragraphs the following section are based on parts of the paper 'Zwischenräume - a case study in the evaluation of interactive sound installations' by Georgios Marentakis and David Pirrò, which appeared in the Proceedings of the International Computer Music Conference 2014

*Zwischenräume* is an interactive sound installation which can be understood as an evolution of the *Interstices* installation (Interstices) and a continuation of the collaboration with Georgios Marantakis. The installation, however, takes a more clear-cut and radical approach addressing some aspects at the core of a dynamical systems-inspired composition of interactive sound environments in an artistic setting.

This work was developed as part of the *Klangraüme* (2013 – 2015) project I was part of together with Georgios Marentakis (Project Leader) and colleague Marian Weger. The research project looked into how evaluation strategies common in HCI and the field of Sonic Interaction Design research could be applied to interactive artistic sound installations, and what kinds of consequences and effects the application of those methods have both on artistic praxis and on evaluation methods themselves.[8] In the context of this research project, the installation was also the object of an evaluation.[9]

The idea behind the installation *Zwischenräume* was that of an interactive environment which would be experienced as an organic entity continuously sensing the space, reacting to sonic events, and providing dynamic sonic spatial perspectives depending on the visitors' actions or their

[8] Unfortunately there is no space here for diving into questions and outcomes of this project. More information however can be found here: http://iem.kug.ac.at/klangraeume/klangraeume.html, accessed on the 17/07/20170

[9] Georgios Marentakis, David Pirrò, and Raphael Kapeller. Zwischenräume - a case study in the evaluation of interactive sound installations. In *Proceedings of the Joint 11th Sound and Music Computing Conference and the 40th International Computer Music Conference*, pages 277-284, Athens, 2014

mere presence. Interaction with the installation is made possible only through sound, which functions both as the input and output channel for the system.

Di Scipio's approach to interactive systems as ecosystemic systems[10] was central to the conception and development of the installation. In this sense, visitors and installation are regarded as equal agents that share the same space. Through their mutual interaction, an evolving dynamical system emerges. Interactivity is conceived as a continuous exchange between these actors; an exchange that affects the state of both of them through an adaptation process that eventually *resonates* in a state, a particular and recognisable *behaviour*. In the context of this work, my working definition of *behaviour* was according to following Arturo Rosenblueth's and Norbert Wiener's formulation:[11]

[10] Agostino Di Scipio. 'Sound is the interface': from interactive to ecosystemic signal processing. *Organised Sound*, 8(3):269-277, 2003

[11] Arturo Rosenblueth, Norbert Wiener, and Julian Bigelow. Bahvior, purpose and teleology. *Philosophy of Science*, 10(1):18 - 24, January 1943

> *By behaviour is meant any change of an entity with respect to its surroundings. This change may be largely an output of the object, the input being then minimal, irrelevant or remote; or else the change may be immediately traceable to a certain input. Accordingly, any modification of an object, detectable externally, may be denoted as behaviour.*

This definition is too broad to be useful, as the author himself suggests. Still, it forms a good basis for further characterisation. In the context of this installation especially, the main focus lay in the composition of behaviour-as-change, which unfolded both in the time domain and in the *spatial* domain, and in the *detectable* quality of this change. That is, the behaviour the installation would expose, would be clearly detectable, or better, sensed by the visitors as a trace of their actions in the space. In the words of Rosenblueth, a *purposeful* behaviour as directed to the attainment of particular condition and opposed to *purposeless* - i.e., random - behaviour.

On this basis, three specific scenarios or separable *eigenbehaviours*[12] were developed. These eigenbehaviours were then recomposed using a dynamical system that orchestrated their temporal and spatial evolution depending on the state of its environment and installation or on visitor interactions.

[12] Heinz Von Foerster. Objects: tokens for (eigen-)behaviors. *Understanding understanding: Essays on cybernetics and cognition*, pages 261-271, 2003

The aforementioned concepts led us to conceiving the installation as a feedback system: sound picked up by microphones would be projected back into the room with a specific delay. Feedback systems exhibit dynamically evolving behaviour, which served as the basis for the eigenbehaviours developed. In particular, varying the time delay created a rich palette of distinct sonic experiences, ranging from feedback tones, to the perception of spaciousness, and to echo effects.

The development revolved around the spatial, temporal, and energy relationships between the location of microphones picking up sound and the loudspeakers projecting it back. Necessary tools were a simple location-detection algorithm, implemented by determining which microphone received the maximum input at any time, and a ring buffer system that allowed an efficient control of the delay and gain of the output of each loudspeaker. All these tools were developed in *rattle*.

The installation was realised using the 48-loudspeaker system already described in Interstices and complemented with an array of 24 microphones. The initial staging decisions related to the placement of the loudspeakers and the microphones. With respect to the loudspeakers, we sought positioning that would structure the space less rigidly, in order to allow the visitor more freedom in choosing which paths to take through the space and installation. Loudspeakers were thus distributed quasi-randomly (see figure A.11), forming small clusters in the exhibition space. Various kinds of objects were used to mount the loudspeakers (music stands, microphone stands, tables, wooden blocks) to underline the playful character of the installation. As a consequence, this configuration provoked spatial heterogeneity and local behaviour as the different loudspeakers clusters projected sound slightly differently. Finally, to emphasise the fact that the installation reacts to the sonic activity in the room, some sound producing objects (a snare drum, some squeaky ducks and a trampoline with some bells attached under it) were distributed in the space.

In contrast to the loudspeakers, the microphones were hung from the ceiling in a very regular fashion. The exhibition area was covered with a regular lattice, on which microphones were placed with a fixed distance between them (see figure A.11). The desktop computer running the installation, the audio interfaces and AD/DA converters, amplifiers, and pre-amplifiers were stacked vertically within a box standing roughly at the centre of the room. Therefore, all signal cables formed a star shaped stem as they connected to the sound system. Although hiding the cabling was appealing to us, for practical as well as aesthetic reasons, we decided to use it as a visual element of the installation and to shape it consciously.

The scenarios for *Zwischenräume* were developed in a preliminary experimentation period aiming at developing a repertoire of clearly separable scenarios or eigenbehaviours yielding interesting and perceptually distinct sonic outcomes. These scenarios were fixed as parametrisations of

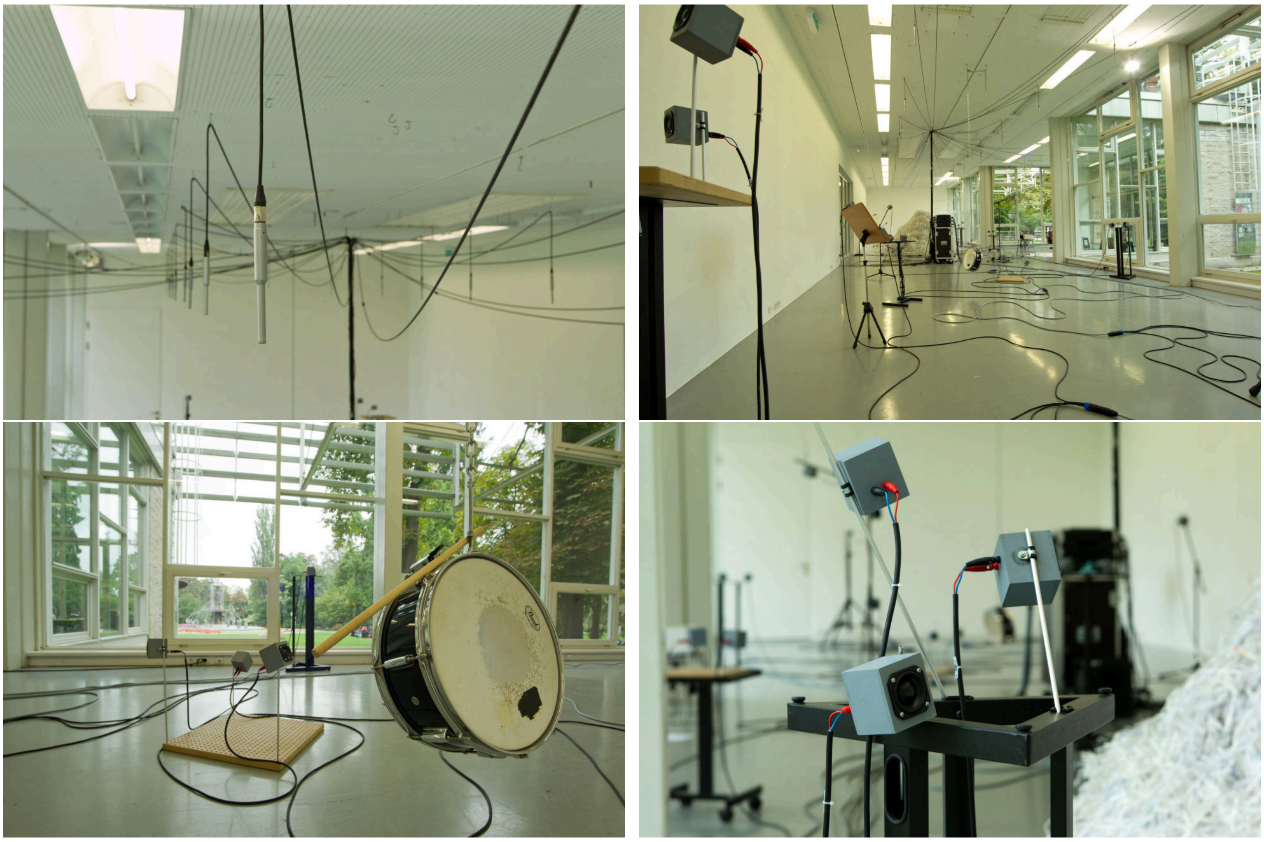

Figure A.11: Photos from the final installation setup in the Forum Stadtpark exhibition space.

the system, exposing a special behaviour with respect to its interactions with the visitor and the environment. Finally, a physical model was conceived that would re-compose these scenarios into a single installation. The model would expose either one or oscillate between two or more scenarios according to the visitors' activity in the room. The three scenarios and the physical model that were eventually chosen for the installation are presented in the next paragraphs.

- The *feedback* scenario directly exploits the feedback phenomenon (i.e., the so-called *Larsen* tones) that occurs when no or very little delay exists between input and output. In the most simple case, feedback manifests as tones, whose frequency depends on the main resonant frequencies of the room and its acoustic characteristics. However, when many loudspeakers with quasi-random orientations and locations are used as output and many microphones as input, more resonant frequencies can be excited simultaneously, producing complex spectra. For there to be spectral variability, however, the main resonant frequencies need to be suppressed, as they would otherwise dominate and lead the system into similar states. This can be achieved using a limiter and a peaking filter bank to control the overall amplitude of the feedback tones and the time needed for these tones to appear. Adapting the filter bank allows direct control over the inertia of the system, that is how much the system is sensible to changes in the environment and therefore the ease with which a transition between different feedback states occurs. Calibrating gain factors, filters, and limiters was challenging, as the feedback system strongly depended on the particular space and position of the loudspeaker and microphone. It was, however, possible to find configurations that produced complex feedback tones whose spectra depended on the listening location and the visitors' presence. In particular, the nearer the visitor was to a loudspeaker (or even holding a hand directly in front of a loudspeaker's membrane), the faster and more dramatically the system reacted. It has to be noted that this is the only scenario in which the installation was producing sound apparently on its own.

- In the *Hall and Echo* scenario the delay between input and output was increased, creating a spatially distributed reverb effect, increasing the perceived acoustic size of the room. With even longer delays, echoes would appear that would propagate onto the loudspeaker leading to an impression of spatial spreading of the sound. Moreover, the feedback of the echoes into the system through the microphones yielded further softer echoes that eventually spread uniformly over the whole array and slowly disappeared. By adjusting the spatial distribution of the loudspeakers, the effect of echoes from specific loudspeakers on specific microphones can be changed leading to the appearance of prominent spatial heterogeneity. It could happen that echoes would 'hang' between some loudspeakers and microphones, never disappearing or even growing louder. To avoid this, we introduced a calibration step by which loudspeaker gains were recomputed so that the maximum *RMS* value from each loudspeaker measured on the microphone array was equalised. This operation allowed more control and more stability in the overall system. Refining this scenario, gain and delay times were chosen so that the delayed signal was just on the threshold of being perceived as an echo. Therefore a reverb effect would emerge for continuous sounds (e.g., whistling), due to the temporal overlap of the sound with the echo onsets. In contrast, for impulsive sounds, the perception of echoes would be accentuated given the temporal distinction between sound offset and echo onset.

- The *Paths* scenario is derived from *Hall and Echo* and restructures it in order to provide the impression of auditory movement; echoes that slowly crawl through space, departing from the sound's originary location and moving along clearly perceivable, dynamic and changing paths through the loudspeaker array. To reinforce echo perception, delays here operate past the echo threshold. Sound captured by the microphone closest to the sound-producing action is recorded and played back with a delay from the nearest speakers. Using an adjustable delay, the same sound is then projected to the one or two loudspeakers closest to the previous with a slightly attenuated amplitude. As this process is repeated, a path of echoes is created, propagating from one loudspeaker to the other and eventually, after a period that depends on an attenuation factor, disappears. We intentionally avoided propagation paths in fixed directions in space (e.g., all paths moving towards one side of the room) and paths that would recirculate between a small number of loudspeakers. In order to minimise the effect of recapturing subsequent repetitions that would obscure the development of the paths in space, the signal from the one microphone receiving maximum energy was used as a source, while the input gain for all other microphones was strongly diminished. Only sound exceeding a specific threshold would be used as sources for this scenario. Particular to this scenario is that it explicitly advocates interaction between the visitor and the installation. In contrast to the previous scenarios, the effect of the acoustic environment is limited, making the behaviour of the installation's response completely dependent on the actions and sonic events produced by the audience.

- The goal of the *physical model* is to operate on the parameter space defined by the previous scenarios and synthesise their behaviour. In the model, both loudspeakers and microphones were defined as elements (masses) placed in locations that resembled their actual positions in the exhibition space, with microphones above the loudspeakers plane. All these objects were connected by forces. The masses representing the microphones exerted gravitational forces on the neighbouring loudspeakers masses. These, in turn, exerted and were affected by fixed spring-like forces exerted by their nearest neighbours. When a microphone received a signal above a certain threshold, it 'pulled' the loudspeakers it was connected with, with a force proportional to the signal's energy, thus exciting the whole system mesh of loudspeakers. This threshold was high enough to allow the whole system to relax when sound in the room was soft. The result was a mesh that, when excited, would behave much like a plate. An excitation would be transmitted to all loudspeaker masses in the model and the whole mesh would slowly wobble back to a resting state within a time frame determined by the inertia of the masses and the attrition we used. Using *rattle*, the simulation of this model was run in real-time at audio rate.

The displacement of the loudspeakers along the $z$ axis (towards the microphones) was used to control the delay with which captured sound would be reproduced by the connected loudspeakers: ranging from zero when at rest position to values appropriate for the hall and echo scenario. Velocity along the $z$ axis was used to control the gain of the loudspeakers: ranging from a lower threshold suitable to the feedback scenario (mass at rest) to a value appropriate for the echo scenario. Speed along the direction connecting one loudspeaker mass to its neighbours (paths scenario) was used to control the amplitude with which the signal was reproduced by the next mass. The displacement of the loudspeaker masses was mapped to the delay factor with which the repetitions were reproduced along the paths.

The effect of these choices was that when the masses were at rest (i.e. when there was little or no activity in the room), the installation would fall into the feedback scenario. As soon as a sound or a feedback tone appeared, the microphone masses would start to 'pull' at the loudspeakers. Feedback tones would slowly disappear as the excitation would spread over the whole mesh and the hall and echos would appear. Louder sounds and much activity in the room would result in greater displacements and speeds of the loudspeaker masses, causing the path scenario to eventually appear. Connecting the real-time physical model's state with the parameters of the scenarios, allowed us to recompose and merge the three single eigenbehaviours into one. Fine-tuning these mappings was a process that took a long time, but they eventually converged into the realisation of one system that would be perceived as coherent, exhibiting a global behaviour that exposed the three scenarios depending on the activity in the space.

In the evaluation phase, visitors were firstly observed during their stay in the installation space and interviewed afterwards. These interviews were then transcribed and analysed using the constant comparisons method and a combination of open and selective coding within the grounded theory framework.[13] All participants perceived not only that the installation was reactive with respect to their presence or actions, but also that it exhibited a sort of identifiable behaviour. Quite often statements classified under the behaviour category overlapped with statements under the interaction category. This is not surprising, as the installation behaviour was meant to manifest itself through interaction with the visitors.

Visitors interacted with the installation in a primarily playful and explorative way. The installation was thus interpreted as a rich medium where different perceptions could be created and observed, a pattern that was also quite evident in the video recordings.

[13] Barney Glaser. *Discovery of grounded theory: Strategies for qualitative research.* Routledge, 2017; and Juliet Corbin and Anselm Strauss. Grounded theory research: Procedures, canons and evaluative criteria. *Zeitschrift für Soziologie*, 19(6):418-427, 1990

# *B*

# rattle *Integration Algorithms*

When it comes to simulation, at the heart of every algorithm lie the developer's choices made about how to perform *numerical integration*. That is, how the problem of calculating the definite integral of a general function $f$ between two limits $a < b$

$$\int_a^b f(x)dx \tag{B.1}$$

is solved. This problem is central in computational physics, the research field dedicated to developing and analysing numerical operations in order to solve systems of differential equations. In this context, the above operation is called is called *quadrature* in order to distinguish this operation from the integration process of analytically solving (i.e. finding the mathematical equations) the above equation.[1]

[1] Steven E Koonin, Dawn C Meredith, and William H Press. Computational physics: Fortran version. *Physics Today*, 44:112, 1991

Many different methods have been devised to tackle the problem, each having different strength and weaknesses. The most important aspect of these numerical operations is that all these methods are *approximations* and thus affected by error. The understanding of how this error affects the found numerical solution identifies two distinguishing characteristics of each method: its *order*, its *stability*, and the *computational effort* it needs:

- the *order* of a method indicates how big the error of the method is with respect to the segment length $b - a$: higher order means the method produces smaller errors;
- the *stability* of the method refers to the property of the algorithm of magnifying (instability) or not (stability) the above error when it is repeatedly applied;
- the *computational effort* refers to the number of operations needed by the algorithm to calculate the result. Specially in a framework like *rattle* (see 4.1 The *rattle* System), where quadrature operations are performed in real-time at audio rate, this aspect plays an important role. In general, less *computational effort*

means smaller *order* and therefore worse approximations, so finding a good balance between these two aspects is central.

In general in this dissertationwe are dealing with *initial value problems* for ordinary differential equations: we are looking for $x(t)$ functions which are solutions to

$$\dot{x}(t) = f(x(t)) \tag{B.2}$$

given the value

$$x(t_0 = 0) = x_0 \tag{B.3}$$

for some initial time $t_0$: it is easy to see that this kind of problem reduces to a similar operation as in eq B.1 as we need to integrate $f(x)$ in order to find $x(t)$. This kind of problem occurs, for example, if we are given the momentum of a particle and its position as time $t_0$ and wish to know it position at some later time. If the function $f(x(t))$ is a continuous function, $x(t)$ is also continuous and can therefore be expressed in terms of its derivatives $\dot{x}, \ddot{x}, \ldots$ using a *Taylor series* to expand it in the neighbourhood of $t = 0$:

$$x(t) = x_0 + \dot{x}\, t + \ddot{x}\frac{t^2}{2!} + \dddot{x}\frac{t^3}{3!} + \ldots \tag{B.4}$$

where the derivatives are evaluated at $t = 0$.

Specifically, in our case, we are interested in the value of $x(t)$ at particular values of $t$ that are integer multiples of some fixed step $h$:

$$x_n = x(t = nh) \quad f_n = f(x_n) \quad n = 0, \pm 1. \pm 2, \ldots \tag{B.5}$$

e.g., $h$ could be, as in the case of *rattle*, the time interval between two audio samples, $1/44100 = 2.26e^{-5}$ seconds for a $44100Hz$ sampling rate. The above expansion in eq B.4 becomes at $nt = \pm 1$

$$x_{\pm 1} = x_0 \pm \dot{x}_0 h + \ddot{x}_0 \frac{h^2}{2} + O(h^3) \tag{B.6}$$

where $O(h^3)$ stands for the terms of order $h^3$ or higher. Assuming that $x$ and its derivatives are all approximately of the same order of magnitude, as is the case in many physical systems, these higher order terms will get smaller and smaller for higher powers if $h$ is chosen small enough.

From the previous equation, focusing only on the lower order terms we can easily derive the *forward and backward difference formulas*:

$$\dot{x}_0 \approx \frac{x_1 - x_0}{h} + O(h) \tag{B.7}$$

$$\dot{x}_0 \approx \frac{x_0 - x_{-1}}{h} + O(h) \tag{B.8}$$

Equation B.7 thus readily leads to *Euler's method*, also called the *forward Euler method*, the simplest of all quadrature algorithms, which for any $n$ and $n+1$ and using eq B.2 becomes

$$\frac{x_{n+1}-x_n}{h}+O(h)=f_n \tag{B.9}$$

and therefore

$$x_{n+1}=x_n+f_n h+O(h^2) \tag{B.10}$$

which gives us a method for calculating the next step of the trajectory $x(t)$ given $x_n$. On the one hand, this method has a very low *computational effort* and thus is very attractive for time-critical applications in audio synthesis. However, on the other hand it is neither very accurate (the step's error is just of second *order* i.e., $O(h^2)$) nor it is very *stable*.

The *numerical stability* of a method is established by applying the method to the numerical solution of a simple differential equation[2]:

[2] Abbas I Abdel Karim. Criterion for the stability of numerical integration methods for the solution of systems of differential equations. *J. Res. NBS*, 718, 1967

$$\dot{x}=\lambda x \tag{B.11}$$

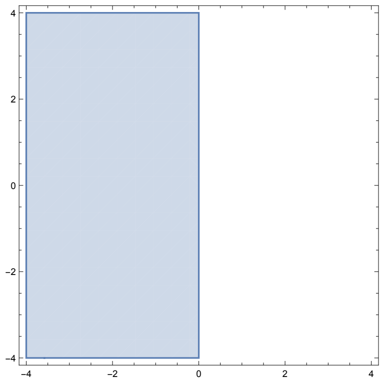

Figure B.1: In blue the region in the complex plane of analytical stability of the solution of equation B.11

which has the analytical solution

$$x=e^{\lambda t}x_0 \tag{B.12}$$

with $\lambda$ a complex number. $x_0=1$ usually. If $Re\{\lambda\}<0$ the solution is analytically stable as all possible trajectories remain bounded as time tends to infinity (see figure B.1).

Applying the forward Euler method to the numerical solution of equation B.11 thus using equation B.10 we get the following iterative rule:

$$\begin{aligned} x_1 &= x_0+\lambda h x_0=(1+\lambda h)x_0 \\ x_2 &= x_1+\lambda h x_1=(1+\lambda h)x_1=(1+\lambda h)^2 x_0 \\ x_3 &= (1+\lambda h)^3 x_0 \\ &\vdots \\ x_n &= (1+\lambda h)^n x_0 \end{aligned} \tag{B.13}$$

Equation B.13 describes a stable system for $n\to\infty$ if

$$|1+\lambda h|<1 \tag{B.14}$$

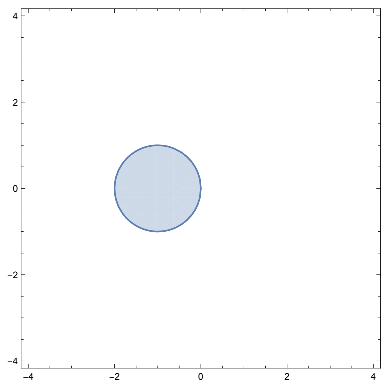

Figure B.2: Region of numerical stability of the forward Euler method in the complex plane $\lambda h$.

which is a disc of radius 1 in the complex plane of $\lambda h$ as depicted in figure B.2. As we can see, the region of numerical stability of the method is very small and does not cover the whole region of stability the analytical solution has. That is, the forward Euler method does a poor job in approximating the analytical solution.

This can be demonstrated with an example. We can, for example, consider the equation B.11 with $k=-2.3$ and $x_0=1$, which gives the stable analytical solution $x=e^{-2.3t}$. Applying the forward Euler method to this problem and choosing $h=0.7$ we would be in the stability region as

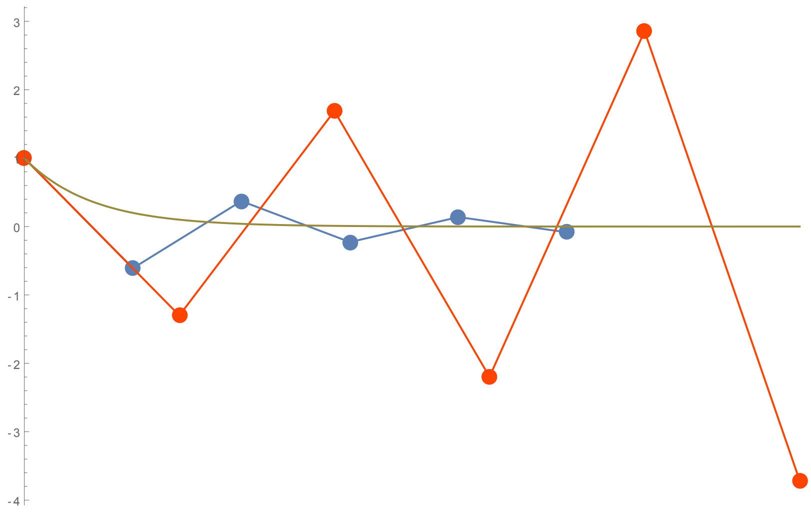


Figure B.3: Plot of the solution to the differential equation $\dot{x} = -2.3x$: in green the exact solution $x = e^{-2.3t}$, in blue the solution computed with the forward Euler method and $h = 0.7$, in orange the solution computed with the forward Euler method and $h = 1$ which is unstable

equation B.13 indicates. As depicted in figure B.3, after a short initial oscillating region, the method would be stable. Choosing instead $h = 1$ would mean being outside the stability region and would therefore be unstable. The method would produce oscillating solutions growing in amplitude, and is thus extremely sensitive to the right choice of the step $h$ - which should be sufficiently small. This instability is particularly evident in oscillatory solutions of the B.11 equations, i.e., when $Im\{k\} \neq 0$. These are of particular interest to us: in this case, even with a very small step size with respect to the frequency of the system, the method would always be unstable, and the energy of the system would grow exponentially.

Taking equation B.8 instead would lead to to a different iterative method, known as *backward* or *implicit Euler*:

$$\frac{x_{n+1} - x_n}{h} + O(h) = f_{n+1} \tag{B.15}$$

and therefore

$$x_{n+1} = x_n + f_{n+1}h + O(h^2) \tag{B.16}$$

Even if this method seems very similar to the previous, it is exhibits substantial differences. The *numerical stability* analysis of this method, applying the previous process, would lead to:

$$\begin{aligned}
x_1 &= x_0 + \lambda h x_1 \Rightarrow x_1 = \frac{1}{(1+\lambda h)} x_0 \\
x_2 &= x_1 + \lambda h x_2 \Rightarrow x_2 = \frac{1}{(1+\lambda h)} x_1 = \frac{1}{(1+\lambda h)^2} x_0 \\
&\vdots \\
x_n &= \frac{1}{(1+\lambda h)^n} x_0
\end{aligned} \tag{B.17}$$

which would be stable if

$$\frac{1}{|1+\lambda h|} < 1 \tag{B.18}$$

As shown in figure B.4, the shape of the stability region of this method is very different to that in the former method. As can be seen, this method is good for approximating solutions for the stable region of the analytical solution. It produces a stable solution even where the analytical solution gives unstable (i.e. growing) solutions in the complex half plane $Re\{\lambda\} > 0$

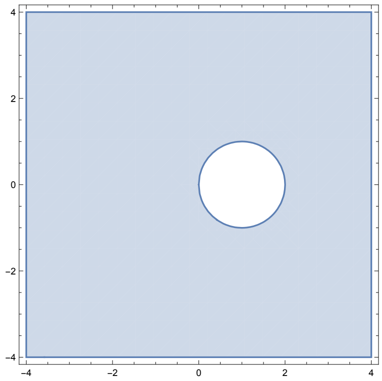

Figure B.4: Region of numerical stability of the backward Euler method in the complex plane $\lambda h$.

Furthermore, the method, as with all other *implicit methods*, presents an ulterior difficulty. In fact, reformulating equation B.16 taking into account that $f_n = f(x_n)$, we see

$$x_{n+1} = x_n + f(x_{n+1}) \tag{B.19}$$

that the term $x_{n+1}$, which we want to find, is on both sides of the equation: this is the fundamental characteristic of all implicit methods. As a consequence, one needs to solve an *algebraic equation* in the unknown $x_{n+1}$: this problem can be reformulated as to find the roots of the function $g(x_{n+1})$:

$$g(x_{n+1}) = x_{n+1} - x_n - f(x_{n+1}) = 0 \tag{B.20}$$

that is the points $x_{n+1}$ for which this function is zero. This can in general be a very difficult problem to solve numerically as $f$ could be any non-linear function. Usually this kind of problem is solved with iterative methods, such as the *Newton-Raphson method* which drastically increase the computational effort. However, neither the Euler methods nor an iterative method are suited for implementation in a software framework that needs to perform fast and stable (i.e., at audio rate) numerical integration.

Of course, the forward and backward Euler are the most simple and error prone numerical methods; still those methods show the basis on which all integration computational methods are constructed. The methods that tend to produce better results are constructed using two principal paths.

*Linear multistep methods* (also known as the *Adam-Bashford methods*) depart from a slightly different formulation, as in equation B.4 to compute $x(t)$. From equation B.2

$$x(t_1) = x(t_0) + \int_{t_0}^{t_1} f(x(t))dt \tag{B.21}$$

that is, in discrete time steps:

$$x_{n+1} = x_n + \int_{n}^{n+1} f(t)dt \tag{B.22}$$

the derivation of these methods follows the idea to approximate better the value of the integral of the function $f$ by taking into account its value at previous time steps and thus producing linear, quadratic, cubic, etc. polynomial approximations of $f$. This leads to a whole family of higher-order explicit or implicit methods. As an example, the explicit methods following the linear and cubic approximation of $f$ would be respectively:

$$\begin{aligned} x_{n+1} &= x_n + h\left(\frac{3}{2}f_n - \frac{1}{2}f_{n-1}\right) \\ x_{n+1} &= x_n + h\left(\frac{23}{12}f_n - \frac{4}{3}f_{n-1} + \frac{5}{12}f_{n-2}\right) \end{aligned} \tag{B.23}$$

and the respective implicit methods would be:

$$\begin{aligned} x_{n+1} &= x_n + h\frac{1}{2}\left(f_n + f_{n-1}\right) \\ x_{n+1} &= x_n + h\left(\frac{5}{12}f_n + \frac{2}{3}f_{n-1} - \frac{1}{12}f_{n-2}\right) \end{aligned} \tag{B.24}$$

The *Runge-Kutta method* family comprises widely used numerical integration algorithms that use higher order expansions of the the Taylor series in equation B.4 to better approximate the integral of the function $f$. The so derived second order method algorithm would be:

$$\begin{aligned} k &= hf(x_n) \\ x_{n+1} &= x_n + hf(x_n + \frac{1}{2}k) \end{aligned} \tag{B.25}$$

and the widely used fourth order method:

$$\begin{aligned} k_1 &= hf(x_n) \\ k_2 &= hf(x_n + \frac{1}{2}k_1) \\ k_3 &= hf(x_n + \frac{1}{2}k_2) \\ k_4 &= hf(x_n + k_3) \\ x_{n+1} &= x_n + \frac{1}{6}\left(k_1 + 2k_2 + 2k_3 + k_4\right) \end{aligned} \tag{B.26}$$

These methods can be very accurate and exhibit better stability properties, but involve the computation of the value of the function $f$ multiple times for each time step.

In *rattle* a different kind of integration scheme is used: a *symplectic scheme*. This particular method can be used in the numerical integration of a special class of problems of the B.1 equation, called *Hamiltonian systems*. These are systems of *coupled* differential equations and grounded on Newton's second law:

$$m\dot{v} = F(x) = -\frac{dU(x)}{dx} \tag{B.27}$$

$$\dot{x} = \frac{dv}{dt} \tag{B.28}$$

which describes how a mass under the influence of the force $F$, or a potential field $U$ accelerates. To understand these methods, a small step backwards into theory is necessary.

*Hamiltonian systems* are dynamical systems which can be described with the Hamiltonian function, embodying Newton's second law of mechanics.[3] These are of utmost interest in physics: they are used to describe most systems found in nature from planetary system to the motion of an electron in an electromagnetic field. These equations depend on the characteristics of the Hamiltonian function $H$, related to position, velocity of the involved elements (masses), and time. The special interest in physics for this function derives from the fact that the Hamiltonian is for these systems the sum of the kinetic and potential energies $T$ and $U$:[4]

$$H(q,p,t) = T(p) + U(q) \tag{B.29}$$

For instance, the Hamiltonian of the simple harmonic oscillator would be:

$$H = \frac{p^2}{2m} + \frac{1}{2}kx^2 \tag{B.30}$$

Thus, usually the Hamiltonian is the energy of the formulated system and for *closed systems*, given the conservation of energy, it is constant and time independent:

$$\frac{\partial H}{\partial t} = 0 \tag{B.31}$$

A principal characteristic of this function is that it describes the *evolution* of the state of the dynamical system, i.e., it describes how the coordinates $q$ and $p$ evolve in time via the so called *Hamilton equations*, a system of differential equations of the general form of equation B.2:

$$\begin{aligned} \dot{p} &= -\frac{\partial H}{\partial q} \\ \dot{q} &= \frac{\partial H}{\partial p} \end{aligned} \tag{B.32}$$

Considering the space spanned by the coordinates $(q,p)$, the *phase-space*, the integration of the former equations results in a so-called *flow* in this space. To any (continuous and differentiable) Hamiltonian corresponds a *flow* $\phi_t$, which describes the time evolution of the system. Given any initial coordinate in the phase-space $(q_0, p_0)$, this returns to the point $(q,p)$ to which the system would evolve at any time $t$:

$$\phi_t : (q_0, p_0) \rightarrow (q(t), p(t)) \tag{B.33}$$

An important characteristic of this function is that, for Hamiltonian systems, it is a so-called *symplectic map* meaning that it is area preserving in the phase

[3] Herbert Goldstein, Charles Poole, and John Safko. *Classical mechanics*. Addison Wesley, 2002

[4] For the sake clarity and conciseness, I'm following a simplified mathematical treatment of this section trying to bring across the most important concepts qualitatively. Further I will use, as in most texts, the *generalised coordinates* notation for position and momenta, $q$ and $p$ respectively. Therefore in the next equation, I assume separable Hamiltonians (the potential $U$ is not dependent of the momentum $q$).

space. In other words, given a section of the phase space, transforming this section with a symplectic map would translate it to different section in the phase-space, which could be different in form, but would have the same area (see figure B.5).[5]

[5] This descends from a 1899 Theorem by Poincarè, published in *Les Methodes Nouvelles de la Mecanique Celeste*.

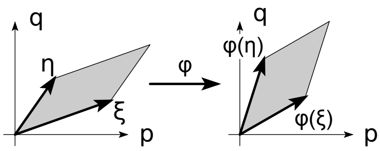


Figure B.5: Simplecticity (area preservation) of the mapping $\phi_t$

This quality of the Hamiltonian systems, which are the systems we are mostly dealing with in *rattle*, is essentially characterising this set of problems and is ultimately related to fundamental principles of physics as *Liouvilles' theorem* and the principle of energy conservation.

It seems therefore obvious to ask that the *symplecticity* property of the exact solutions of Hamiltonian systems should also be embodied and respected by the numerical integration methods.[6] That is, any numerical method $\Phi_h$ approximating the flow of the exact solution such that

[6] Ronald D Ruth. A canonical integration technique. *IEEE Trans. Nucl. Sci.*, 30 (CERN-LEP-TH-83-14):2669-2671, 1983

$$(q_{n+1}, p_{n+1}) = \Phi_h(q_n, p_n) \tag{B.34}$$

given any point $(q_n, p_n)$, should be a symplectic transformation. None of the methods described above, whether explicit and implicit, are symplectic independently from the order they could reach. Nor, could any of the above methods be guaranteed to respect fundamental characteristics of dynamical systems as the conservation of energy. This can be depicted on the basis of the Euler methods we introduced above. Recalling the explicit Euler method, we know it would tend to expand the energy of the system (solutions grow in energy) and the section of the phase space would grow in area. The implicit Euler method, meanwhile, would tend to reduce it (solutions would tend towards stability even if analytically they would not). This behaviour is depicted graphically in figure B.6, considering the example phase space flow generated by the Hamiltonian system of the simple pendulum.[7]

[7] Ernst Hairer, Christian Lubich, and Gerhard Wanner. *Geometric numerical integration: structure-preserving algorithms for ordinary differential equations*, volume 31. Springer Science & Business Media, 2006

The symplecticity request leads to the formulation of a new family of *symplectic methods* which guarantee conservation of energy and area when applied to the integration of a dynamical system. The first of these methods is the *symplectic Euler* method, which can be equivalently expressed in two ways:[8]

[8] Rene de Vogelaere. Methods of integration which preserve the contact transformation property of the hamiltonian equations. *Department of Mathematics, University of Notre Dame, Report*, 4:30, 1956

$$\begin{aligned} p_{n+1} &= p_n - h\,\frac{\partial H(q_n, p_{n+1})}{\partial q} \\ q_{n+1} &= q_n + h\,\frac{\partial H(q_n, p_{n+1})}{\partial p} \end{aligned} \tag{B.35}$$

or

$$\begin{aligned} p_{n+1} &= p_n - h\,\frac{\partial H(q_{n+1}, p_n)}{\partial q} \\ q_{n+1} &= q_n + h\,\frac{\partial H(q_{n+1}, p_n)}{\partial p} \end{aligned} \tag{B.36}$$

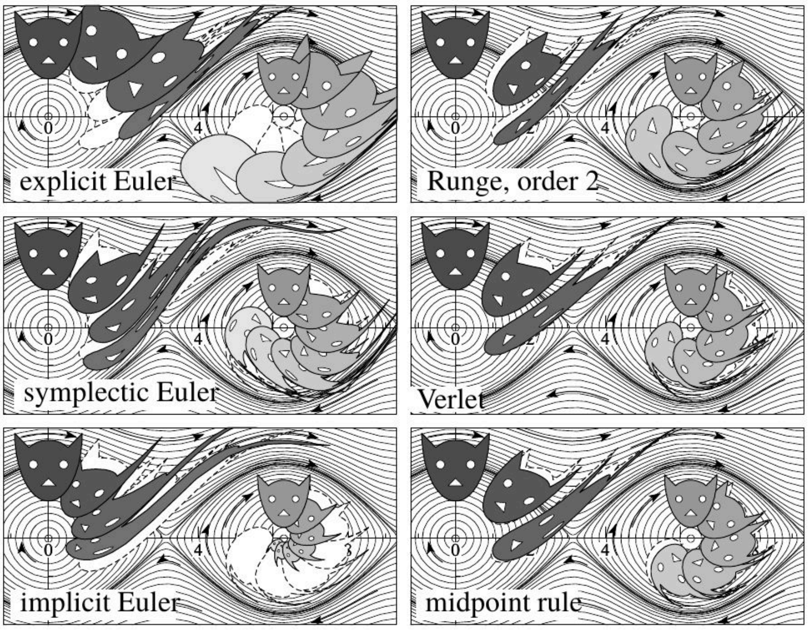


Figure B.6: Area preservation behaviour of various numerical integration methods on the basis of the a phase space of the simple pendulum. Same initial areas (and values) are chosen

which, recalling that:

$$-\frac{\partial H(q,p)}{\partial q} = -\frac{\partial U(q)}{\partial q} = f(q) \tag{B.37}$$

where $f(q)$ is the force acting on the mass and

$$\frac{\partial H(q,p)}{\partial p} = \frac{p}{m} = v \tag{B.38}$$

the former reduce to

$$\begin{aligned} p_{n+1} &= p_n + hf(q_n) \\ q_{n+1} &= q_n + h\frac{p_{n+1}}{m} \end{aligned} \tag{B.39}$$

and the equivalent:

$$\begin{aligned} q_{n+1} &= q_n + h\frac{p_n}{m} \\ p_{n+1} &= p_n + hf(q_{n+1}) \end{aligned} \tag{B.40}$$

That is, each of these methods uses an implicit method for the evolution of one state variable and the explicit method for the other. The performance of these two methods, even if only of first order, is already much more stable as is depicted in figure B.7.

One of the most far-reaching consequences of the *symplecticity* of Hamiltonian systems is that a geometrical way of thinking about the numerical integration of such systems' evolution is made possible. In fact, these integration methods are usually also referred to as *geometric integrators*.

This geometric perspective is the basis of further development of those methods, given the following observations:

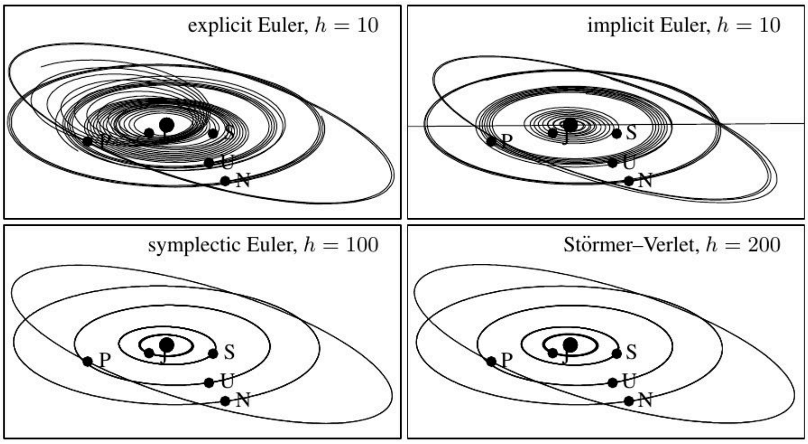


Figure B.7: Solution to the outer solar system as computed with the explicit, implicit, and symplectic Euler and Strömer-Verlet methods. The graphic is taken from the book by E. Hairer: *Geometric numerical integration: structure-preserving algorithms for ordinary differential equations*

- *Composition*: Numerical methods can be composed in the same way functions can be composed. That is if $\Phi_h$ and $\Psi_h$ are two different numerical methods of order $r$ and $s$ respectively for the same problem, their composition $\Phi_{\frac{h}{2}} \circ \Psi_{\frac{h}{2}}$ is also a method $X_h$ for the same problem with order $r + s$.

- *Symmetry*: The exact flow of a dynamical system $\phi_t$ usually satisfies the relation $\phi_t^1 = \phi_t$: This property is in general not satisfied by the flow $\Phi_h$ of a numerical method. The *adjoint method* $\Phi_h^*$ is defined as equal to the inverse method with reversed time.

$$\Phi_h^* = \Phi_{-h}^{-1} \tag{B.41}$$

  and a method is called *symmetric* if is is equal to its adjoint $\Phi_h^* = \Phi_h$. Further, the adjoint of an adjoint method is the original method $(\Phi_h^*)^* = \Phi_h$ and the adjoint of a composition, if the composition of the single adjoint methods is reversed in order $(\Phi_h \circ \Psi_h)^* = \Psi_h^* \circ \Phi_h^*$. *Symmetry* is an important quality of flows which is related to the reversibility of dynamical systems, a fundamental characteristic of all conservative systems and is therefore a quality that a numerical method should provide.

- *Splitting*: A flow in phase space, i.e., a vector field, can be split into the sum of two (or more) simple flows along one of the dimensions of the phase space. The total flow is then the composition of the two flows (see figure B.8). For instance, the first symplectic Euler method $\Phi_h$ formulated in equation B.40 could be split into two flows

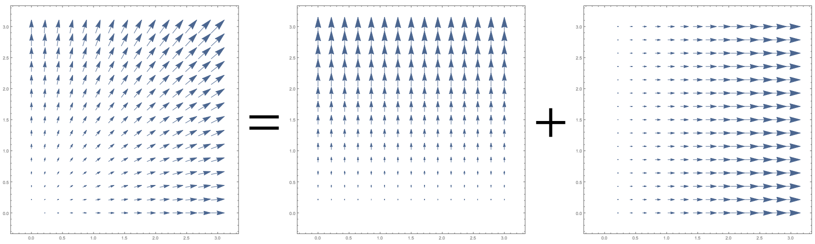

Figure B.8: The splitting of a flow in two-dimensional phase space is expressed as the sum of two more simple flows

$\phi_h^{[1]}$ and $\phi_h^{[2]}$ respectively along the $p$ and $q$ dimensions:

$$\begin{aligned} &\phi_h^{[1]} \\ q_{n+1} &= q_n \\ p_{n+1} &= p_n + hf(q_n) \\ &\phi_h^{[2]} \\ q_{n+1} &= q_n + \frac{h}{m} p_n \\ p_{n+1} &= p_n \end{aligned}$$

so that

$$\Phi_h = \phi_h^{[1]} \circ \phi_h^{[2]} \tag{B.42}$$

Combining principles of *composition*, *symmetry*, and *splitting*, a general rule for the generation of symmetric symplectic methods of high order can be formulated.[9] As an example, we can see the Euler method in equation B.40, split into two flows and compose with its adjoint, and simplified. We obtain:

$$\begin{aligned} \Phi^*_{\frac{h}{2}} \circ \Phi_{\frac{h}{2}} &= (\phi_{\frac{h}{2}}^{[1]} \circ \phi_{\frac{h}{2}}^{[2]})^* \circ (\phi_{\frac{h}{2}}^{[1]} \circ \phi_{\frac{h}{2}}^{[2]}) \\ &= \phi_{\frac{h}{2}}^{[2]} \circ \phi_{\frac{h}{2}}^{[1]} \circ \phi_{\frac{h}{2}}^{[1]} \circ \phi_{\frac{h}{2}}^{[2]} \\ &= \phi_{\frac{h}{2}}^{[2]} \circ \phi_{h}^{[1]} \circ \phi_{\frac{h}{2}}^{[2]} \end{aligned} \tag{B.43}$$

This is a symmetric method of the second order. The above equation may also be rewritten as:

$$\begin{aligned} q_{n+\frac{1}{2}} &= q_n + \frac{h}{2m} p_n \\ p_{n+1} &= p_n + hf(q_{n+\frac{1}{2}}) \\ q_{n+\frac{1}{2}} &= q_{n+\frac{1}{2}} + \frac{h}{2m} p_{n+1} \end{aligned} \tag{B.44}$$

This is also known as the *Strömer-Verlet method*.

By reapplying *composition* and *splitting* to the above equation B.43, we can deduce higher-order symmetric integration schemes. Furthermore, these methods can be generalised and be applied to multi-dimensional dynamical systems where the flow of the system can be reformulated as a composition of simple flows along each dimension. For instance, a $n$ dimensional dynamical system governed by the

[9] Gilbert Strang. On the construction and comparison of difference schemes. *SIAM Journal on Numerical Analysis*, 5 (3):506-517, 1968; and Robert I McLachlan and G Reinout W Quispel. Splitting methods. *Acta Numerica*, 11:341-434, 2002

flow $\Phi_h$:

$$\begin{aligned} \dot{x}_1 &= f_1(x_1, x_2, \ldots, x_n) \\ \dot{x}_2 &= f_2(x_1, x_2, \ldots, x_n) \\ &\vdots \\ \dot{x}_n &= f_n(x_1, x_2, \ldots, x_n) \end{aligned}$$

can be reformulated as a splitting into $n$ first-order flows

$$\Phi_h = \phi_h^1 \circ \phi_h^2 \circ \cdots \circ \phi_h^n$$

and therefore, using the adjoint, a second-order symmetric method would be:

$$\Phi^*_{\frac{h}{2}} \circ \Phi_{\frac{h}{2}} = \phi^n_{\frac{h}{2}} \circ \cdots \circ \phi^2_{\frac{h}{2}} \circ \phi^1_h \circ \phi^2_{\frac{h}{2}} \circ \cdots \circ \phi^n_{\frac{h}{2}} \tag{B.45}$$

This is the integration method I used in the second formulation of *rattle* for integrating arbitrary multi-dimensional dynamical systems. To formulate a fourth-order symmetric and symplectic integration method of the above, one would simply use composition and write the method:

$$\Phi^*_{\frac{h}{4}} \circ \Phi_{\frac{h}{4}} \circ \Phi^*_{\frac{h}{4}} \circ \Phi_{\frac{h}{4}} \tag{B.46}$$

and so on for higher orders.

# C

# *Phase Space Experiment*

This Appendix contains a detailed technical description of a small case study called *Phase Space Experiment* previously introduced in 4.3.1 Phase Space Thinking: An Experiment. The aim of the experiment was to put to test the consequences of an understanding of interaction environments in terms of *dynamical systems thinking* and their possible realisation. In this setup, two performers were asked to interact with a computer music system whose sound output is modulated by the evolution of a simple dynamical system, which is, in turn, perturbed and influenced by their playing.

For this experiment, a simple type of two-dimensional dynamical system has been chosen, resulting in one of the most prototypical attractor types: the centre attractor (the attractor of the simple harmonic oscillator). This dynamical system (DS in the following) produces a flow in the two-dimensional phase plane, inducing each state, identified by its abscissa and ordinate values, to vary and move when time advances: its evolution will inscribe trajectories in the phase-space accordingly to its specific attractor. The subsequent $x$ and $y$ coordinates a trajectory will traverse in time are used to modify salient characteristics of the output sound. In the following we have decided to use the value for controlling the *transposition* factor of the projected sound with the abscissa.

The involved performers are musicians who are asked to react to the sound produced by the computer music system by playing their instruments. The instrument's sound is picked up by a microphone, analysed, and recorded: a contact microphone is used in order to allow the musician to move and to keep a coherent recording level throughout the experiment. Thus, *only sound* is used as input and output in this experiment – no other sensing technology as motion tracking is employed.

A crucial element is how the coupling between the performer and computer music system is formulated. That is, how does the analysed sound of the musician influence

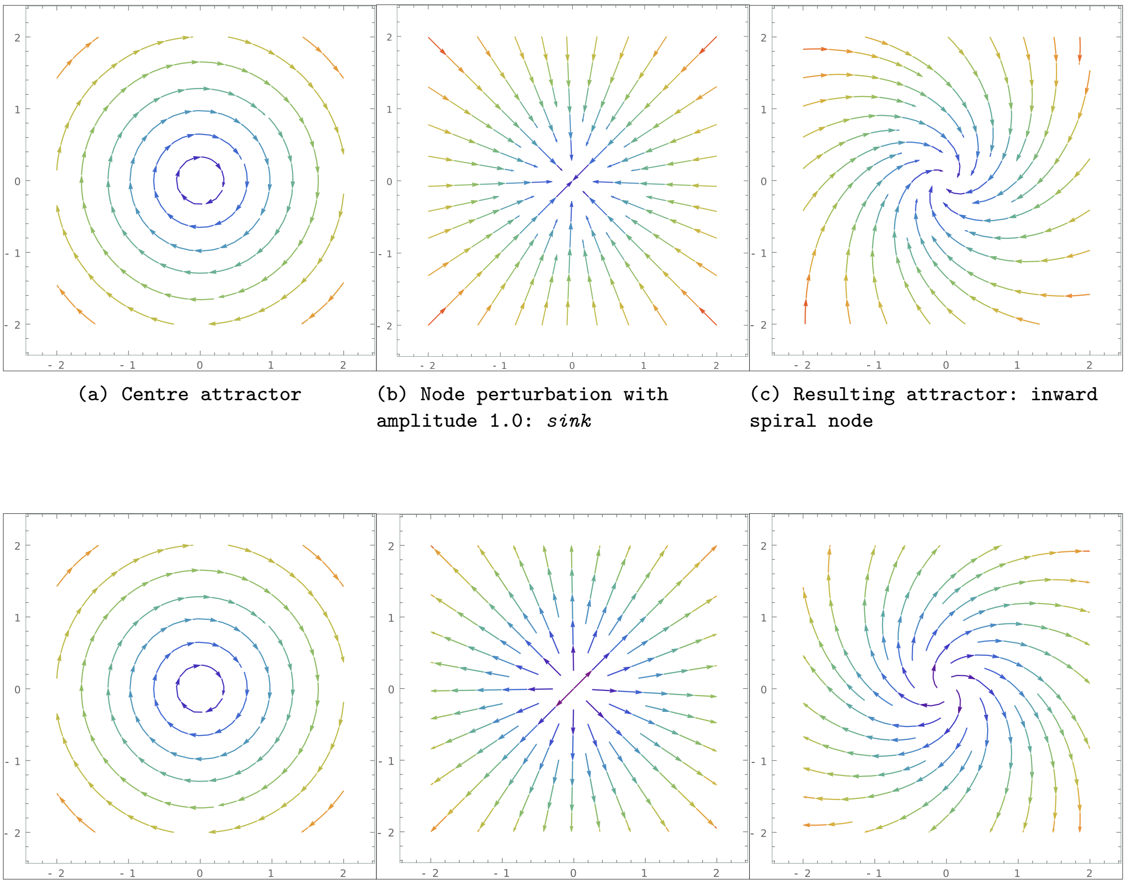

(a) Centre attractor

(b) Node perturbation with amplitude 1.0: *sink*

(c) Resulting attractor: inward spiral node

(d) Centre attractor

(e) Node perturbation with amplitude -1.0: *source*

(f) resulting attractor: outward spiral node

the DS's evolution? In the first implementation, the coupling was understood as a second *perturbing* DS, which then modified the unperturbed DS. The magnitude of of this system's influence is modulated in dependence of the input sound's features. This perturbing DS is a *node*-type attractor. Exemplifying the effect of combining the two attractors, in figure C's top row, a centre attractor is perturbed with a node attractor with magnitude 1.0: the resulting attractor is an asymptotically stable inward spiral, in which the phase space trajectories spiral down towards the origin. Alternately, in figure C's bottom row, the centre attractor is perturbed with a *source*, i.e. a node attractor with magnitude −1.0: in this case the result is the asymptotically unstable outward spiral, which causes trajectories to spiral out from the origin.

This DS could be expressed mathematically using the Jacobi matrices formalism with the following system:

$$\begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} = \left[ \begin{pmatrix} 0 & a \\ -a & 0 \end{pmatrix} + p \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \right] \begin{pmatrix} x \\ y \end{pmatrix}, p \in [-1, 1] \tag{C.1}$$

Here $p$ we see the perturbation's magnitude, while $a$ is the harmonic oscillator's period of oscillation. At first, this factor was chosen such that its time interval was at ca. $2.6s$, four times the maximal '*salience of pulse sensation*'[1], which lies approximately at $600ms$. This choice would ensure that the musicians could easily hear the period of the inherent oscillation produced by the unperturbed DS. However, this value has been left variable in order to allow adjusting during the experiment.

[1] Richard Parncutt. A perceptual model of pulse salience and metrical accent in musical rhythms. *Music Perception: An Interdisciplinary Journal*, 11(4): 409-464, 1994

In this experiment, the value of $p$ was made dependent on the input sound's instantaneous *RMS variation*. The RMS value is computed over a variable time window ranging from $10ms$ to $1s$. After having computed its variation, this value is scaled, mapped, and clipped in the value range from $-1.0$ to $1.0$ through a specialised sigmoid function.

To this end, the input signal $s$ is written into a ringbuffer. Furthermore, the *RMS* of the current input signal $s[n]$ is computed using the following algorithm which allows for fast sample per sample calculation:

$$sumSquared[n] = sumSquared[n-1] - s[n - rmsSize]^2 + s[n]^2 \tag{C.2}$$

$$rms[n] = \sqrt{\frac{sumSquared[n]}{rmsSize}} \tag{C.3}$$

where *rmsSize* is the chosen RMS window size. Next the variation of the *RMS* with respect to its value *rmsDel* samples before is computed and passed through a sigmoid function:

$$drms[n] = sigmoid(rms[n] - rms[n - rmsDel], p, g) \tag{C.4}$$

where $rmsDel = 512$. The specialised sigmoid function is implemented using the formula:

$$sigmoid(x, p, g) = \left(1 + \frac{2}{exp(p) - 1}\right)\left(\frac{2}{exp\left(\frac{p|x|^g}{x}\right)} - 1\right) \tag{C.5}$$

which allows for a small 'gating' region around the origin in dependence of the factor $g$ (see figure C.1). In this implementation $g = 3.0$ and $p = 7.0$.

The value of $drms[n]$ is again smoothed using an integrator (also know as lagin, for example, the *SuperCollider* programming language) with a $0.1s$ seconds $t_{60}$ time constant:

$$lagRms = exp\left(\frac{log10(0.001)}{10 * 44100}\right)$$

$$drmsL[n] = drms[n] + (drmsL[n-1] - drms[n]) * lagRms \tag{C.6}$$

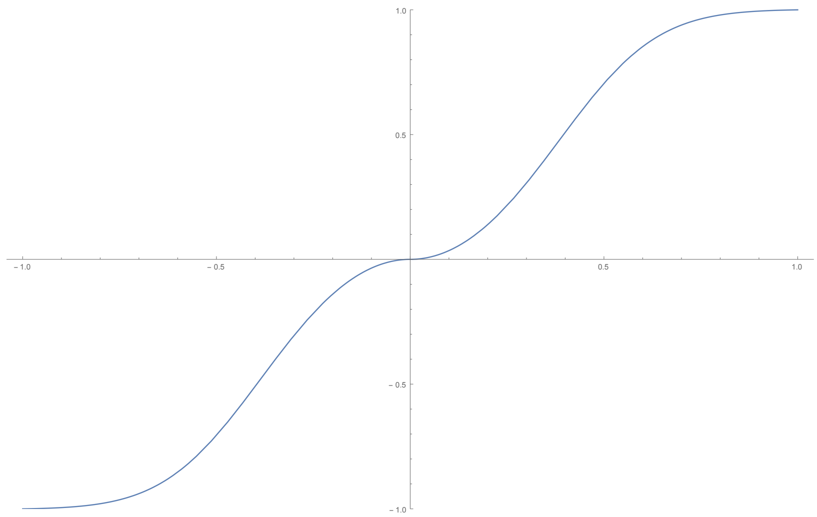


Figure C.1: specialised sigmoid function with $p = 7.0$ and $g = 3.0$

This is eventually rescaled to be used as the magnitude $p$ of the attractor perturbation as in equation C.1

$$p = magPer * drmsL[n] \tag{C.7}$$

The parameters for the input signal conditioning stage have been chosen such that slow crescendi or decrescendi would have the maximum effect in perturbing the underlying attractor, whereas short or impulsive changes in the input signal's amplitude have a minimal impact on the evolution of the dynamical system.

In detail, a constant crescendo leads to a constant positive derivative of the RMS. As this crescendo is slow the variation and thus the derivative will not be very big, and its value would map in the linear positive region of the sigmoid. A fast change in the input RMS would lead to very high value returning to a very small value after a very short time. The subsequent integrator step would then minimise the effect even more.

In order to avoid that the state of the system growing too large due to the perturbations, an additional flow field has been applied, which would drag the current state towards the plane origin when its distance from it is greater than a certain threshold.

$$\texttt{lim}(x, y) = \begin{cases} \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}, & \text{if } r \geq thresh \\ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, & \text{otherwise} \end{cases} \tag{C.8}$$

where $r = \sqrt{(x^2 + y^2)}$ is the state's distance to the origin and $thresh = 2.0$. Added to the flows resulting from the system as in equations C.1, this vector field would ensure that the system's state would stay mostly within the region with $r \leq 2.0$ and not grow indefinitely.

The plane origin is a singular point for both dynamical systems. Especially in the scenario of the centre attractor, once the state of the system reaches this point of asymptotic stability, it would be impossible for it to leave this position as the flow in this point, since even with perturbation it will always be $(0, 0)$. In order to avoid this situation, which would eventually stop the evolution of the system, a second fixed flow field has been added to the attractor flows.

$$\mathbf{push}(x, y) = \begin{cases} \begin{pmatrix} rand() & 0 \\ 0 & rand() \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}, & \text{if } r \leq floor \\ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, & \text{otherwise} \end{cases} \tag{C.9}$$

where $rand()$ stands for a random number between $-1$ and 1 generated anew at each frame and $floor = 0.03$.

Everything has been implemented in a Fortran version of the *rattle* framework, and centred around the idea of phase space construction or the geometrical representation of dynamical systems (see appendix B *rattle* integration algorithms).

Taking the RMS input signal as the only interaction parameter was a choice motivated by the desire to reduce the experiment's complexity and to provide an intuitive and simple parameter for the performer: RMS should be tightly related to the felt effort or the intensity of the playing. All other choices made during the implementation phase of the experiment were taken with the aim of helping the musicians find out how the system works, for it to react to their play while interacting with it, and then for them to consciously engage with it.

The sound produced by the computer music system and heard by the musician is generated using two models. The musician's instrument sound is recorded and played back with a short delay of $5s$ using granular synthesis. The sound is transposed according the DS state's abscissa value remapped exponentially in the range from 0.5 to 2.0 (i.e. $\pm 1$ octave transposition). Thus, the sound would both give information about the DS's state evolution and about the musician's input.

With the previous implementation of coupling, initial informal tests showed that the system would be very difficult to cope with. In particular, there was clear tendency for the system to grow in energy as interaction

Figure C.2: Musicians Joel Diegert (left) and Lorenzo Derinni (right) while engaging with the phase space experiment setup.

with the performer through the node type of attractor pushed the paths of the system towards bigger orbits. Also, if decreasing RMS variation brought the system to states near the origin of the phase space, a substantially bigger effort was needed to bring the system once again to a path with a more sensible evolution.

We therefore sought a different, simpler and more directly controllable implementation of the coupling. This second implementation, in terms of a Jacobi matrix formulation, could be written as:

$$\begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} = \begin{pmatrix} 0 & a \\ -a & 0 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} 0 \\ p \end{pmatrix}, p \in [-1, 1] \tag{C.10}$$

where $p$ remains the same as in the previous version.

With this modification, a positive (increasing RMS) perturbation would push the current DS state towards the positive $x$ axis, while a negative perturbation (decreasing RMS) towards the negative (see figure C). In addition, a weak attractor of the node type has been added to the base DS. The resulting effect would be a small but constant loss of energy from the system, which would then slowly spiral down towards the phase plane origin. That is, this attractor would act as a sort of attrition factor to the whole system. The magnitude of this attractor would be modulated with the system's state distance from the origin,

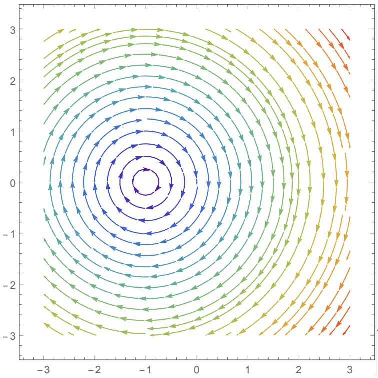


(a) Centre attractor perturbed according to equation C.10 with $p = -1.0$

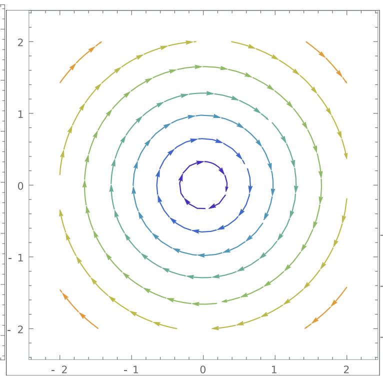


(b) Unperturbed centre attractor

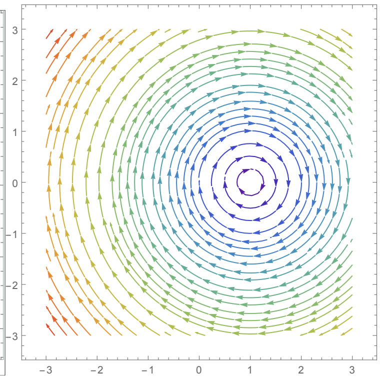


(c) Centre attractor perturbed with positive $p = 1.0$

so that attrition would be turned off when the system's own oscillations would fall under a certain threshold. This means that the system would never entirely die out (phase state at the phase plane origin) but rather always preserve some activity of its own.

As a consequence of these changes, the performer could, on the one hand, bring the system into resonance by applying the right push at the right moment during the system's evolution: i.e., producing an increase of the RMS when the system state is in the $x < 0.0$ half-plane or a decrease of RMS when the system is in the $x > 0.0$ region. On the other hand, with the action of the 'attrition' factor, the system would prevent any non-controlled growth of its energy, continuously digesting input energy while retaining a base amount of activity.

# D

# DBAP and ADBAP

In some of the works presented here, I make use of a simple algorithm to spatialise a sound source over a loudspeaker array in known positions. The algorithm is an extended and modified version of *Distance Based Amplitude Panning* (DBAP) algorithm.[1] As this algorithm makes no *a priori* assumption of the effective loudspeaker setup, and no assumptions as to where the listeners are situated in the venue, this spatialisation method can be used very flexibly - a fact enhanced by its relative low computational cost. It was therefore a natural choice when working with non-standard speaker distributions, which are necessary in spaces where predefined speaker layouts cannot be applied or, as in some of the works I present here, the speaker layout itself becomes part of and artistic endeavour.

[1] Trond Lossius, Pascal Baltazar, and Théo de la Hogue. Dbap-distance-based amplitude panning. In *Proceedings of the Internatinal Computer Music Conference*, pages 489-492, 2009

The method is a panning algorithm modulating the amplitude $a_i$, by which a sound source is projected in an inverse dependence of the Cartesian distance $d_i$ of the (virtual) sound source $s$ to the a loudspeaker $i$.[2]

[2] Although original formulation of the method is based on a two-dimensional spatial representation of source and loudspeaker positions, the extension implements a a three-dimensional version.

$$a_i = \frac{k}{d_i^p} \qquad \text{(D.1)}$$

where $p$ is an exponent coefficient calculated from the rolloff $R$ in decibels per doubling of distance

$$p = \frac{R}{20\,log_{10} 2} \qquad \text{(D.2)}$$

Setting $R = 6dB$ equals to the inverse distance law for free-field sound propagation.

Extending the principle of constant intensity stereo panning, the original DBAP method assumes that overall intensity is constant over the whole array regardless of the virtual source's position. Therefore the sum of all squared amplitudes should be normalised to 1.

$$\sum_i a_i^2 = 1 \qquad \text{(D.3)}$$

and the factor $k$ in equation D.1 is then computed accordingly so that this normalisation holds.

$$k = \frac{1}{\sqrt{\sum_i 1/d_i^{2p}}} \tag{D.4}$$

which also ensures that the loudspeaker amplitudes remain in the range $0 < a_i < 1$ for any distance, including $d_i = 0$.

Eventually, a *blurring factor* $b$ is introduced in the calculation of the distances in order to adjust for too-sharp changes in the amplitude distribution, i.e., the spatial spread when some $d_i = 0$. If $(x_s, y_s, z_s)$ is the three-dimensional position of the virtual sound source and $(x_i, y_i, z_i)$ the position of the $i$ loudspeaker:

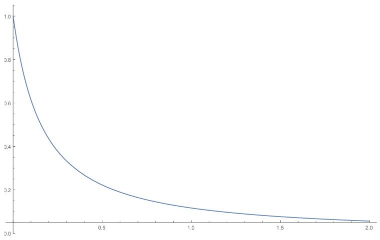

Figure D.1: One speaker DBAP amplitude as a function of distance without blur

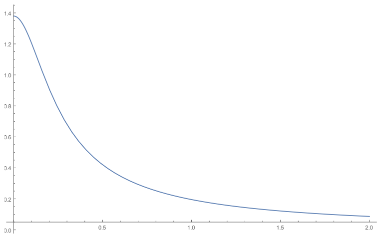

Figure D.2: One speaker DBAP amplitude as a function of distance with blur

$$d_i = \sqrt{(x_i - x_s)^2 + (y_i - y_s)^2 + (z_i - z_s)^2 + b^2} \tag{D.5}$$

that provides smoother variations around $d_i = 0$ (see figures D.1 and D.2).

Having specified the rolloff and blurring coefficients, the distances from the source to the associated loudspeaker objects are computed and used to determine the relative amplitudes of the projected sound. As a consequence of implementing the principle of constant intensity, positions outside the loudspeaker field cannot be clearly rendered: in this region, the relative amplitude differences tend towards zero with increasing distance, while the overall intensity remains constant, resulting in a spatially undifferentiated sound output. The resulting overall intensity is constant, regardless of the position of the source.

During the course of our case studies, it became necessary to spatialise sources that could also travel out of the loudspeaker field and completely disappear. To achieve this, we modified the DBAP algorithm, removing the constant intensity condition. Sound spatialisation is achieved by defining a distribution of absolute (rather than relative) amplitudes. This creates sources that move sufficiently far from the loudspeaker array to fade out. Furthermore, the trajectory of moving sounds appears more clearly shaped, or sharper, compared to the unmodified DBAP algorithm. Lacking a more explanatory name, we have called this simplified version of the DBAP algorithm *Absolute Distance Based Amplitude Panning* (ADBAP).

Using the blurred distance method introduced in equation D.5, the ADBAP would then compute the amplitude of the $i$-th loudspeaker as:

$$a_i = \left(\frac{b}{d_i}\right)^p \tag{D.6}$$

which ensures that $0 \leq a_i \leq 1$ for any distance. In this case, one can imagine the effect of the blur as widening the source; even with this slight modification to the natural laws of sound propagation,the ADBAP method would continue to provide the correct behaviour of the loudspeaker amplitudes with respect to the inverse distance law. In fact, the slope of the amplitude function D.6 is proportional to the derivative of the same function without blurring. That is, the variations in distance would produce similar variations with or without blurring factor, in particular for distances $d_i \gg b$ . This is particularly important for moving sound sources, which is the case in most of the works I refer to here.

Figure D.3 shows the slope of the function D.6 for different blur factors.

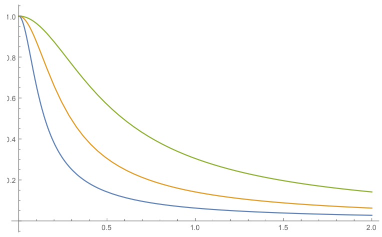

Figure D.3: Behaviour of the function D.6 in dependence of the distance $d_i$ for different blur factors: $b = 0.1$ corresponds to the blue function, $b = 0.2$ to the orange and $b = 0.4$ to the green

# E

# Own Publications

I collect here a list of the research I've worked on and has been published in the course of this dissertation.

D. Pirrò. Staging collisions: On behaviour. In Michael Schwab, editor, *Transpositions*, number 1 in Aesthetico-Epistemic Operators in Artistic Research. Leuven University Press, forthcoming in 2018.

Marian Weger, David Pirrò, and Robert Höldrich. Evaluation of an acoustic interface for tremor analysis. In *Proceedings of the 14th Sound and Music Computing Conference*, pages 234-241, Espoo, Finland, 2017.

Georgios Marentakis, David Pirrò, and Marian Weger. Creative evaluation. In *Proceedings of the 2017 Conference on Designing Interactive Systems*, pages 853-864. ACM, 2017.

Hanns Holger Rutz and David Pirrò. Anemone actiniaria. In *Proceedings of the 2016 xCoAx Conference, Computation, communication, Aesthetics and X*, pages 404-408, 2016.

Marian Weger, David Pirrò, Alexander Wankhammer, and Robert Höldrich. Discrimination of tremor diseases by interactive sonification. In *5th Interactive Sonification Workshop (ISon)*, volume 12, 2016.

Georgios Marentakis, David Pirro, and Raphael Kapeller. Zwischenräume-a case study in the evaluation of interactive sound installations. In *Proceedings of the International Computer Music Conference ICMC*, 2014.

Katharina Vogt, David Pirrò, and Robert Höldrich. Waveguides for model-based sonification. In *Proceedings of the 9th Sound and Music Computing Conference*, Copenhagen, Denmark, 2012.

Gerhard Eckel, Martin Rumori, David Pirrò, and Ramón González-Arroyo. A framework for the choreography of sound. In *Proceedings of the 38th International Computer Music Conference*, pages 404-511, Ljubljana, Slovenia, 2012.

Georgios Marentakis and David Pirrò. Exploring sound and spatialization design on speaker arrays using physical modelling. In *Proceedings of the 9th Sound and Music Computing Conference*, pages 55-60, Copenhagen, Denmark, 2012.

David Pirrò, Alexander Wankhammer, Petra Schwingenschuh, Alois Sontacchi, and Robert Höldrich. Acoustic interface for tremor analysis. In *Proc. of the 18th International Conference on Auditory Display*, pages 210-213, Atlanta, Georgia, 2012.

Katharina Vogt, David Pirrò, Martin Rumori, and Robert Höldrich. Sounds of simulations: data listening space. In *Proceedings of the 38th International Computer Music Conference*, pages 525-530, Ljubljana, Slovenia, 2012.

David Pirrò and Gerhard Eckel. Physical modelling enabling enaction: an example. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 461-464, Oslo, 2011.

Katharina Vogt, David Pirrò, Ingo Kobenz, Robert Höldrich, and Gerhard Eckel. Physiosonic-evaluated movement sonification as auditory feedback in physiotherapy. In *Auditory display*, pages 103-120. Springer, 2010.

Katharina Vogt, Robert Höldrich, David Pirrò, and Christof Gattringer. Spin quartets, sonification of the xy model. In *Proceedings of the 16th International Conference on Auditory Display*, pages 97-102, Washington, DC, 2010.

Katharina Vogt, Robert Höldrich, David Pirrò, Martin Rumori, Stefan Rossegger, Werner Riegler, and Matevz Tadel. A sonic time projection chamber, sonified particle detection at cern. In *Proceedings of the 16th International Conference on Auditory Display*, pages 103-108, Washington, DC, 2010.

Anna Saranti, Gerhard Eckel, and David Pirrò. Quantum harmonic oscillator sonification. In Sølvi Ystad, Mitsuko Aramaki, Richard Kronland-Martinet, and Kristoffer Jensen, editors, *Auditory Display*, volume 5954 of *Lecture Notes in Computer Science*, pages 184-201. Springer Berlin Heidelberg, 2010.

Gerhard Eckel and David Pirrò. Motion-enabled live electronics. In *Sound and Music Computing Conference*, pages 36-41, Porto, Portugal, 2009.

Gerhard Eckel and David Pirrò. On artistic research in the context of the project embodied generative music. In *Proceedings of the 35th International Computer Music Conference*, pages 541-544, Montréal, 2009.

# *Bibliography*


Ralph Abraham and Christopher D. Shaw. *Dynamics-the geometry of behavior*. Addison-Wesley, Advanced Book Program, Redwood City, Calif., 1992.

Newton Armstrong. *An enactive approach to digital musical instrument design*. PhD thesis, Princeton University, 2006.

Karen Barad. *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning*. Duke University Press, Durham & London, 2007.

Xabier E Barandiaran, Ezequiel Di Paolo, and Marieke Rohde. Defining agency: Individuality, normativity, asymmetry, and spatio-temporality in action. *Adaptive Behavior*, 17 (5):367-386, 2009.

Randall D Beer. A dynamical systems perspective on agent-environment interaction. *Artificial intelligence*, 72(1-2):173-215, 1995.

Alain Berthoz. *The brain's sense of movement*. Harvard University Press, 2000.

Frédéric Bevilacqua, Jeff Ridenour, and David J Cuccia. 3d motion capture data: motion analysis and mapping to music. In *Proceedings of the workshop/symposium on sensing and input for media-centric systems*, 2002.

Bert Bongers. Physical interfaces in the electronic arts. In M.M. Wanderley and M. Battier, editors, *Trends in gestural control of music, Paris, IRCAM/Centre Pompidou*, pages 41-70. IRCAM-Centre Pompidou, 2000.

Jonas Braasch. A loudspeaker-based 3d sound projection using virtual microphone control (vimic). In *Audio Engineering Society Convention 118*. Audio Engineering Society, 2005.

Herbert Brün. *über Musik und zum Computer*. G. Braun, 1971.

Claude Cadoz and Marcelo M. Wanderley. Gesture - music. In M.M. Wanderley and M. Battier, editors, *Trends in gestural control of music, Paris, IRCAM/Centre Pompidou*, 2000.

Claude Cadoz, Annie Luciani, and Jean Loup Florens. Cordis-aniama: Modeling and simulation system for sound and image synthesis - the general formalism. *Computer Music Journal*, 17(1):19 - 29, Spring 1993.

John Cage. John cage: An autobiographical statement, 1990. URL http://johncage.org/autobiographical_statement.html. Accessed on 29/10/2017.

John Cage and Roger Reynolds. An interview with john cage on the occasion of the publication of silence. *Generation - The University Inter-Arts Magazine*, pages 40-51, November 1961.

Italo Calvino. *If on a Winter's Night a Traveler*. Houghton Mifflin Harcourt, 1981.

Nicolas Castagné and Claude Cadoz. Genesis: a friendly musician-oriented environment for mass-interaction physical modeling. In *ICMC 2002-International Computer Music Conference*, pages 330-337. MPublishing, 2002.

Joel Chadabe. Interactive composing: An overview. *Computer Music Journal*, 8(1):22-27, 1984.

Joel Chadabe. The history of electronic music as a reflection of sructural paradigms. *Leonardo Music Journal*, 16:41-44, 1996.

Joel Chadabe. *Electric Sound:The Past and Promise of Electronic Music*. Prentice-Hall, Upper Saddle River, New Jersey, 1997.

Joel Chadabe. The limitations of mapping as a structural descriptive in electronic instruments. In *Proceedings of the 2002 conference on New interfaces for musical expression*, pages 1-5. National University of Singapore, 2002.

Andy Clark. *Being there: Putting brain, body, and world together again*. MIT press, 1998.

Andy Clark and David Chalmers. The extended mind. *Analysis*, 58(1):7-19, 1998.

Nick Collins. The analysis of generative music programs. *Organised Sound*, 13(3):237-248, 2008.

Juliet Corbin and Anselm Strauss. Grounded theory research: Procedures, canons and evaluative criteria. *Zeitschrift für Soziologie*, 19(6):418-427, 1990.

Rene de Vogelaere. Methods of integration which preserve the contact transformation property of the hamiltonian equations. *Department of Mathematics, University of Notre Dame, Report*, 4:30, 1956.

Ezequiel A Di Paolo. Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the cognitive sciences*, 4(4): 429-452, 2005.

Agostino Di Scipio. Iterated nonlinear functions as a sound-generating engine. *Leonardo*, 34(3):249-254, 2001.

Agostino Di Scipio. 'Sound is the interface': from interactive to ecosystemic signal processing. *Organised Sound*, 8(3):269-277, 2003.

Agostino Di Scipio. Listening to yourself through the otherself: on background noise study and other works. *Organised Sound*, 16(2):97-108, 2011.

Luc Döbereiner. Models of constructed sound: Nonstandard synthesis as an aesthetic perspective. *Computer Music Journal*, 35(3):28-39, 2011.

Christopher Dobrian and Frédéric Bevilacqua. Gestural control of music: using the vicon 8 motion capture system. In *Proceedings of the 2003 conference on New interfaces for musical expression*, pages 161-163. National University of Singapore, 2003.

Paul Dourish. *Where the Action is: The Foundations of Embodied Interaction*. The MIT Press, 2001.

Hubert L Dreyfus. The current relevance of Merleau-Ponty's phenomenology of embodiment. *The Electronic Journal of Analytic Philosophy*, 4:1-16, 1996.

Jon Drummond. Understanding interactive systems. *Organised Sound*, 14(2):124-133, 2009.

Gerhard Eckel. Embodied generative music. In Deniz Peters, Gerhard Eckel, and Andreas Dorschel, editors, *Bodily expression in Electronic Music*, chapter 10, pages 143 - 151. Routledge, 2012.

Gerhard Eckel and David Pirrò. On artistic research in the context of the project embodied generative music. In *Proceedings of the 35th International Computer Music Conference*, pages 541-544, Montréal, 2009.

Umberto Eco. *Opera aperta*. Harvard University Press, 1989.

Umberto Eco. *How to write a thesis*. MIT Press, Cambridge, Massachusetts, 2015.

Alfred O Effenberg. Movement sonification: Effects on perception and action. *IEEE multimedia*, 12(2):53-59, 2005.

Richard P Feynman, Robert B Leighton, Matthew Sands, et al. *The Feynman lectures on physics, Vol. 2*. Addison-Wesley, 1964.

Richard Phillips Feynman, Robert B Leighton, and Matthew Sands. *The Feynman lectures on physics. Vol. 1*. Addison-Wesley, 1963.

Guy E Garnett. The aesthetics of interactive computer music. *Computer Music Journal*, 25(1):21-33, 2001.

James J Gibson. *The ecological approach to visual perception: classic edition*. Psychology Press, 2014.

Bob Gilmore. Five maps of the experimental world. *Artistic Experimentation in Music: An Anthology*, pages 23-29, 2014.

Barney Glaser. *Discovery of grounded theory: Strategies for qualitative research*. Routledge, 2017.

Rolf Inge Godøy. Motor-mimetic music cognition. *Leonardo*, 36(4):317-319, 2003.

Michael Gogins. Iterated functions systems music. *Computer Music Journal*, 15(1):40-48, 1991.

Herbert Goldstein, Charles Poole, and John Safko. *Classical mechanics*. Addison Wesley, 2002.

Michael Gurevich and Jeffrey Treviño. Expression and its discontents: toward an ecology of musical creation. In *Proceedings of the 7th international conference on New interfaces for musical expression*, pages 106-111. ACM, 2007.

Ernst Hairer, Christian Lubich, and Gerhard Wanner. *Geometric numerical integration: structure-preserving algorithms for ordinary differential equations*, volume 31. Springer Science & Business Media, 2006.

Martin Heidegger. *Being and time: A translation of Sein und Zeit*. SUNY press, 1996.

Cyrille Henry. Physical modeling for puredata (pmpd) and real time interaction with an audio synthesis. In *Proc. of the Sound and Music Computing Conference*, October 2004.

Lejaren Arthur Hiller and Leonard M Isaacson. *Experimental Music; Composition with an electronic computer*. Greenwood Publishing Group Inc., 1979.

Andy Hunt, Marcelo M Wanderley, and Ross Kirk. Towards a model for instrumental mapping in expert musical interaction. In *Proc. of the 2000 International Computer Music Conference*, pages 209-211, 2000.

Don Ihde. *Technology and the lifeworld: From garden to earth*. Indiana University Press, 1990.

Jonathan Impett. Interaction, simulation and invention: a model for interactive music. In *Proceedings of ALMMA 2001 Workshop on Artificial Models for Musical Applications*, pages 108-119, Cosenza, Italy, 2001.

Robert Irwin. The state of the real. In Beatrice Hohenegger, editor, *Notes Towards a Conditional Art*, chapter 7, pages 49 - 53. Getty Publications, 1972a.

Robert Irwin. Re-shaping the shape iof things. In Beatrice Hohenegger, editor, *Notes Towards a Conditional Art*, chapter 8, pages 54 - 60. Getty Publications, 1972b.

Sergi Jorda. *Digital Lutherie Crafting musical computers for new musics' performance and improvisation*. PhD thesis, Department of Information and Communication Technologies, 2005.

Ajay Kapur, George Tzanetakis, Naznin Virji-Babul, Ge Wang, and Perry R Cook. A framework for sonification of vicon motion capture data. In *Conference on Digital Audio Effects*, pages 47-52, 2005.

Abbas I Abdel Karim. Criterion for the stability of numerical integration methods for the solution of systems of differential equations. *J. Res. NBS*, 718, 1967.

Jin Hyun Kim and Uwe Seifert. Embodiment and agency: Towards an aesthetics of interactive performativity. In *Proceedings of the 4th Sound and Music Computing Conference*, pages 230-237, 2007.

Gottfried Michael Koenig. Kompositionsprozesse. In *Ästhetische Praxis*, volume 3 of *Texte zur Musik*, pages 191-210. PFAU Verlag, Saarbrücken, 1993.

Steven E Koonin, Dawn C Meredith, and William H Press. Computational physics: Fortran version. *Physics Today*, 44:112, 1991.

Susan Kozel. Embodying the sonic invisible. In Deniz Peters, Gerhard Eckel, and Andreas Dorschel, editors, *Bodily expression in Electronic Music*, chapter 10, pages 61 - 70. Routledge, 2012.

Tellef Kvifte. On the description of mapping structures. *Journal of New Music Research*, 37(4):353-362, 2008.

Marc Leman. *Embodied music cognition and mediation technology*. MIT Press, 2008.

Marc Leman. Music, gesture, and the formation of embodied meaning. In Rolf Inge Godøy and Marc Leman, editors, *Musical gestures: Sound, movement, and meaning*, pages 126-153. Routledge New York and Abingdon, England, 2010.

Marc Leman and Rolf Inge Godøy. Why study musical gestures. In Rolf Inge Godøy and Marc Leman, editors, *Musical gestures. Sound, movement, and meaning*, pages 3-11. Routledge New York, NY, 2010.

George E. Lewis. Too many notes: Computers, complexity and culture in "voyager". In *Leonardo Music Journal*, volume 10, pages 33-39, 2000.

Sol LeWitt. Paragraphs on conceptual art. *Artforum*, 5(10): 79-83, 1967.

Edward N Lorenz. Deterministic nonperiodic flow. *Journal of the atmospheric sciences*, 20(2):130-141, 1963.

Trond Lossius, Pascal Baltazar, and Théo de la Hogue. Dbap-distance-based amplitude panning. In *Proceedings of the Internatinal Computer Music Conference*, pages 489-492, 2009.

D.G. Luenberger. *Introduction to Dynamic Systems: Theory, Models, and Applications*. Wiley, 1979.

Peter Manning. *Electronic and computer music*. Oxford University Press, 2013.

Georgios Marentakis, David Pirrò, and Raphael Kapeller. Zwischenräume - a case study in the evaluation of interactive sound installations. In *Proceedings of the Joint 11th Sound and Music Computing Conference and the 40th International Computer Music Conference*, pages 277-284, Athens, 2014.

Humberto R Maturana and Francisco J Varela. Autopoiesis: the organisation of the living. In *Autopoiesis and cognition*, pages 73-135. Springer, 1980a.

Humberto R Maturana and Francisco J Varela. Biology of cognition. In *Autopoiesis and cognition*, pages 2-58. Springer, 1980b.

Humberto R Maturana and Francisco J Varela. *The tree of knowledge: The biological roots of human understanding*. New Science Library/Shambhala Publications, 1987.

Robert I McLachlan and G Reinout W Quispel. Splitting methods. *Acta Numerica*, 11:341-434, 2002.

Dylan Menzies. Composing instrument control dynamics. *Organised Sound*, 7(3):255-266, 2002. ISSN 1355-7718.

Maurice Merleau-Ponty. *Phenomenology of Perception*. Routledge, 2002.

Eduardo Reck Miranda and Marcelo M Wanderley. *New digital musical instruments: control and interaction beyond the keyboard*. AR Editions, Inc., 2006.

Robert A Moog. Voltage-controlled electronic music modules. In *Audio Engineering Society Convention 16*. Audio Engineering Society, 1964.

F Richard Moore. The dysfunctions of midi. *Computer music journal*, 12(1):19-28, 1988.

Roberto Morales-Manzanares, Eduardo F Morales, Roger Dannenberg, and Jonathan Berger. Sicib: An interactive music composition system using body movements. *Computer Music Journal*, 25(2):25-36, 2001.

Tom Mudd, Simon Holland, Paul Mulholland, and Nick Dalton. Dynamical interactions with electronic instruments. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 126-129. Goldsmiths, University of London, 2014.

Frieder Nake. Paragraphs on computer art, past and present. In *Proceedings of CAT 2010 London Conference*, pages 55-63, 2010.

Gerhard Nierhaus. *Algorithmic composition: paradigms of automated music generation*. Springer Science & Business Media, 2009.

Alva Noë. Experience and experiment in art. *Journal of Consciousness Studies*, 7(8-9):123-136, 2000.

Alva Noë. *Action in Perception*. The MIT Press, 2004.

Alva Noë. What would disembodied music even be? In Deniz Peters, Gerhard Eckel, and Andreas Dorschel, editors, *Bodily expression in Electronic Music*, chapter 3, pages 53 - 60. Routledge, 2012.

Donald A Norman. *The psychology of everyday things.(The design of everyday things)*. Basic Books, 1988.

J Kevin O'Regan and Alva Noë. A sensorimotor account of vision and visual consciousness. *Behavioral and brain sciences*, 24(5):939-973, 2001.

Bob Ostertag. Human bodies, computer music. *Leonardo Music Journal*, 21:19-23, 2006.

Garth Paine. Interactivity, where to from here? *Organised Sound*, 7(3):295-304, 2002.

Garth Paine. Towards unified design guidelines for new interfaces for musical expression. *Organised Sound*, 14 (2):142-155, 2009.

Garth Paine. Interaction as material: The techno-somatic dimension. *Organised Sound*, 20(1):82-89, 2015.

Luciana Parisi. *Contagious architecture: computation, aesthetics, and space*. MIT Press, Cambridge, MA, 2013.

Richard Parncutt. A perceptual model of pulse salience and metrical accent in musical rhythms. *Music Perception: An Interdisciplinary Journal*, 11(4):409-464, 1994.

Jana Parviainen. Seeing sound, hearing movement. In Deniz Peters, Gerhard Eckel, and Andreas Dorschel, editors, *Bodily expression in Electronic Music*, chapter 5, pages 71 - 81. Routledge, 2012.

Deniz Peters. Touch. real, apparent, and absent: On bodily expression in electronic music. In Deniz Peters, Gerhard Eckel, and Andreas Dorschel, editors, *Bodily expression in Electronic Music*, chapter 1, pages 17 - 34. Routledge, 2012.

Nils Peters, Tristan Matthews, Jonas Braasch, and Stephan McAdams. Spatial sound rendering in max/msp with vimic. In *Proceedings of the 2008 International Computer Music Conference*, 2008.

Robert F Port and Timothy Van Gelder. *Mind as motion: Explorations in the dynamics of cognition*. MIT press, 1995.

Jeff Pressing. Nonlinear maps as generators of musical design. *Computer Music Journal*, 12(2):35-46, 1988.

Hans-Jörg Reinberger. Experimental systems: Historiality, narration, and deconstruction. *Science in Context*, 7(1): 65-81, 1994.

Eleanor Rosch, Lydia Thompson, and Francisco J Varela. *The embodied mind: Cognitive science and human experience*. MIT press, 1991.

Arturo Rosenblueth, Norbert Wiener, and Julian Bigelow. Bahvior, purpose and teleology. *Philosophy of Science*, 10(1):18 - 24, January 1943.

Robert Rowe. *Interactive music systems: machine listening and composing*. MIT press, 1992.

André Ruschkowski. *Elektronische Klänge und musikalische Entdeckungen*. Reclam, 1998.

Ronald D Ruth. A canonical integration technique. *IEEE Trans. Nucl. Sci.*, 30(CERN-LEP-TH-83-14):2669-2671, 1983.

Andrea Schiavio and Damiano Menin. Embodied music cognition and mediation technology: a critical review. *Psychology of Music*, 41(6):804-814, 2013.

Norbert Schnell and Marc Battier. Introducing composed instruments, technical and musicological implications. In *Proceedings of the 2002 conference on New interfaces for musical expression*, pages 1-5. National University of Singapore, 2002.

Julia H Schröder. Emergence and emergency: Theoretical and practical considerations in agostino di scipio's works. *Contemporary Music Review*, 33(1):31-45, 2014.

John R Searle. *Mind: a brief introduction*. Oxford University Press, 2004.

Wayne Siegel and Jens Jacobsen. The challenges of interactive dance: An overview and case study. *Computer Music Journal*, 22(4):29-43, 1998.

Dan Slater. Chaotic sound synthesis. *Computer Music Journal*, 22(2):12-19, 1998.

Gilbert Strang. On the construction and comparison of difference schemes. *SIAM Journal on Numerical Analysis*, 5(3):506-517, 1968.

Steven H. Strogatz. *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. Westview press, 2014.

Dag Svanæs. *Understanding interactivity: steps to a phenomenology of human-computer interaction*. PhD thesis, Norges teknisk-naturvitenskapelige universitet, 2000.

Esther Thelen. Time-scale dynamics and the development of an embodied cognition. In Robert F. Port and Timothy van Gelder, editors, *Mind As Motion - Explorations in the Dynamics of Cognition*, chapter 3, pages 69-100. MIT Press, 1995.

Leon S Theremin and Oleg Petrishev. The design of a musical instrument based on cathode relays. *Leonardo Music Journal*, 6(1):49-50, 1996.

Evan Thompson. *Mind in life: Biology, phenomenology, and the sciences of mind*. Harvard University Press, 2010.

Evan Thompson and Francisco J. Varela. Radical embodiment: neural dynamics and consciousness. *Trends in Cognitive Sciences*, 5(10):418-425, October 2001.

Tim Van Gelder. The dynamical hypothesis in cognitive science. *Behavioral and brain sciences*, 21(5):615-628, 1998.

Katharina Vogt, David Pirrò, Ingo Kobenz, Robert Höldrich, and Gerhard Eckel. Physiosonic-evaluated movement sonification as auditory feedback in physiotherapy. In *Auditory display*, pages 103-120. Springer, 2010.

Heinz Von Foerster. Objects: tokens for (eigen-)behaviors. *Understanding understanding: Essays on cybernetics and cognition*, pages 261-271, 2003.

Heinz Von Foerster. *Understanding understanding: Essays on cybernetics and cognition*. Springer Science & Business Media, 2007.

Marcelo M Wanderley. Gestural control of music. In *International Workshop Human Supervision and Control in Engineering and Music*, pages 632-644, 2001.

Marcelo M. Wanderley and Philippe Depalle. Gestural control of sound synthesis. *Proceedings of the IEEE 2004*, 92(4): 632 - 644, April 2004.

Simon Waters. Performance ecosystems: Ecological approaches to musical interaction. *EMS: Electroacoustic Music Studies Network*, pages 1-20, 2007.

David Wessel. An enactive approach to computer music performance. *Le Feedback dans la Creation Musical,Lyon: Studio Gramme, France*, pages 93-98, 2006.

Alfred North Whitehead, David Ray Griffin, and Donald W Sherburne. *Process and reality: An essay in cosmology*. University Press Cambridge, 1929.

Todd Winkler. Motion-sensing music: Artistic and technical challenges in two works for dance. In *Proceedings of the International Computer Music Conference*, pages 261-264, 1995.

Todd Winkler. *Composing Interactive Music*. MIT Press, 1998.

Terry Winograd and Fernando Flores. *Understanding computers and cognition: A new foundation for design*. Intellect Books, 1986.

Iannis Xenakis. *Formalized music: thought and mathematics in composition*. Pendragon Press, 1992.