Algorithmically driven synthesis and feedback systems:

an investigation into the aesthetic and technical processes in my compositional practice.

Sten-Olof Hellström

A portfolio of original compositions and commentary submitted to the University of Huddersfield in partial fulfilment of the requirements for the Doctor of Philosophy

January 2013

Contents

Figure Index

Acknowledgements

There are many people that deserve credit for having supported, guided and inspired me and I cannot thank them enough. Please forgive me for only mentioning a selected few people that stick out in my memory as I am writing this and they are:

My Parents Lotte and Perre Hellström; my Partner, fellow composer and performer, workmate, lover and wife Ann Rosén; my supervisors, Prof. Monty Adkins and Dr Pierre Alexandre Tremblay; and in random order without titles: Denis Smalley, Rolf Enström, John M. Bowers, Kai Mikael Jää-Aro, Lars Bröndum, Lennart E Fahlén, Bennett Hogg, Anton Hellström, Staffan John, Lennart Nilsson, Simon Waters, Jan W Morthenson, Simon Vincent and everyone else. Thank you ever so much.

Abstract

This commentary and portfolio of works documents my compositional research into algorithmically driven synthesis and feedback systems. The foundation on which all my work rests consists of self-reference, feedback, emergent properties and algorithmically mediated interaction. Feedback is the lowest common denominator in this process and is for me an essential part of evolution and life and provides a means to create organic sonic forms that generate complex behaviours and phenomena. The predominant approach taken to realise such behaviours is by means of sound synthesis. The preference for sound synthesis is also a means by which compositional elements can interact with each other in a purely sonic manner. To me working with synthesis is like working with organic and living material instead of piling static sounds on top of, or next to each other. In order to work with complex synthesis in an intuitive and direct way I have investigated and developed new methods to control the large number parameters required. To achieve this my main tool has been algorithmically mediated interaction in combination with simple input devices such as a joystick.

When sound files are utilised in a piece they are means to engender a way of working that is outside of my comfort zone. This 'painting oneself into a corner' is an important practical and aesthetic choice and is demonstrated in a number of documented instances in the creative portfolio.

Even though a large part of my work consists of solo work my main source of inspiration comes from working with other people whether it is performing, composing, rehearsing, or building sound installations. All the collaborations I have been part of have been of great importance for me but I have chosen to discuss those that have been significant in the development of my aesthetics, my performance and studio set-up as well as my working methods.

List of works presented for PhD in the order of appearance in the thesis

Sine Cera by Hellström and Bowers (John M Bowers)	43:00	2009
D-Dist by Ann Rosén and Sten-Olof Hellström	5:45	2003
TOE by As Good as it Gets (Hellström and Lars Bröndum)	41:18	2011
SEQUEL	7:39	1995
PREQUEL	5:42	1996
EQUAL	8:26	1997
ON	8:09	1990
JUST	6:24	2007
FLUGAN II	8:04	1990
AMBIENT EDGE	7:04	2006
SPIRAL	8:55	2010
IN MY OWN WORDS	8:36	2012

Chapter 1

1.1 Early influences and background

My earliest musical influences stem from my parent’s record collection of which I remember The Beatles, Vanilla Fudge, Ekseption, Rolf Harris (his didgeridoo, the sound effects he created with metal sheets and various items and of course his Stylophone synthesizer), Karl-Birger Blomdahl's opera *Aniara* (especially the electronic part *Mimans sång*) and a record featuring the Bell Laboratory digital synthesizer. Another possible early influence that I do not recall myself was the Swedish electronic music pioneer Bengt Hambreus who was a friend of our family. When my parents took me to symphony orchestra concerts I was completely blown away by the quality of the sound, that you could hear so many small details and that you were immersed in a fantastic and complex sound world. This was especially true when the orchestra tuned their instruments (which of course was my favourite part). Another important factor in my awakening to the ‘sonic’ world was that when I was a child and staying with my grandparents, my bed was situated next to the ventilation shaft that started behind their television. This meant that I would lie in the dark listening to sound-only versions of crime dramas, TV plays and many other programmes, using my own imagination to fill in the blanks. I also used my grandparents’ radio to create my own sounds by turning the tuning dial and the knobs controlling treble and bass.

Karlheinz Stockhausen was my first musical influence that I discovered on my own. I heard *INTENSITÄT* (Intensity) for ensemble from *Aus den sieben Tagen* (From the Seven Days) on the radio and was fascinated by the fact that so much could be communicated musically using essentially a hammer and nails. At the first opportunity I had, I took the train to the city (Stockholm) and searched the record shops for a copy of what I heard on the radio. I tried to describe the music as music for hammer and nails by someone I am ashamed to say, I remembered as Kalle-ketchup (the closest English equivalent would be Charlie-Ketchup). Of course I knew that this was not his real name but I could not remember anything but Kalle-ketchup. I made a complete fool of myself in almost all of Stockholm's record shops before I found someone who knew exactly what I meant. He laughed so much that he actually fell of his chair when, after a few seconds, he understood what I meant (at least that is how I remember it). He did not have the record but he wrote down the title and the composer for me, and more importantly gave me the vinyl with *Kontakte*[1] on it. I recall it as he gave it to me as a gift but I do not remember much between getting the vinyl in my hand and actually listening to it at home. *Kontakte* made such an impression on me that later in life when I was in ear training class at University, I used *Kontakte* as an aide-memoire to recognise different musical intervals instead of the more traditional nursery rhymes and children songs.

As a teenager I turned my interest to Frank Zappa, Jimi Hendrix, Pink Floyd and Hawkwind to mention a few. In my later teens I discovered The Residents, Henry Cow, Laurie Anderson, Robert Fripp and was also introduced to the Fylkingen Institute – the most important cultural centre in Swedish musical life for over 75 years - where I was directly confronted with Harry Partch and Conlon Nancarrow. Fylkingen opened a door into a new and strange world that I am still exploring. Despite this plethora of influences, what really got me going was the music of Anthony Braxton. Lenart Nilsson, a popular science author among many other things, was at the time somewhat of a mentor to me introducing me to some of the more important things in life namely – heavy drinking, philosophy, superstring theory, Go (Asian board game), Mah Jong, poker and Anthony Braxton.

What really attracted me to Anthony Braxton was that his music managed to surprise me even after listening to the same piece many times. His music had many bizarre twists and turns and sometimes made impossible leaps that made sense only in retrospect. It was like a wild and almost uncontrollable energy that was lurking under the surface erupting seemingly at random yet at the perfect time. Of particular importance to me were Braxton's duets with another of my sources for inspiration, Derek Bailey. Other musical influences from my early twenties include James Blood Uhlmer and Peter Brötzman. At this time I rediscovered Karlheinz Stockhausen, then Bartok, Webern, Bouléz and the avant garde, particularly Iannis Xenakis. Xenakis' music and thoughts have had great influence on me, especially since I was fortunate enough to be able to attend a number of his lectures and workshops. By this time I had started to work at EMS making my own music and had attended the Darmstadt summer school twice, taking classes from Brian Ferneyhough, Morton Feldman, Clarence Barlow, Helmut Lachenmann amongst others. During my early years of composing, two figures that influenced me immensely were Rolf Enström and Denis Smalley. Although both were at different periods, my composition teacher, their main influence came from actually hearing their music.

From the time I started to work at EMS, my music listening habits changed drastically. I started to attend and play at festivals, concerts and conferences around the world and have, ever since, been constantly confronted with new musical impulses. Thanks to festivals such as Perspectives, Sounds of Stockholm and the Norberg festival in Sweden, the Darmstadt festival and Transmediale in Germany I have had the opportunity to hear a lot of inspiring music from all kinds of genres such as noise, clicks'n'cuts, glitch, drone and also techno, ambient, dub-step, metal etc. Two artists of importance that I discovered through attending such festivals were Sunn O)))) and Sachiko M. Even though festivals continue to be a great source of listening to new music, now, most of my inspiration comes from working with others - which I will return to in more depth later.

1.2 Origins of a style

Self-reference, feedback, emergent properties and algorithmically mediated interaction all constitute the foundation for both my artistic work as well as my research. Feedback is, however, the lowest common denominator in my work and is where it all started, which is hardly unique since feedback is an essential part of evolution and life. Feedback has been a very useful tool for me to understand, describe and create complex behaviours and phenomena.

1.3 Feedback

My first feedback encounter happened when I was four or five years old. Using my parents' record player, according to my grandmother, I managed to create a violent and scary sound that could 'wake up the dead and at the same time kill everything living'. Apparently I turned up the volume because I could not hear anything, then remembering that I should press the start button. In doing so I created my first feedback sound. The feedback was probably initiated by the power up surge when I pressed the start button.

As a teenager I came to the conclusion that I would never learn to play the guitar but I eventually became quite good at creating various feedback effects. I was, up until then, completely unaware of the record player incident above. When my grandmother visited us she once remarked that my guitar playing did not seem to generate any music. In the discussion that followed she retold the story of the record player incident. This of course spurred me on to try and recreate what had happened which turned out not to be very difficult. I soon learned that I could control and alter the feedback by varying the stylus pressure or for instance manually rotating the turntable using intentionally scratched and in other ways modified vinyl records. The turntable setup, pictured below (see Figure 1.1), became the first step on my journey towards building an instrument for performance and composition.

Figure 1.1: My first true instrument. Note the loudspeaker lying flat on its side with the turntable sitting directly on top. This made it possible to create direct mechanical feedback, which I used alongside with acoustic feedback.

Douglas Hofstadter's book 'Gödel Escher Bach: An Eternal Golden Braid'[2] gave me a wider and deeper understanding of feedback and its kin, self-reference. The question Hofstadter tries to answer in his book is how does a self arise out of things that do not have a self? How do we go from meaningless symbols to symbols that refer to themselves? By talking about mathematics, art and music, Hofstadter presents a convincing explanation of how cognition and thinking emerge from well-hidden neurological mechanisms. The book provided me with a theoretical platform and introduced me to recursion, formal systems, isomorphism, chaos theory, logistic maps and emergent property.

1.4 Emergent Property

If feedback is the lowest common denominator in my work then emergent property is one of the goals for my artistic work. The concept of emergent property is fairly old. It was, for example, used by Aristotle in his *Metaphysics* (Book H 1045a 8-10) in which he writes that 'the totality is not, as it were, a mere heap, but the whole is something besides the parts', i.e., the whole is greater than the sum of the parts. The term emergent is, however, relatively new and was coined by psychologist G. H. Lewes, who wrote that:

> Every resultant is either a sum or a difference of the co-operant forces; their sum, when their directions are the same -- their difference, when their directions are contrary. Further, every resultant is clearly traceable in its components, because these are homogeneous and commensurable. It is otherwise with emergents, when, instead of adding measurable motion to measurable motion, or things of one kind to other individuals of their kind, there is a co-operation of things of unlike kinds. The emergent is unlike its components insofar as these are incommensurable, and it cannot be reduced to their sum or their difference.[3]

The definition has been updated over the years by, among others, Jeffrey Goldstein at the School of Business at Adelphi University. The underlying message that 'the whole is greater than the sum of the parts' still stands as a good definition and description of emergence. An obvious example would be consciousness that could be described as an emergent property of the brain.

My reason for almost never using sequencers, samples, sound files, LFOs (low frequency oscillators) etc. is that I try to construct synthesis models that have emergent properties generating the behaviours that I want to work with. Since I am trying to create instruments or sound models I obviously need to control various parameters of the models but I want the 'heart' of the sound model's behaviour to come from within and not be a result of enforced and pre-programmed behaviour.

The best way to demonstrate emergent property in my work is through an early synthesiser experiment I carried out. The first synthesiser I came in contact with was an analogue Buchla modular synthesiser and once I understood how it worked I tried to recreate my record player feedback instrument by replacing the turntable with a Buchla oscillator (see Figure 1.2). After much trial and error I ended up with something that sort of worked in the way I had intended it to. I took an oscillator and fed it straight to the studio speakers. Then a microphone picked up the signal in the room, which was then fed, via a microphone amplifier, to the oscillator's FM-input. The result was an oscillator with chaotic behaviour as an emergent property. It was possible to morph between a sine wave and white noise by varying the amplification of the signal going to the FM-input (different versions of this synthesis technique are demonstrated in the video Oscmovie and the MaxMSP patch S-O Osc demo in Appendix A2). This technique is in many ways the essence of this thesis and my work since it consists of a feedback chain and has emergent property. It can also be used as a crude algorithm for algorithmically mediated interaction. By using low frequency and a long delay (in the digital version) one input can generate two distinctively different but fully controllable outputs. Algorithmically mediated interaction is described below.

1.5 Algorithmically mediated interaction

Complex sound synthesis models have many parameters that potentially could be controlled in real-time. The problem is how can this be managed in a performance setting? In non-improvised situations, a 'score' may be of some help but it does not really address the problem of handling many parameters. Controlling a large number of parameters, particularly when improvising, can become a serious problem.

Figure 1.2 - My first sound synthesis experiment with the Buchla Modular Synthesizer.

When John Bowers and I worked together, in both an artistic (which we still do) as well as in a research capacity at KTH [Centre for User-Oriented IT-Design (CID), Royal Institute of Technology (KTH), Stockholm, Sweden] one of our main concerns was how do you manage a large number of parameters in a performance setting? Our idea was to map the outputs from, for instance, a joystick (two outputs, x, y) and map them in such a way that they could control, for example, twelve parameters of a sound synthesis model. Could a simple hardware control surface with a limited number of DOF (degrees of freedom = number of output parameters) in an intuitive way control complex sound synthesis in real-time? Our solution was to map the output from a given control surface (such as a joystick) through various algorithms thus generating a large number of control parameters in order to control the sound synthesis. Algorithmically mediated interaction has become a thread throughout our collaborations.

Chapter 2 Collaborations

2.1 Collaboration with John Bowers

2.1.1 The Lightwork performance

The *Lightwork* performance is based on virtual reality technology. Virtual reality (VR) is a term that applies to computer-simulated environments that can simulate physical presence in places in the real world, as well as in imaginary worlds. The performance consists of a number of performed visual and sonic virtual environments. The visual virtual environments were displayed on a large projection screen behind the performers and the sonic virtual environments were conveyed through a multi-speaker system surrounding the audience. Even though we were using known techniques, to our knowledge, this was the first time ever that virtual environments were generated, controlled and manipulated in real-time on stage by the performers (as opposed to performers navigating through and interacting with pre-programmed virtual environments). The *Lightwork* performance (see Figure 2.1) was our first collaboration and we were aided by the programmer Kai Mikael Jää-Aro (also at KTH).

The *Lightwork* performance is realised by two performers. One performer improvises a response to the projected virtual environment by generating, processing and spatialising sonic elements. The same performer also controls a pre-composed soundtrack by starting sections, adjusting pitch and volume. It is also possible to adjust the spatialisation of the soundtrack, but generally it is controlled by the viewpoint in the visual virtual environment. An interaction device comprising two joysticks and a pair of pressure sensitive gloves were under development during the time of the first performance but conventional MIDI-faders were used for the premiere because of reliability. The other performer interacted with the algorithms, which generate visual virtual environment content by playing an electronic wind instrument—the Yamaha WX-11—and using footswitches to trigger virtual environment modifications and select the algorithm to be used. By using a musical instrument to determine virtual environment interaction and a device usually employed for virtual environments (joysticks) to control music, we wanted to reverse conventional associations.

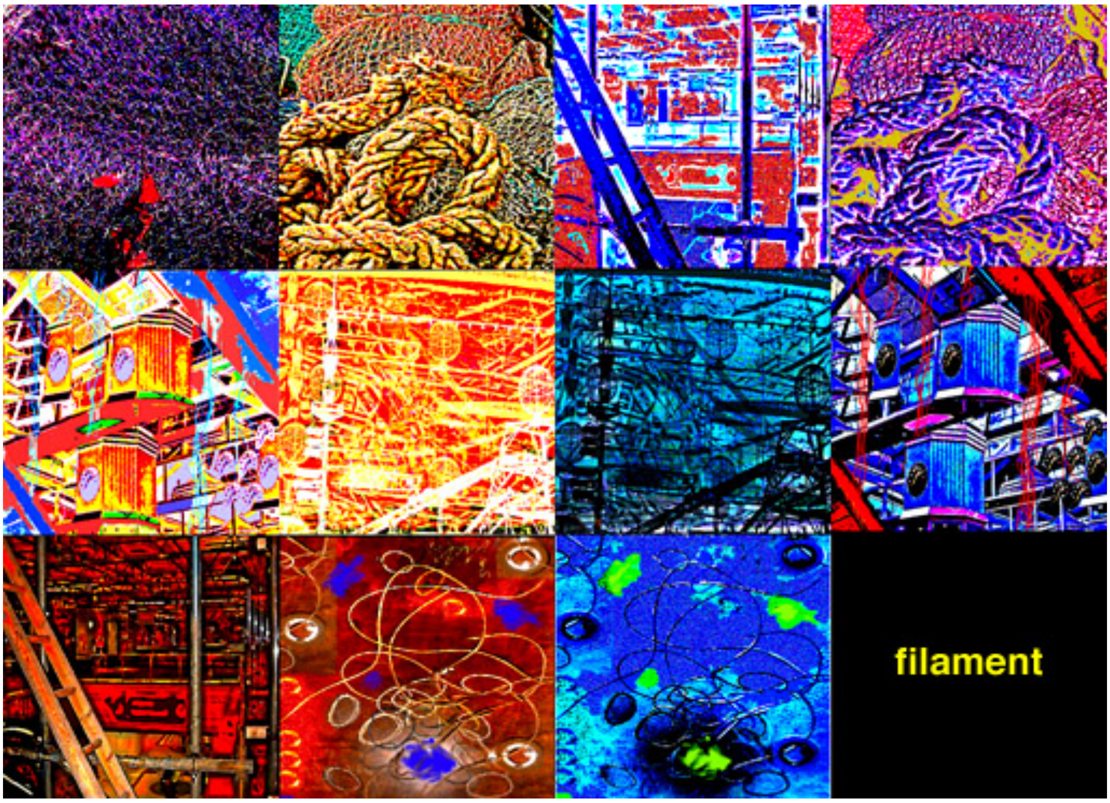

Figure 2.1 - Various Images of the *Lightwork* visual environment. (Copyright: John Bowers 1997-2013)

It is worth noting that the performers do not directly manipulate the virtual environment content. Rather, interaction is mediated through algorithms, some constructing content, others governing the creation of virtual objects, yet others control the navigation of the viewpoint in the virtual environments. The *Lightwork* performance was, for us, an important exploration of *algorithmically mediated interaction.* The first performance took place at KTH in 1997 and has been fairly well documented in research reports (see appendix paper 'The *Lightwork* Performance: Algorithmically Mediated Interaction for Virtual Environments' and published in the Proceedings of CHI98, Los Angeles, USA, New York: ACM Press.) The performed 3-d sonic virtual-environment was documented as a stereo recording and is included in the sound-example appendix A3 as LightworkPerformance.aif

2.1.2 Collaborations with John Bowers leading up to our performance at GEM days 2009

Algorithmically mediated interaction and the development of the 'ultimate' set up for performance and possibly also for composition has been a common ground in our collaborations. In many ways the foundation for this work was laid during the work with the *Lightwork* performance presented above. Since then, we have both been striving to build our own instrument or performance set-up bouncing ideas off each other. The result of this work is documented in Chapter 2.2 and in the paper eRENA Deliverable *D6.4 Tessellating Boundaries for Mixed Realities* (see the appendix with the same name) where we describe our programs *Geosonus* (Bowers) and *SO2* (mine). *SO2*, or as later versions are named, *SO-2* is a two-dimensional parameter space which you can move around in by, for example, using a trackpad or a joystick. *SO-2* is further described in the sections below about *SEQUEL* and *EQUAL*. These sections also describe my work towards creating the optimal performance environment (set-up) or as I see it building an instrument.

2.1.3 The GEMdays 2009 performance *Sine Cera* with John Bowers

In our performance at the GEMdays 2009 we had both reached quite advanced stages in our work towards a performance setup or an 'instrument' for performance. We had both, in our own way, optimized our ideas from *Lightwork* and our applications *Geosonus* and *SO-2* respectively. We both had very minimal but effective set-ups. On the surface they were very similar in that they both included a Clavia micro modular synthesiser and additional analogue equipment. However, Bowers' setup also included a computer running PD (Pure Data which is a graphical programming environment for creating audio applications very similar to MaxMSP). Both setups allowed 3-6 control inputs that were algorithmically mediated. I used two photo-resistors and three knobs on my Clavia micro modular synthesiser as a control surface. My micro modular contained both the sound engine (based on the self-modulating oscillator described in the introduction) and a simplified version of *SO-2*, which mediated the interaction (an interpreting layer between the control inputs and the sound engine). My setup also included a Moogerfoogger MF-103, a 12-stage phaser, which I used to modify and spatialise the audio. We improvised with the idea of staying outside of our comfort-zones and to encourage mistakes and spontaneous choices. Of course we had our own ideas and plans before our performance but before the event itself we agreed that as an essential part of our aesthetic for the GEMdays, we were only allowed to communicate our pre-performance ideas through our playing. The GEMdays performance was documented and is included in the submitted portfolio as *Sine Cera*.

2.2 Collaboration with Ann Rosén

2.2.1 Developing a sound based motion detection system.

In my work with Ann Rosén, feedback in the form of genetic algorithms, logistic maps, and sensor and analysis techniques has played an important part in both our sound installations and performances. In one of our first collaborative experiments we tried to create an artwork where sound was at one and the same time: 1) the audible surface, 2) the content, and 3) the theoretical basis for the underlying framework and building material. We started by trying to build an ultrasonic motion detector from the drawing below (see Figure 2.2) replacing the inaudible ultrasound with a sound well within the range of human hearing.

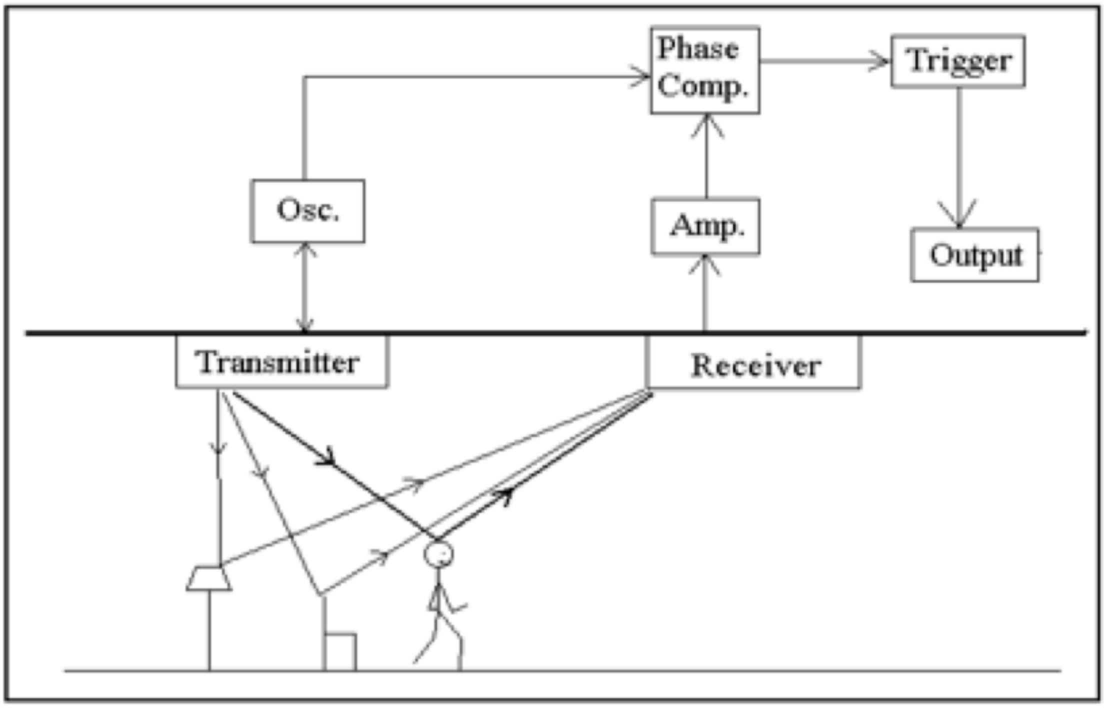

Figure 2.2 Ultrasonic Motion Detectors

A computer generated the sound, which was then fed into a room via high quality loudspeakers (transmitter). A microphone (receiver) at the other end of the room in respect to the loudspeakers picked up the sound and sent it back to the computer. The incoming sound was then subtracted from the outgoing sound and the resulting difference was used as a one of the parameters that controlled the outgoing sound itself. The result was a boring prototype that fulfilled our initial goals and it did, in a sluggish fashion, react to movement. Even if we did not get the prototype to work in a useful way without using ultrasound the prototype had three really interesting features.

The first interesting feature of the system was that the prototype had a self-stabilising behaviour. This meant that if there was the same sort of movement in the room over a period of time then the prototype adapted to that type of movement regarding it as the steady state resulting in a static sound. If the movement stopped the sound was altered until 'no movement' became the steady state once again. The second feature was a kin to memory. If there was a lot of erratic movement over a given period of time the system behaved as though the movement energy was being accumulated by the prototype causing it to continue its alteration of the output sound for some time after the movement in the room had stopped. Finally, the whole idea and setup could be described in quite an elegant way namely: The prototype is a feedback loop that starts in the virtual world where the sound is digitally generated and then past out to the 'real world'. After travelling through the real world the sound returns to the virtual world where the resulting difference in volume and phase-shifts is measured. The result is past on, and in a way added to the sound that is being generated, thus completing the feedback loop. *The sound that is heard is not merely a result of the system but is the sound of the system. The audible surface, the content driver and the technological implementation are resultant from the initial audio.*

Even though the first prototype never directly resulted in an art exhibition we kept on developing and refining the software and hardware. Eventually we got the motion detection working the way we wanted to and it became an essential part in several installation pieces such as *Noise Tank*, *Sound Inserts* and *Conveyor*.

2.2.2 *Sound Inserts*

During the summer and autumn of 2007, Ann Rosén and I were artists-in-residence at the Interactive Institute in Stockholm, Kista. Together we carried out a number of 'sound-inserts' with the help of a portable sound installation that we placed at different locations in Kista and Sätra. A sound-insert is a sound action that takes place in public spaces such as train platforms, parks, markets, entrances and squares. We wanted to study how people move under the influence of sound. What happens when the sound environment is changed? Will our movement and behaviour change according to the sound? The installation itself consisted of eight loudspeakers hanging on two mobile coat racks (see Figure 2.3) and our motion detection system described above.

Figure 2.3 - In the foreground 4 of the 8 speakers hanging from a mobile coat rack. In the background the other 4 speakers and in between the two sets of speakers a computer running the installation software and a pressure zone microphone used for our motion detection software.

The sounds used were derived from the actual sites and were triggered, altered, modified and spatialised by the movement of passers-by and visitors. We documented the installation and the reactions of passers-by and we also invited choreographers Ingo Reulecke, Moa Hanssen and Maija Hirvanen as observers to conduct their own on-site studies of how people reacted and behaved when they came in contact with the installation.

Figure 2.4 - By-passers interacting with *Sound Inserts* at one its sites.

2.2.3 *Conveyer*[4]

Our studies formed the foundation for the performance *Conveyer*. This work was a synthesis of live electroacoustic music, choreographed audience, dance and interactive soundscapes.

2.2.4 *Träffas 1983-2003*[5]

Collaborating on all levels in a totally democratic manner, agreeing, sharing and doing everything together is rarely achievable, a utopian aim in collaborative work, but when working on *Träffas 1983-2003* Ann Rosén and I did just that. The starting point was to create a sound installation that evolved over time, explored spatiality and from a sonic perspective, utilised sound objects with physical properties such as weight, shape, height, depth, width, colour etc.

We started by researching various forms of artificial intelligence in order to create an installation that was always changing and evolved. The book 'Virtual Music' by David Cope[6], which focuses on his experiments in musical intelligence (EMI) and computer music composing programs, turned out to be very important for us. The book was a not only a great source for inspiration in general but it also described SPEAC analysis[7] which looked really interesting from our point of view. We modified and implemented SPEAC in Max/MSP (see Figure 2.5) After a lot of tweaking our modified and probably misinterpreted implementation turned out to be a perfect tool for generating the type of sound progressions we wanted.

- SPEAC analysis is based on the analysis of duration, location, vertical tension etc. SPEAC identifiers are defined as follows;
- "S = *statement;* stable–a declaration of material. Statements can precede or follow any SPEAC function.
- P = *preparation;* active–an introductory gesture. Preparations precede any SPEAC functions though more typically occur prior to statements (S) and antecedents (A).
- E = *extension*; stable–a continuance of material or ideas. Extensions usually follow statements (S) but can follow any SPEAC function
- A = *antecedent*; very active–requires a subsequent consequent (C) function.
- C = *consequent*; conclusive–must be preceded directly or indirectly (with intervening extensions) by antecedents (A)."

Our implementation of SPEAC was an important building block for our installation, another was a logistic map originally intended to model the insect population in a pond:

(1) $x_{n+1} = rx_n(1-x_n)$

where:

xn is a number between zero and one, and represents the ratio of existing population to the maximum possible population at year *n*, and hence x_0 represents the initial ratio of population to the maximum population (at year 0);

r is a positive number, and represents a combined rate for reproduction and starvation.

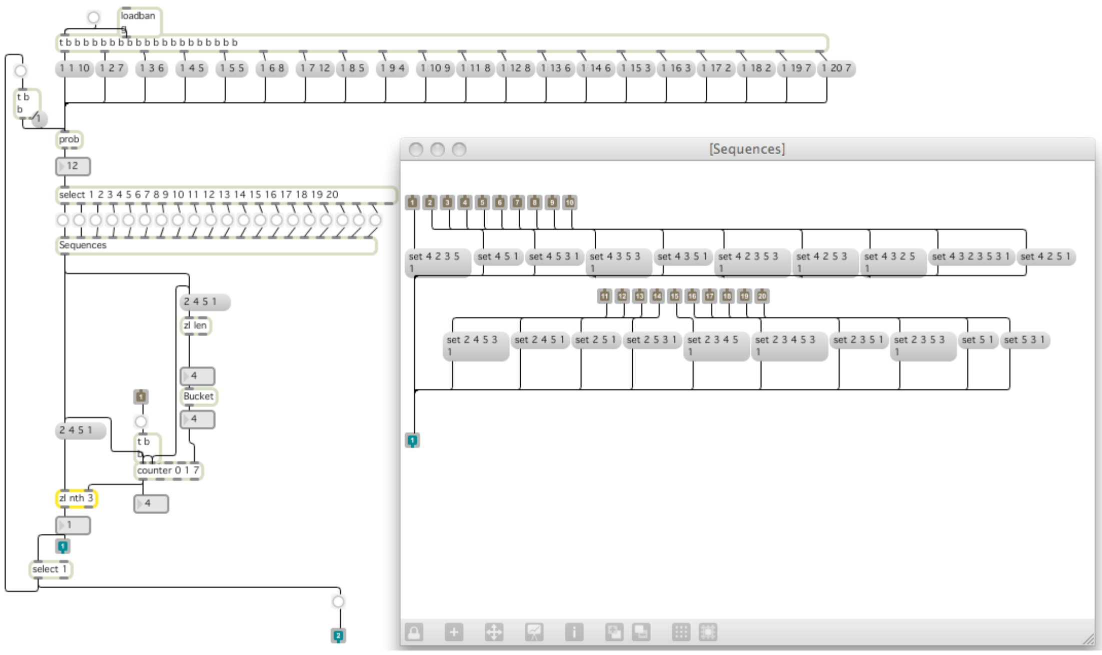

Figure 2.5: Adapted SPEAC implementation in Max/MSP

This nonlinear difference equation is intended to capture two effects:

1) reproduction - where the population will increase at a rate proportional to the current population when the population size is small;

2) starvation (density-dependent mortality) - where the growth rate will decrease at a rate proportional to the value obtained by taking the theoretical 'carrying capacity' of the environment less the current population.

Our interest in equation (1) was the chaotic behavior displayed in the graph below (see Figure 2.6).

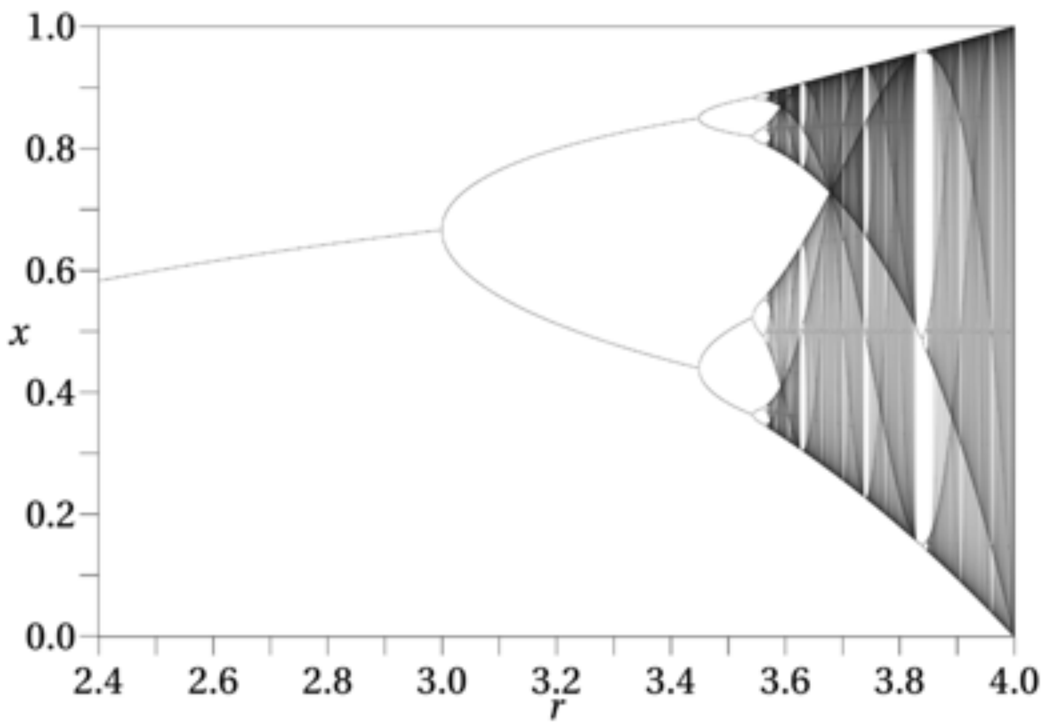

Figure 2.6: A graphical representation of equation (1) displaying chaotic behaviour. In the darker regions of the graph a minimal change of *r* will cause *x* to change dramatically.

The logistic map controlled the generation of sound events that were structured by our SPEAC implementation. At this early stage we had a working prototype with a few of the features implemented that we were interested in. From this point we started to experiment freely by altering and adding fragments of each other's software code and exchanging ideas and fragments of code with each other on the fly. During this developmental stage we had no clear result in mind. As long as the prototype kept on working, in spite of us playing Dr Frankenstein with it, we were happy to investigate the potential and limits of the system.

Working with a logistic map intended, on one level, to model a small insect society made us think about our installation as a society of noise creatures with their own personality. Our wild experimentation, misinterpretation and combination of SPEAC and logistic maps actually did in the end result in a space or virtual world inhabited by noise creatures.

A noise creature's personality consisted of the following parameters (see Figure 2.7):

- Sound colour (pink, blue, red, brown noise)

- Volume variation (static, erratic, slowly pulsating, rapidly oscillating etc.)
- Movement in the space (see volume variation)
- Life span
- Popularity as a mate (to generate future sound creatures)
- A variable determining how much of their features were going be passed on to their children

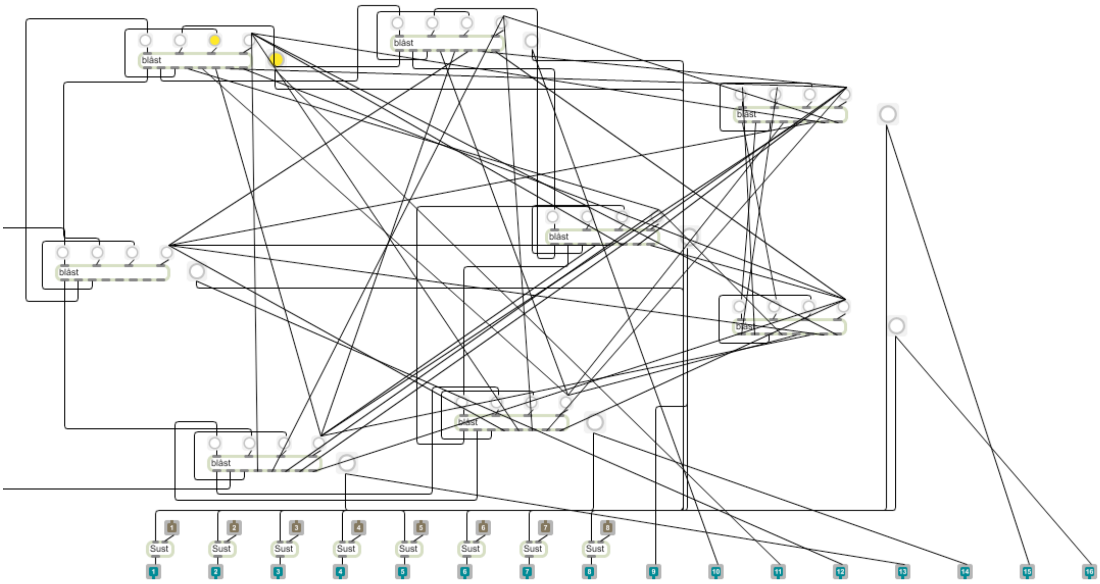

Figure. 2.7: The part of a noise creatures 'brain' that determine spatial behaviour. It is actually a map of the exhibit space with each rectangular 'blåst' box representing a loudspeaker. The six 'blåst' boxes, more or less in a circle, represent six speakers mounted in a circle with a 3m radius at normal ear level. The two 'blåst' boxes on the far right represent two speakers mounted in the ceiling.

After consulting with researchers working within the field of artificial intelligence we were informed that we had more or less developed a messy but working genetic algorithm having the following somewhat surprising features:

- A child could have any number of parents;
- Unpopular noise creatures had a longer lifespan giving a slightly enhanced chance to contribute to the gene-pool;

- If noise-creatures move together, stop in the same place at the same time more than three times in a row they were mating and giving rise to a new noise-creature;
- The noise-creatures die relatively soon after mating;
- A noise-creature learns and evolve by interacting with other noise-creatures;
- The population is small and varies between 2-6.

In the first trials of the prototype the noise-creatures interacted with visitors by responding to the visitor's movement. However, we found that when we removed the interactive part the noise-creature's movement in space became much less erratic. It was a lot easier to hear where in the room a noise-creature was and it became fairly easy to keep track of their movements. If you were quick, you could actually run after a noise-creature for a short while (which a lot of children really enjoyed).

Figure 2.8: A still from the installation

Another observation was that visitors perceived the noise-creatures to be interactive when they actually were not and vice versa - that when the noise-creatures actually were interactive the visitors did not feel that they had any contact with them. This led to a situation where we dropped the interactive part altogether. *Träffas* was documented as a stereo recording and is called *D-Dist* (it is included in my portfolio of submitted works).

Another part of Ann Rosén's and my collaborations is our joint composition work, which has resulted in the fixed media compositions *Skivad* and *Lagrad*. We also perform together under the name HR. *Skivad* and *Lagrad* and HR's performance 'Live at Dr John's' are included in appendix A3.

2.3 Collaborations with Lars Bröndum

Lars Bröndum and I formed the duo *As Good As It Gets* as a platform for developing and exchanging ideas. We started out by discussing music theory, aesthetics and philosophy. Eventually we wanted to try out our ideas practically and we started improvising and composing together. The first outcome of this work was *Dodecachordon Phase II*, which is included in appendix A3. The second outcome was *TOE,* which is included in the portfolio of submitted works.

2.3.1 *TOE* (Theory of Everything)

The Theory of Everything is a presumptive theory of theoretical physics that, if it existed, could explain and link together all existing physical phenomena. String theory and loop quantum gravity are two current candidates for the Theory of Everything title. The piece *TOE* takes its starting point in a music-theoretical discussion between Lars Bröndum and myself that has been on going for as long as we have known each other.

We wanted to explore consonance and dissonance relationships both within each timbre or time-slice as well as between different sound events or timbres separated in time (this also relates to my submitted compositions *On* and *EQUAL* where I explore vertical and horizontal time). Plomp and Levelt[8] and Sethares[9] played an important role in how we defined consonance and dissonance (see also MaxMSP patches in the Sethares folder in appendix A2). In our discussions we talked about rules, systems and structuring methods. We looked for the role that traditional harmony played in atonal music (as a non-present background or contrast and in some cases present residues) and searched for links between serialism, spectral music, and counterpoint etc. Basically we travelled around the world of music theory a few times without really getting anywhere and never even coming close to the TOE of music theory. However, we did gravitate to certain texts and quite often returned to 'The Concept of Unity in Electronic Music' by Karlheinz Stockhausen and Elaine Barkin[10], 'Formalized Music' by Iannis Xenakis[11], 'Improvisation' by Derek Bailey[12] and Wendy Carlos' writings regarding tonality, tuning and the supernatural scale[13].

Becoming frustrated by not getting anywhere, we started to meet once a week to explore different musical concepts practically, slowly building a practice-based platform for our discussions. Even though we knew we never reached the Theory of Everything of music theory we realized that, without trying, we had built a solid ground for making music and playing together. After the hands-on exploration of various ideas and concepts one at a time, *TOE* had started to transform from theory in to a set of loosely defined rules for a piece.

The first set of rules:

1. All frequencies used must be derived from the dissonance curve generated by the supplied Max patch.

2. All pulses and beats must also be derived from the same dissonance curve. This means that a base pulse is decided and all subdivisions or multiplication are done by using the ratios extrapolated from the dissonance curve.

3. Just one sound and one run of the Max patch generating the dissonance-curve was allowed.

4. There is no rule 4! (We thought we needed a rule 4 but never did but for programing reasons this remain as a placeholder).

5. The dissonance curve could only be generated just before the piece was going to be performed.

6. A fixed overall structure.

7. The piece was to consist of four parts described only by words such as anger, calm, dynamic, adaptive etc.

8. The order, length and interpretation of each part was done by each performer individually and kept secret from the other performer.

9. No samples, sound-files, prepared sounds or sequencers where allowed.

10. Avoid all comfort zones, fail rather than use favorite tricks to get out of difficult situations.

11. Avoid any ‘call-response’ behavior.

We worked for a few weeks with the above rules and used our networked laptops to share information about how we used the dissonance curve, which frequencies were to be used. A base pulse was also distributed over the network.

After a lot of experimentation and practice we found that we did not need to share information over a network Although it was useful to have the networked information as a support in the beginning, later, when we sped things up and did not stop to correct or discuss things during the process of making music, we found that we never had the time to take in and interpret all the information being sent. In addition, we had also, of course, learned what to listen for and adapt on the fly. The same thing could also be said about the rules, they were great for practicing and served as really good creative tools but they quickly became loose guidelines. When actually putting *TOE* together and rehearsing for the first performance, rules 1-3 and 10-11 became our central focus again and were more or less implemented de facto in the final work.

2.3.2 Time

We started the actual work on *TOE* by developing our beat mismatching technique where we try to create the maximum disturbance of an existing pulse or beat only using material that could be derived from rule 2 above. During the first trials we used a shared clock (part of the network described above) but we quickly abandoned that and decided on a base pulse beforehand and tried to implement rule 2 (above) by ear. This might sound harder than it actually was, but what happened was when we 'failed' we landed on the agreed 'virtual' grid of beats/sync-pulses and nothing was really added (no beat mismatching). It also had the added advantage that the result became more interesting since at least one of our internal/mind clocks drifted. Our 'good' mistakes were built into the system and added an extra dimension to the performance, whereas our 'bad' mistakes were more or less inaudible and passed unnoticed. In the beginning, this took quite a lot of rehearsal time to achieve because everything just sounded like a morass of sound and was unfocused. Eventually we pushed the beats produced into the background (most of the time inaudible for a listener) but it is still there and manifests itself in how the piece is structured in time.

2.3.3 Frequency

In my piece *On* I first created a harmonic progression or a tonal background that lasted the duration of the work. I then used this as a grid over which I placed/fitted/adjusted the actual sound material of *On* before removing this background sound file. This technique is somewhat reminiscent of the techniques used in *TOE* except that in the latter work we had one frequency grid or a set of allowed frequencies governed by rule 1. This was of course impossible to implement in actuality and we prepared a number of sounds beforehand that fulfilled the necessary requirements violating some of our predetermined rules in doing so. Having these sounds as support we could actively mismatch some aspects of each other's sound creating some very inharmonic sounding timbres that still had quasi-harmonic properties, which meant they could be used to create consonance and dissonance.

2.3.4 Overall structure

The overall structure was supposed to be derived from the sound used to generate the dissonance curve by stretching it out in time (a never implemented rule 12). This has never happened since it turned out to be much more fun to construct it manually before each concert taking into account the different acoustics, type of audience and event at each particular venue.

2.3.5 Gestures

The gestures used in *TOE* are completely governed by a palette derived from our rehearsals. We never actually discussed this process in detail as a gestural language emerged from the beginning of the process and developed throughout successive rehearsals. We did, however, decide to (if possible) actively avoid clichés such as call and response, anacrusis etc. in order to explore the gestural palette in more considered ways.

2.4 Collaborations with Paulina Sundin

Paulina Sundin and I had both worked at EMS Stockholm for a considerable period and discovered that we shared an interest in microtonality, particularly the works of William A. Sethares and his ideas about the relationship between timbre and scale. In his book 'Tuning, Timbre, Spectrum, Scale'[14], Sethares presents a method to analyse a sound and use the analysis to derive a scale. This scale should, in theory, work with the analysed sound in much the same way as a just major scale would work together with a sound whose timbre is made up of harmonic partials. I implemented Sethares' methods in MaxMSP and Paulina and I used this resource in our respective work. We also started to work on a composition together which resulted in the not yet finished piece *The Harmony of Found Objects*.

2.4.1 *The Harmony of Found Objects*

This work is a site-specific composition that is composed live on stage by two performers that communicate by sending bits of code, instructions and analysis to each other via networked laptops. Live coding, real-time analysis and real-time synthesis are the basic building blocks in *The Harmony of Found Objects*. The piece starts by bringing site-specific found objects to sound (by for instance striking them or using a bow). The resulting sounds are fed into a matrix of delay lines that constitute the heart of the sound synthesis. The initial sounds are also analyzed using algorithms developed by William A. Sethares and the result is used as a basis for the on-the–fly composition and the real time sound synthesis. By using the actual concert room as one of the main delay lines, in the matrix, almost every action taken by the performers has a direct audible result, which means that *The Harmony of Found Objects* could be seen as a sonification of the creation process of *The Harmony of Found Objects* as well as being the actual piece. The MaxMSP Patches developed for this project are included in appendix A2 in the Sethares folder. Although the work has yet to be completed, the resulting research and MaxMSP patch has been utilized in my own compositions.

Chapter 3 Individual Works

3.1 *SEQUEL*

3.1.1 Physical Modelling Concepts

In my compositional practice I have never consciously adopted an aesthetic that has a purist approach, one that pushes things to the limit or one that refines elements beyond recognition as its goal. Nevertheless, the aforementioned approaches are important in my work. This is best illustrated in my work *SEQUEL*. I borrowed the term ‘sequel’ from the film world where it represents a movie that in one way or another is a continuation of an earlier movie. The structure of *SEQUEL* was derived from two main elements: a) the analysis of a number of movies, focusing on the story line and dramatic episodes, and b) investigating how sound was used to provide a sense of transition between scenes as well as gluing the story together. *SEQUEL* was composed by first making a skeletal structure. This provided a framework of the piece. Using my implementation of physical modelling, I made the sounds and gestures for the framework. The framework was then ‘fleshed out’ with sounds that I had recorded, processed, shaped, and then mixed together in the studio at EMS, Stockholm.

Physical modelling (PM) is a synthesis method where, instead of trying to imitate instruments or natural occurring sounds one tries to replicate the way the sounds are naturally generated. This is normally done with computers due to the computation involved though there are analogue and hybrid examples such as the first incarnations of the Musse synthesiser (see Figure 3.1) built at KTH in the mid 1970s and used to generate a singing human voice. Generally the term PM is used when physical phenomena are replicated in computer software (simulation of turbulence in air or water for example).

In *SEQUEL* I aimed to use digital waveguide synthesis, which is a type of PM where digital delays are used to emulate the bore of an instrument, for example a saxophone or a clarinet i.e. the space, and media through which sound waves propagate.

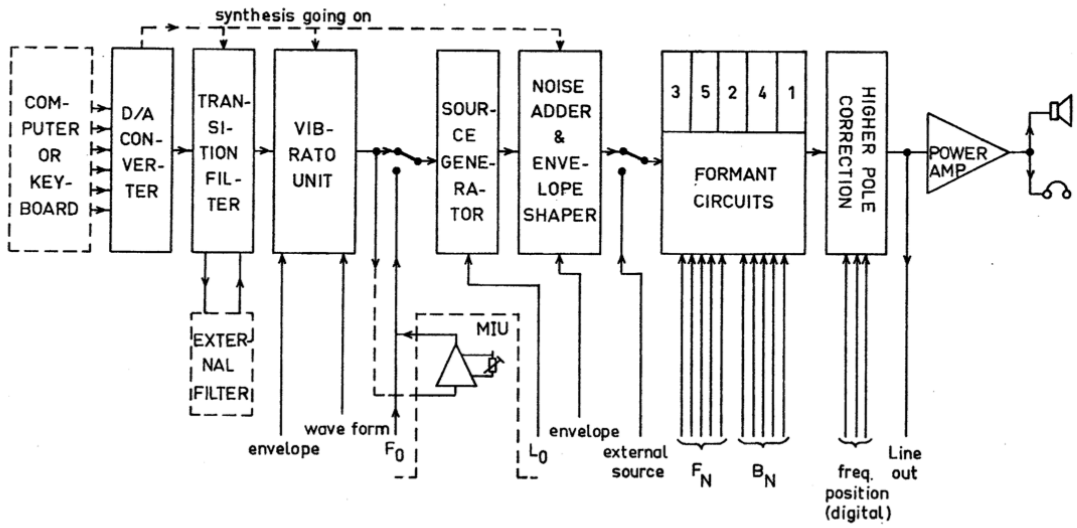

Figure 3.1 Block diagram of Musse.

Digital waveguide synthesis can be seen as a number of delays hooked up in an attenuated feedback loop. By injecting noise into the feedback loop it is excited and resonates producing a pitched sound. A typical waveguide model of a saxophone consists of an excitation part that emulates air pressed through the mouthpiece exciting the reed (causing the reed to vibrate). White noise can be used for simple and rudimentary excitation. The sound generated by this excitation part is fed through the first of four delay lines. The first delay line emulates the upper half of the saxophone bore where sound waves propagate outwards (towards the bell). Interaction between outwards propagating sound waves and the sound waves reflected by the bell (inwards propagating sound waves) occur in a scattering junction. A scattering junction is basically a combination of a mixer and extremely short delays (causing phase shifts). After the scattering junction the sound waves propagate through a second delay emulating the lower half of the saxophone bore after which the waves reach a lo-pass filter emulating the bell of the saxophone. A part of the sound exits the saxophone model resulting in audible sound, the other part is phase inverted and fed through a third delay. This delay is modelling the lower half of the saxophone bore through which the sound that is reflected back by the bell propagates. After passing through the scattering junction, the reflected sound goes through a fourth delay (the first half of the saxophone bore). The reflected sound is then once again phase inverted and mixed together with the sound from the excitation model. This is done to emulate the reflection occurring at the mouthpiece of a saxophone.

In trying to create my own digital waveguide model I was influenced by a number of key articles[15]. I spent a lot of time assimilating the principles of PM and made many attempts to implement the examples contained in the above articles. However, in the end I constructed a model that was my own simplified description of the above. Another big difference from the methods outlined in the articles was that I used hardware delays, mixers, filters and a computer running Max (pre MaxMSP) controlling the delays. The setup consisted of:

- an Ensoniq DP4 for the bore parts of the model (four delay lines and some phase shifting);
- a Buchla 200 system (analogue modular system);
- an excitation model consisting of two Programmable Complex Waveform Generator Model 259 frequency modulating each other;
- a scattering junction consisting of two Frequency Shifter/Balanced Modulator Model 285, two Mixer/Preamplifier Model 207, one of the below mentioned Digitech effect units and a studio inline mixing desk;
- a Bell model consisting of a Quad Voltage-Controlled Lo-pass Gate Model 292 and a Programmable Spectral Processor Model 296;
- two Digitech DSP 128 multi-effect units were used, one in combination with the above mentioned frequency shifters and mixers for the scattering junction, and the other was used for modelling the reflection occurring at the mouthpiece.
- a 24 channel automated inline mixing desk for the above mentioned scattering junction and also connecting all the above mentioned equipment together;
- a MacPlus computer running Max controlling the Ensoniq DP4 and the two Digitech DSP 128
- a dedicated computer controlling the mixer;
- a 24-track tape recorder used to store the control signals for the mixer.

3.1.2 - Assembling *SEQUEL*: The foundation and framework.

Due to this very awkward and complex setup, I had to complete each part of the piece in one go as even the slightest change to any of the parameters could result in a completely different sound. Even changing a cable could completely alter the sound. Prior to the actual composition of the work, I created a pool of material. I created approximately seventy sounds between 1-4 seconds; about one hundred sounds and sonic (music) gestures with a duration of 5-30 seconds; and finally twenty atmospheric and ambient sonic backgrounds (some of them were mixed together resulting in more foreground material). These sounds were then mixed and assembled together into a sonic foundation and framework.

3.1.3 – 'Fleshing Out' the Framework.

Right from the start of my work on *SEQUEL* I started to carry field-recording equipment in my backpack (4 kilos including extra batteries) wherever I went. This made it possible for me to record sounds outside of the studio as I encountered them - for instance straight after watching a movie or on my way home after a framework studio session. Once I had gathered a number of acoustic sounds I carefully mounted them on the framework (i.e. mixed them together with the framework). The main bulk of the sounds I recorded outside the studio could be divided into two categories: 1) spaces and environments that essentially were 'silent' but reacted to existing sounds by resonating, reverberating and reflecting them; 2) acoustic sounds that were fairly 'dry' i.e. sounds recorded in non-reverberant environments which made them fairly easy to fit together with the framework. The ambience in *SEQUEL* is either constructed in the studio using artificial reverb or assembled using the recorded environments as building blocks.

As the piece started to come together big gaps in the framework suddenly became apparent. This meant that I would have to book three studios at EMS again, take the equipment from two of the studios and cram everything into the third, transport all of my gear into EMS in order to finish the framework. I was completely overwhelmed and beside the fact that there were other composers who wanted to use the studios, I could never go through the same process again especially since this time around, I knew what I had ahead of me. I took a break from EMS and revisited the University of East Anglia to do a presentation and a concert. After the concert I came across the magazine *Sound on Sound* where Turnkey advertised that they slashed the price on a fantastic synthesiser - a monophonic digital synthesiser that even after the 70% price cut still cost as much as any polyphonic rack-mountable synthesiser-workstation. This really peaked my interest, a digital monophonic synthesiser, and an anomaly in itself, in this price range (not to mention its price tag the previous year) must be something extraordinary. I had found what I then perceived to be my holy grail, the Yamaha VL-1m, which was a digital waveguide model of a saxophone in a very practical and small box (even though today I regard it as quite heavy, extremely bulky and something I would never consider carrying anywhere outside my studio).

3.1.4 - Experimenting with the Yamaha VL-1m

The Yamaha VL-1m has a very open structure allowing the user to access an overwhelming number of parameters via MIDI. The downside is of course that if you start experimenting there are a proportionately small number of parameter settings that do not make any sound at all. In some cases there are parameters that have to vary constantly or else the VL-1m does not produce any sound. To address this I built a software controller in Max called SO-2 (see appendix eRENA *Deliverable D6.4 Tessellating Boundaries for Mixed Realities* chapter 2.2).

3.1.5 - SO-2

SO-2 is a two dimensional parameter space in which you can move around in by moving the cursor. The version used in *SEQUEL* and later in *EQUAL* can handle up to 40 parameters. With SO-2 connected to the VL-1m one could experiment with the 40 parameters by either setting a static value, choose between different LFOs (low frequency oscillators), use a simple logistic map or use a random walk function. Once you were satisfied with a sound, the parameter settings or behaviours could be stored and then allow you to develop a new sound in the same way. When you had four sets of parameter behaviours (or static settings) you could allocate each set to one of the corners of the computer screen. SO-2 would interpolate between the four sets of parameter behaviours when the cursor was moved across the screen letting the user navigate in a two dimensional parameter/behaviour-space.

3.1.6 - Finishing *SEQUEL*

The *SEQUEL* framework was finished using SO-2 and the VL-1m. With the VL-1m editor I constructed a giant 6m long saxophone that had an ultra-thin reed. With the right virtual air pressure (one parameter) in the mouthpiece (also of course virtual) the ultra-thin virtual reed would start to behave chaotically (which an actual reed would also do if it is thin and blown hard). This meant that I could create the low frequency and slightly chaotic sounds I needed to finish the framework for *SEQUEL*. I also used the VL-1m together with SO-2 to create the missing sounds to 'flesh out' the framework instead of using recorded sounds.

3.2 *PREQUEL*

Making *SEQUEL* was an highly intellectual process requiring careful planning, following ideas through to the end, doing things according to the score/script, learning new synthesis methods, building a software controller etc. The process of making *PREQUEL* was almost the opposite - it was all about doing, spontaneous composing and solving problems on the fly. This was, however, only possible because of all the work behind *SEQUEL* which gave me a set of tools, sounds and musical ideas etc. It was as if the work with *SEQUEL* had created a giant wave and working with *PREQUEL* was only a matter of surfing on this wave. 'Prequel' is again a term from the film world and it refers to a movie that is a follow up to an earlier movie but the story takes place in a time before the one in the earlier movie. In *PREQUEL* I harvested methods, tools and ideas that grew out of the work with *SEQUEL*.

In spirit, however, I returned to an idea for a composition I had had a long time before any ideas regarding *SEQUEL* had entered my mind. The idea was to base a piece purely on having fun and spontaneity. I wanted to follow the sounds wherever they would take me and at in between working on the piece I wanted to visit the Stockholm Film Festival as much as possible. So the underlying concept for *PREQUEL* was to make a piece by just playing around and to go to the film festival during my breaks (or maybe the other way around). *PREQUEL* is in fact the only piece I have composed where I followed my plan from the start to the end.

The sounds in *PREQUEL* consisted of the results from playful experiments with the Yamaha VL-1m and sounds that I had previously recorded. I did not tailor make or record any extra sounds for *PREQUEL* and I did not process any of the recorded sounds. The compositional process was one of generating sounds then composing a part that suited those sounds and so on in an iterative manner. The exception was that I did use a few of the processed sounds that were left over from *SEQUEL*. Since *PREQUEL* was made during the ten days of the film festival, all the different parts where created in the same spirit and therefore automatically related. The assembling of the materials was therefore not particularly onerous as the different parts fitted together more or less effortlessly.

3.3 *EQUAL*

EQUAL is the synthesis of *SEQUEL* and *PREQUEL* - where *SEQUEL* is the thesis and *PREQUEL* the antithesis. The three pieces all relate to film in different ways and in *EQUAL* the use and displaying of time in movies is an important source of inspiration. When I started work on the composition of *EQUAL* I realised that this way of looking at time has been a part of my work from the start. Following this realisation, I read through all of my old notebooks (the oldest was from 1978) and found that the thinking behind *ON*, composed in 1990, could be expanded and implemented in *EQUAL*.

The aesthetic goal in *ON* was to utilise the interplay with what I call horizontal and vertical time. The idea behind vertical time can be explained as the musical equivalent of a sculpture. Everything is present all the time, it is up to the visitor to chose where to focus, from which side and in which order to examine the details. An example could be a chord progression where the chords are stacked on top of each other separated by octaves instead of time. The listener could then, in theory, hear the progression by focusing on the different octaves in a predetermined order. In reality this is, of course, very hard or even impossible, but it gives an idea of what I mean by vertical time - that everything happens at the same instance and that the separation of different events depends on our perception. Horizontal time is what we normally mean by time when listening to music and other time-based media – the progression of time passing.

After the first performance of *ON* I was made aware of Jonathan Kramer's book 'The Time of Music: new meanings, new temporalities, new listening strategies'[16]. I was intrigued by the fact that to my understanding, his concept, which he came up with before me, was identical to mine - or to be more precise, I could not pinpoint any clear differences. The only thing I can point to is that my use of vertical time might have more of an extra musical perspective. However, when I started to make *EQUAL* I found that Kramer's book was an excellent companion to my own notes and helped me to refine my ideas.

In chapter 2 of 'The Time of Music' Kramer starts to define linearity and non-linearity and towards the end of the chapter he states,

> Phrases have, until recently, pervaded all Western music, even multiply-directed and moment forms: phrases are the final remnant of linearity. But some new works show that phrase structure is not necessarily a component of music. The result is a single present stretched out into an enormous duration, a potentially infinite "now" that nonetheless feels like an instant. In music without phrases, without temporal articulation, with tonal consistency, whatever structure is in the music exists between simultaneous layers of sound, not between successive gestures. Thus, I call the time sense invoked by such music "vertical." (Kramer, 1988: 55)

Throughout his book Kramer refines and updates his 'time definitions' which could be summarized as (compare with a summary posted by Phillip Henderson on his blog http://verticaltemporality.blogspot.se/2008/12/review-time-of-music-new-meanings-new.html)

> **Absolute time**: There is a linear procession of moments.
>
> **Clock time**: The linear procession of moments as described by the reading of a clock.
>
> **Gestural time**: The significant gestures within a piece that determine the continuity.

Goal-directed time: The musical events move toward foreseeable conclusions.

Non-directed time: The musical events move toward unforeseeable conclusions.

Moment-time: A point in time, a moment that is defined as a 'self-contained (quasi-) independent section, set off from other sections by discontinuities' (Kramer, 1988: 453). My own use of moment time refers to the perceived time within a moment (see vertical and virtual time).

Multiply-directed time: The musical movements occur in multiple directions.

Social time: The system of time organisation reliant on timetables and plans.

Vertical time: There is no differentiation between past, present and future. There may be no separate events.

Virtual time: The subjective time felt by a person absorbed in listening to a piece of music.

For my purposes I divided everything in the above list into two groups: Horizontal time and Vertical time. Horizontal time consisted of Absolute time, Clock time, Gestural time, Goal-directed time and Social time. Vertical time consisted of Non-directed time, Moment-time and Virtual time. Multiply-directed time could belong to either or both groups depending on the situation, but I never used the term in my notes or in the score for *EQUAL* since, from my point of view it was better to use a combination of the other terms to describe different types of Multiply-directed time.

A book that is related to these thoughts and somewhat bridges the gap between film and music is 'Wagner and Cinema' edited by Jeongwon Joe and Sander L. Gilman[17]. This book has, however, not influenced the composition of *EQUAL* as such, but has been of help in structuring my thoughts when writing, as well as influencing my work with Lars Bröndum.

Contrast and lack of contrast has an important role in *EQUAL*. By using synthesised sounds and acoustic spaces I intended to generate very diverse sounds and make them appear as if they were almost the same by using spatialisation, reverb, and mixing. I also tried to make sounds that were almost impossible to separate from each other, appear as diverse as possible using the same techniques. I had an image of sound objects living in different realms or spaces being able to interact with each other because of a kinship that reaches over borders or diverse sound objects interacting with each other because of close proximity in space. My goal was to create a dualism between a completely upside-down, inside-out, chaotic sonic world and a second completely normal sound world that actually made no sense at all.

In the making of *EQUAL* I wanted to avoid my own compositional clichés and push myself to come up with new solutions to all the problems that emerged when composing. The structure for *EQUAL* can be described as a string of eleven blocks consisting of vertical time interconnected with lines of horizontal time. One of the ideas behind the vertical time blocks is that they are static versions of the whole piece. Inside a vertical time block everything slows down, freezes, the piece stops to move giving the listener the opportunity to explore the piece in their own time. The lines of horizontal time generate the forward motion in the piece, moving the listener between vertical time blocks. They resemble story lines during which musical ideas could evolve, events unfold, develop, and die out. I intended to take the listener for a ride along the horizontal time lines. When arriving at a vertical time block I would let them explore, rest and contemplate on their own for a while. Then the journey would continue along the horizontal time lines until the next block of vertical time and so on.

The working process is always important for me and I normally first make some sort of score and a plan for the work process. After a few sessions in the studio, the working process slowly changes leading to a revision of the score. In preparing *EQUAL*, the score and the plan for the work process merged together into a sort of manual or set of instructions for creating *EQUAL*.

The instructions specified the following work process; first a session in the studio, improvising using my software control surface *SO-2* and the Yamaha VL-1m synthesiser focusing on vertical time (see the section about *SEQUAL* for more information about *SO-2* and the VL-1m). In the next stage the results were to be analysed leading to another studio session. In this session I improvised using the same tools and materials as in the previous studio session but now focusing on the development of these materials in horizontal time. In the following session the results from the previous sessions were to be fitted together by creating sounds that could bridge the gap between the materials and 'glue' them together or alternatively cut up the material in segments and then splice the segments together always using segments from alternating sessions (first a segment from the improvised vertical time session then a segment from the horizontal time session and so on).

In the actual working process, one studio session could consist of anything from a quarter of an actual stage up to 3 or even 4 whole stages. On some occasions I also returned and revised material from previous sessions. Even though I thought that my original instructions incorporated enough space for revision and adjustment it turned out that in order to follow the spirit of my instructions I had to do a lot of work that was not following the letter of those same instructions. The transitions that were to glue or splice together horizontal and vertical time sessions were the most problematic and often led to a number of 'in between sessions' where I alternated between improvising and gluing and also edited the improvised material before the splicing and gluing.

The last step in the instructions for making *EQUAL* was to remove excessive material without altering or in any way changing the remaining material. The idea was to carefully uncover what I hoped would be, the hidden gem beneath. To put it another way, having created a somewhat organic block of material, I intended to sculpt out a piece following the grain of the sound and utilising the structure of the material. Having adhered to the process described above the final 'uncovering' stage was realised on a 24-track Dyaxis system (similar to Pro-tools). I only needed two studio sessions to remove and filter out any unwanted material.

Throughout the composition process, I used sounds created with the Yamaha VL-1m physical modelling synthesiser and *SO-2*, my control surface for the VL-1m programmed in MaxMSP. The only exception to this and also where I deviated from the instructions for the final stage was in my use of recorded acoustic sounds that were left over from *SEQUEL* on top of what became *EQUAL*. These sounds were, according to my notes, someone chewing a sweet, a drop of water hitting a water surface and a soprano voice. There were a few more but they are hardly noticeable and in my notes they are marked as acoustic 'spice' sounds.

3.4 *FLUGAN II*

From a chronological point of view *Flugan II* should be placed first among my individual pieces but from an aesthetic and conceptual point of view it is somewhat related to *SEQUEL* and *EQUAL* and quite close to *PREQUEL*. *Flugan II* shares some sound material with *EQUAL*. My work with the piece started with the creation of a sonic framework working along the lines of *SEQUEL* and then I 'fleshed out' the framework by adding improvisations on the Buchla synthesiser. As I was working in the studio I noticed a big housefly zooming around and irritating me but I tried to ignore it and carry on working. After a while the fly disappeared and I was preparing to add one of the final tracks on a 24-channel 2-inch tape. When I started the tape recorder it sounded like the pulling out of adhesive tape from a big roll. It was a horrible crackling sound and my first thought was that the back-coating of the tape was coming loose, a common problem at the time, but upon closer examination it turned out that it was the fly that had completely ruined the tape. It was actually now a very thin 235 meters long and 2 inch wide fly (I still have the tape which has been exhibited at an art gallery). Since I had a deadline, I quickly assembled the piece with little prior planning or the production of a pre-compositional score. Like *PREQUEL*, I worked swiftly and spontaneously allowing my intuition and ear to guide the sonic result. I was very pleased with the result and it is for many reasons one of my favorites.

3.5 *JUST*

There were many diverse musical and non-musical ideas that came together in the creation of *JUST*. The common thread linking them is about me challenging myself. The first idea was to create a challenge and an environment that could best be described as juggling with too many objects i.e. a task that is essentially impossible but with total focus you can manage it for a short time. The reason for this is that I wanted to create a situation that is so demanding that you will have to give it everything - you are pushed to the edge. It is like trying to surf on the lip of a wave (the crest or the edge of the wave). The lip will not carry your weight but if you have enough speed you can surf up the face of the wave to the lip where you hang weightless for a few moments before going down the face again (commonly termed - 'Hit the lip' or 'Off the lip'). *JUST* is all about hanging on to these few moments as long as possible.

I wanted *JUST* to be all about performance, something that could only be done on stage with my adrenaline levels peaking. After a lot of experimentation and trials I created a network of delay-lines and filters with a large number of feedback possibilities. Basically, it was a large feedback loop consisting of a lot of smaller feedback loops. None of the loops could self-oscillate together or individually without adding external energy but if a short noise burst was injected into the network the different loops could be brought to oscillation. This meant that I could mix in or out the different oscillating loops as well as retune them by altering the delay time on the delays or the cut-off/resonance frequency on the filters. If at least one of the loops were oscillating, it could be fed into another loop which, if tuned correctly, could also be brought back to life (oscillation). After a lot of practice I could keep the energy level high enough in the network, outputting sounds for approximately 6-7 minutes after the initial noise burst. It was of course possible to set up a fairly static sound by always having at least one or two loops sounding, carefully adjusting the parameters and thereby keeping the network sounding for as long as you like. The whole point of the work however, was to vary the dynamics, to use gestures and pauses. This meant that the procedure would be the following:

1. Inject a noise burst into the delay and filter network, which will automatically bring at least one or probably several of the loops into oscillation. You will then be at stage 3 or 4 in this list;
2. Direct residual energy to one loop and bring it to oscillation;
3. Use the energy created by the oscillation to get more loops oscillating;
4. Start shaping the sound by varying the volume, altering the pitch mixing the sounding loops. Eventually the energy starts to drop and you have to start at 2 again. If the energy drops too much you will not be able to bring any loops into oscillation which means that the piece has ended.

If you have lots of energy i.e. many or very loud oscillating loops you can feed the resulting sound to dead/non-oscillating loops and thereby reviving them (bringing them into oscillation). If you have little energy you have to tune the receiving loop (by adjusting the delay time or cut-off frequency) to match the frequency of the incoming sound in order to make the receiving loop active (oscillating) again. There is a special case if you have too little energy to start an oscillation but you still have some residual energy left you can jumpstart an oscillation by altering as many parameters as quickly as you can except for the ones controlling the loop you are trying to jumpstart. This process causes zipper noise (due to that the relative low resolution of the control-parameters) and this in combination with the residual energy can in some cases be enough to jumpstart a loop. The zipper noise is, however, not enough on its own to bring any loops back to life. Another important factor is that the more energy you have the longer pauses you can make. To me, performing *JUST* is a sort of chicken race where you challenge silence. I try to make the most of what momentum I have by, for example, creating a phrase or a gesture perhaps followed by a short pause, a moment of near silence before I have to gain momentum again. A long phrase or gesture will drain momentum and a pause even more so. If there is no momentum left one cannot do anything and the piece has ended.

The second idea underlying *JUST* was to create a minimal performance setup, an instrument that implemented as many of my ideas as possible regarding playability, ergonomics, quick response and was most of all designed to utilize my personal performance skills. The first setup consisted of a small laptop (a note-book computer) running the editor software for the micro modular giving me onscreen access to all the sound parameters. In combination with the four potentiometers and three push-buttons on the micro-modular this set-up proved to be a good bench test. However, it was slow and tricky to use and did not feel at all like an instrument. Eventually, I arrived at a design that felt very comfortable, fast to work with, and something that felt like a performance instrument. It was a hardware control surface that consisted of four joysticks, one crossfader, and a pressure-sensor and modulation wheel all generating MIDI. I also equipped it with my own version of Michel Waisvisz's Cracklebox. The Kraakdoos (or Cracklebox) is a small box with six metal contacts on top, which, when pressed by fingers will generate electronic tones, sounds and all kinds of noise effects. The human body becomes an integrated part of the sound-generating circuit and factors such as pressure; muscle tension and which contacts you press determine the range of sounds possible. The reason for a cracklebox was that when I was finished with my control surface where everything was placed for optimal ergonomics and playability there was lot of unused space. I thought that it would be the perfect place to implement my cracklebox that I had been working on. This turned out really well and even if I mainly use it for performing other pieces and improvisations I use it to generate the initial noise-burst in *JUST*. My version of the cracklebox is actually really close to Michel Waisvisz's original design, but I have larger capacitors and some extra transistors, which in the end did not make a big difference. My goal was to be able to generate more low frequency sounds than was possible with the original model which I had some limited success in achieving. The most important difference was that I had an extra metal contact, which I added to the circuit. This extra metal contact probably changed the sound more than my carefully thought out modifications.

3.6 *Ambient Edge*

My collaborations with Ann Rosén has led to many spin-off projects and one of the more spectacular sound-installations we did was *Spatial Silences*[18] - exhibited at the Royal Academy of Fine Arts in 2005. In a large room filled with noise we created a silent sphere (using active noise cancelation techniques), with a diameter of approximately 1.5 meters that floated around in the room. The visitors could either go around trying to find it or just stand in the room and suddenly be engulfed by the sphere of silence. If this happened it was quite spectacular with most people experiencing it as a very weird and strange sensation. The idea of silence and silent sounds stayed with me for a while and eventually I thought of completely reversing the process of *Spatial Silences*. I decided to use inaudible sounds to create a piece so that the sounds that you hear when listening to the piece are actually not the sounds on a hard drive or sounds coming from a CD or soundcard, but the interference effects from these sounds.

The structure of the piece grew out of my work in the studio (which in many way is my normal way of working). Many of the sounds I created this way, however amazing they sounded in the studio did not work when they where recorded and then replayed. To counteract this I used these sounds as parameters to control other sounds and then mixed them together. An example of this was a slowly pulsating low-frequency sound that hovered around in the studio but when recorded and played back it just had lost all its character and was completely useless. Instead of removing this sound I animated it using a recording of water being stirred with a spoon in a glass. I synced the stirring to the pulsation and used convolution to map the volume curve and frequency content to the water sound. Then I mixed the original, the processed and the synced but unprocessed sounds together and thereby I managed to capture the spirit of the original sound in order to us it in *Ambient Edge*.

When I think of the process of making *Ambient Edge* I remember it as a very simple piece to make and I think of it as coming together in three weeks from start to finish. When I re-read my notebooks however, I realized that the idea had in actual fact been gestating for a considerable time. The only difficulty in the composition process was reanimating the sounds I really liked but could not use straight from the recording sessions (one such instance is described above). The ending, or the last 'silent' 30 seconds of *Ambient Edge* are intentional as I wanted to point out that the ending was not originally silent but consisted of sounds that were almost inaudible to the human ear, however still detectable, and one could sense that there was something present in the studio. These sounds were actually the result of inaudible sounds creating an audible interference in the room but I did not manage to record these sounds with a microphone since ventilation noise and computer fans masked them. In addition, the 'silent' sounds themselves were completely out of range for normal recording techniques. Even if conceptually the work necessitated sounds being 're-animated', the silent ending functions as a tribute to my original idea. I do not mind if the ending is left out in concerts but for me it is important that it is still there.

3.7 *Spiral* a continuation of *Ambient Edge*

The first time I worked in the SPIRAL studio at Huddersfield University I revived my re-animation ideas and some of the material from *Ambient Edge.* The SPIRAL studio has 24 speakers arranged in three equal sized rings above each other with each ring consisting of eight speakers. This speaker setup made it possible for me to listen to several parallel audio streams of material at the same time without them masking each other. The SPRIAL studio made it possible for me to move sounds both in space and time. By using a multi-track sound editing software I created a 24-track set-up assigning each track to an individual and unique speaker. I would then loop the 24 tracks for about three minutes. This meant I could take a sound and place it in any of the 24 speakers and also move it in time during the three-minute loop. After having added a number of sounds to the loop I then changed the loop length and carried on in the same fashion. When I eventually ended up with a whole piece (*Spiral*) I made a stereo mix for portability.

3.8 *In My Own Words*

In one of my earliest pieces, *Off the Lip* (1987), I was inspired by surfing and in particular the concept of manoeuvring the surf board on the edge of a breaking wave (the lip). This is a concept that has stayed with me and in my music. I continue to explore breaking points, event horizons, metamorphosis and transition points in time and space. Related to this concept is the on-going quest I have musically to investigate the relationship between the very large and the very small alongside the very slow and the very fast. I want to explore the relationship between parallel co-existing events that are only separated by size or scale. One such example that is very important to me is the relationship between sound synthesis and composition or the smallest sine wave present in the piece and the whole composition itself.

In My Own Words is fundamentally about two things: the first is the transitions between different states such as solid and liquid; the second is the relations between the atoms in melting ice cubes and the ice-cubes themselves. I started with the idea of six ice cubes on a flat warm metal surface. The water gradually undergoes a transition from solid (ice) to liquid (water) and finally to gas (steam). The ice-cubes move around in a certain fashion as they melt, the water molecules within each ice cube have another movement pattern and the atoms yet another and the different movement patterns interact with each other even though they are separated by their size.

Back in 1998 I worked together with Aatto Laaksonen, Professor in Physical Chemistry, Stockholm University and Kai-Mikael Jää-Aro research engineer, KTH (Royal Institute of Technology in Stockholm) at KTH on visualising and sonifying water molecules behaviour in different conditions.[19] Based on this work and a data set (provided by Aatto Laaksonens) containing computer-generated data for water particles at different temperatures, I made a very simple computer model of a microscopic cube of water (0.00000005 ml or 1670000000000000 molecules). I combined 27 instances of this cube into a larger cube. This gave me a rough computer-simulated cube of water that was large enough for me to use as a base for my composition. The simulation was fast enough for real-time control and large enough to let the temperature vary in different parts of the cube. When running the simulation the virtual water-cube is initially frozen, the simulation starts when heat is applied to one corner of the cube. The raised temperature slowly propagates through the cube and consequently melting it. A probe inside the virtual water cube samples data from the nearest molecule. The sampled data is then used to control the sound synthesis. One way of controlling the sound during a performance is to move the probe around inside the simulated melting cube of water. Each water molecule consists of two hydrogen atoms bound to one oxygen atom.

The sound for *In My Own Words* is generated in real-time during performance by four analogue oscillators connected so that oscillator 1 frequency modulates oscillator 2 which in turn frequency modulates oscillator 1. Oscillator 3 and 4 were connected in the same way, thus creating two pairs of two cross-modulating oscillators. The distance between the first hydrogen atom and the oxygen atom controls the pitch of oscillator 1 and 2 and the distance between the second hydrogen atom and the oxygen atom controls the pitch of oscillator 3 and 4. The bond angle (the angle between the two hydrogen atoms bonds to the oxygen atom) controls the modulation level. The four analogue oscillators were also able to modulate each other so that oscillator 2 frequency-modulates oscillator 3 and oscillator 4 frequency-modulates oscillator 1. Oscillator 2 could also be hard-synced to oscillator 1, as could oscillator 4 to oscillator 3.

To control the sound, one can change the clock that controls the speed of the virtual model, vary the temperature and move the probe within the model. Other aspects of the sound such as the individual component volumes and filtering were left open for improvisation. A further layer of sound was generated before the first performance. This sonic layer was produced by recording the sounds created when contact microphones are dipped into water as it freezes (see Figure 3.2).

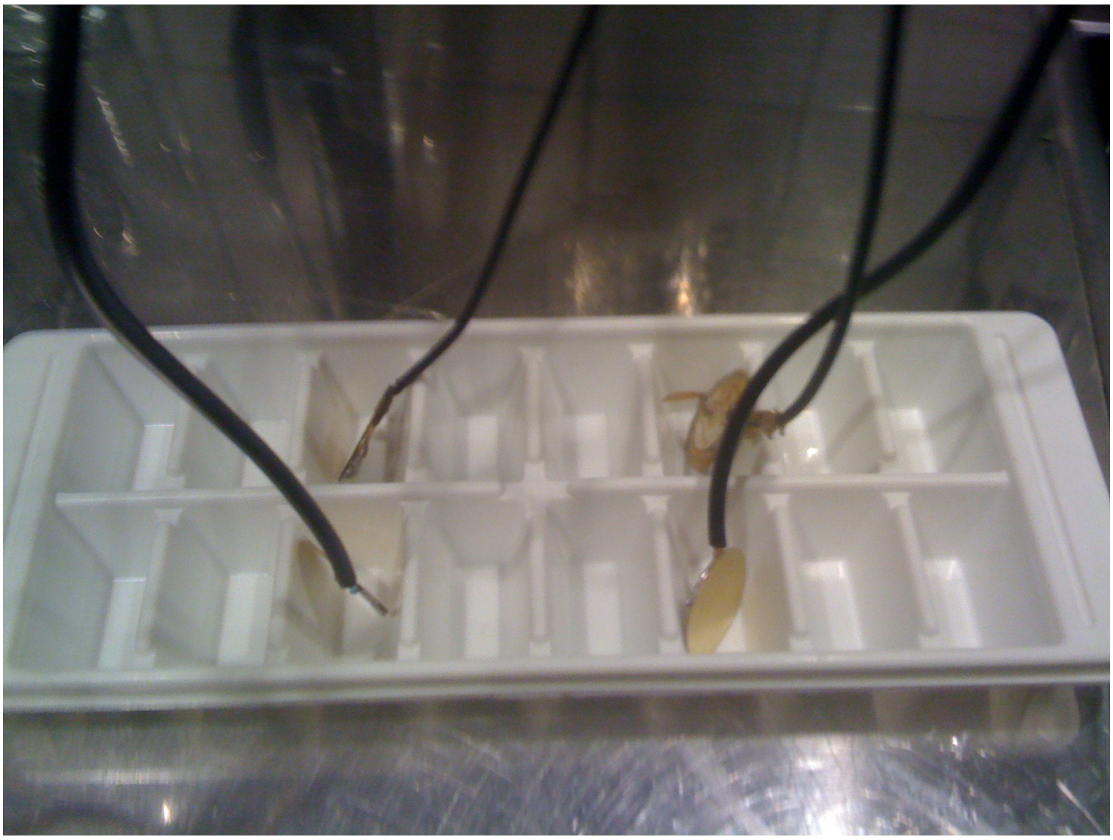

Figure 3.2 - Contact microphones in freezing water.

These sounds were then spatialised using 16 small loudspeakers and then rerecorded with a pair of high quality microphones. This made it possible to capture a small three-dimensional sound world in a stereo format (see Figure 3.3) and then move the whole miniature sound world around in a large concert diffusion set-up. Six of these sound-worlds were used and spatialised for *In My Own Words*. The spatialisation during the creation of the miniature sound-worlds and to some extent during the performance was done with a MaxMSP patch that I made for this purpose. The patch automatically distributes the incoming sounds out to different speakers depending on the phase and frequency content of the incoming sound. During the performance the spatialisation was mostly improvised mainly because the MaxMSP patch affected the sound quality too much.

In My Own Words was complicated on both a technical level and conceptually. Nevertheless, with a lot of hard work and careful planning the work came to fruition successfully. As such it sums up my approach of combining technical, aesthetic and conceptual planning and their integration in the creation of new work.

Figure 3.3 - Creating a small 3-dimensional sound world that itself could be moved around in a large diffusion set-up.

Chapter 4 Conclusion

My composition methods might seem somewhat complicated and convoluted. However, for me, these methods are a way of finding and generating structure in my compositions as well as a means for generating the sounds I use. As mentioned in the introduction of this commentary, emergent property is important to me, I like to put things together so that interesting events, structures, gestures, shapes, behaviours emerge from the assembly. For example, instead of composing with sounds as such, I want to compose with synthesis models that have their own will or inherent behaviours. It is like composing for instruments in an ensemble (instead of composing for sampled sounds) except that the instruments I like to use behave very different when compared to acoustic instruments. Even though many of my synthesis models are based on acoustic instruments I normally push these models to the limit, which often makes them behave in new and surprising ways. All my synthesis models are based on feedback and they can be pushed into chaotic behaviour. The same model that I use to generate sound can be used to generate the overall structure. In most cases I start in the middle of a piece with one synthesis idea and then work outwards with the structure growing from the process. Most of my pieces are based on several instances of the same sound model, *JUST* is based on one instance of one model and at the other end of the scale is *In My Own Words* which is based on several instances of several sound models.

Exploring things outside of my own comfort zone is another important factor in the way I work. When I was studying composition I discovered that the bits that were really problematic during the composition process were the bits I liked the most when reflecting on the finished piece. This made me realize that I really enjoyed solving problems and creating compositional challenges for myself. Soon I found that I often 'painted myself into corners' on purpose because being forced to solve an almost unsolvable problem made me think 'outside the box'. I was forced to leave my comfort zone since my favourite tricks seldom worked and I had to address the problem by a lot of hard work, creative thinking and in return often learning something new in the process. The more I became aware of this the more challenges became a part of my working process. The two most extreme examples of this are *In My Own Words* and *JUST*.

In *In My Own Words* I connected the atoms in freezing water to both sound-synthesis and to the spatialisation of the sounds constituting the work. In *JUST,* a real-time performance, I have to keep the dying and resonating echoes from an initial noise-burst alive for as long as possible.

As my techniques and my performance set-up (my instrument) have developed over the years, I make less and less distinction between what is a fixed media piece and what is a live performance piece. *In My Own Words* is a fixed media piece, although it can be performed live, whilst *JUST* is a piece intended for live performance but has been played as a fixed media piece.

To some extent this is also true for how I look at improvisation and composition since I improvise a lot while I compose and also in some of my pieces use structured improvisations. My current performance set-ups are the same as I use in the studio and they all originate from or are derived from the set-up I created for *JUST*. The journey to this set-up started with the software (a MaxMSP patch) I developed to control the many parameters in physical modelling synthesis and has since taken many turns in both the digital and physical world. This process of moving back and forth between different domains is one of my most important creative tools and the deeper I go into some aspect of creating a piece the more important it is for me not only to maintain an overall perspective but also to constantly change my perspective. A hands-on example is to develop a sound synthesis model in software and then adapt and implement the same model using analogue equipment. This forces me to change perspective and probably discover improvements, which then could be used in the software version and so forth. To move over boundaries, to work with the same ideas in different domains and always to have a fresh perspective is the key to everything for me. This is how I refine things as well as generate new ideas. For example when I formatted the endnotes for this thesis I got a new perspective on my writing, which gave me the idea for this final sentence!

Notes

1 Stockhausen, K. (1959) *Kontakte* Wergo (WER 60009)

2 Hofstadter, D. *Gödel Escher Bach: An Eternal Golden Braid* (Common Books 1979)

3 i) Lewes, G. H. *Problems of Life and Mind* (First Series), 2, London: Trübner, 1875, p.412
ii) Blitz, D. *Emergent Evolution: Qualitative Novelty and the Levels of Reality*. (Dordrecht: Kluwer Academic, 1992).

4 Pictures and video clips from the performance can be found at:
http://schhh.se/ljudingrepp/conveyer.html
an excerpt of the music is included (conveyercompilation.aif)

5 Documented online at http://www.annrosen.nu/hvt/traffas/traffas.html

6 Cope, D. *Virtual Music*. (Massachusetts: MIT Press, 2001)

7 Ibid. 2, (p. 129 -137)

8 Plomp, R. and W. J. M. Levelt, 'Tonal Consonance and Critical Bandwidth', *Journal of the Acoustical Society of America*.38, 548-560, 1965

9 Sethares, W. *Tuning Timbre Spectrum Scale*, (Springer, 2004)

10 Stockhausen, K. and E. Barkin, 'The Concept of Unity in Electronic Music' *Perspectives of New Music* Vol. 1, No. 1, 39-48, Autumn, 1962

11 Xenakis, I. *Formalized Music: Thought and Mathematics in Composition.* (Bloomington and London: Indiana University Press, 1971)

12 Bailey, D. *Improvisation: Its Nature and Practice* (Da Capo Press, 1980)

13 Carlos, W. 'Tuning at the crossroads' *Computer Music Journal* Vol. 11, No. 1, Spring, (MIT Press, 1987)

14 Ibid. 9

15 i) Smith III, J. O. 'Physical Modelling using Digital Waveguides', *Computer Music Journal*, Vol. 16, Number 4, Winter (MIT Press, 1992)
ii) Smith III, J. O. 'Musical Applications of Digital Waveguides', Stanford University Center for Computer Research in Music and Acoustics, STAN-M-39.
iii) Sullivan, C. R. 'Extending the Karplus-Strong Algorithm to Synthesize Electric Guitar Timbres with Distortion and Feedback', *Computer Music Journal*, Vol. 14, Number 3, Fall (MIT Press, 1990)
iv) Cook, P. 'SPASM, a Real-Time Vocal Tract Physical Model Controller' *Computer Music Journal*, Vol. 17, No. 1, Spring (MIT Press, 1993)
v) Karplus, K. and A. Strong, 'Digital Synthesis of Plucked String and Drum Timbres', *Computer Music Journal*, Vol. 7, Number 2, Summer (MIT Press, 1983)
vi) Keefe, D. H. 'Physical Modelling of Wind Instruments', *Computer Music Journal*, Vol. 16, Number 4, Winter (MIT Press, 1992)

16 Kramer, J. D. *The Time of Music: new meanings, new temporalities, new listening strategies* (New York, Schirmer, 1988)

17 Joe, J. and S. L. Gilman *Wagner and Cinema* (Indiana University Press, 2010)

18 http://www.annrosen.nu/sch/RT.html

19 Hellström, S.O., Jää-Aro, K-M., & A. Laaksonen, 'Perceptualisation of molecular dynamics data' in *Highlights in Molecular Dynamics* (2000)

Bibliography

Aristotle *Metaphysics* (Book H 1045a 8-10)

Bailey, D. *Improvisation: Its Nature and Practice* (Da Capo Press, 1980)

Carlos, W. *Tuning at the crossroads*. Computer Music Journal Vol. 11, No. 1, Spring, (MIT Press, 1987)

Cope, D. *Virtual Music* (MIT Press, 2001)

Hellström, S-O. 'The *Lightwork* Performance: Algorithmically Mediated Interaction for Virtual Environments' published in the Proceedings of CHI98, Los Angeles, USA, (New York: ACM Press, 1998)

Hofstadter, D. *Gödel Escher Bach: An Eternal Golden Braid* (Common Books, 1979)

Joe, J. and S. L. Gilman *Wagner and Cinema* (Indiana University Press, 2010)

Kramer, J. D. *The Time of Music: new meanings, new temporalities, new listening strategies* (New York, Schirmer, 1988)

Lewes, G. H. *Problems of Life and Mind* (First Series), 2, (London: Trübner, 1875)

R. Plomp and W. J. M. Levelt, "Tonal Consonance and Critical Bandwidth," *Journal of the Acoustical Society of America*.38, 548-560, 1965

Sethares, W. *Tuning Timbre Spectrum Scale* (Springer, 2004)

Stockhausen, K. and E. Barkin, 'The Concept of Unity in Electronic Music' *Perspectives of New Music* Vol. 1, No. 1, 39-48, Autumn, 1962

Xenakis, I. *Formalized Music: Thought and Mathematics in Composition.* (Bloomington and London: Indiana University Press, 1971)