2012

SOUND SYNTHESIS WITH CELLULAR AUTOMATA

Serquera, Jaime

http://hdl.handle.net/10026.1/1189

http://dx.doi.org/10.24382/4193
University of Plymouth

SOUND SYNTHESIS WITH CELLULAR AUTOMATA

by

JAIME SERQUERA

A thesis submitted to the University of Plymouth

in partial fulfilment for the degree of

DOCTOR OF PHILOSOPHY

School of Humanities and Performing Arts

Faculty of Arts

July 2012

Sound Synthesis with Cellular Automata

Jaime Serquera

Abstract

This thesis reports on new music technology research which investigates the use of cellular automata (CA) for the digital synthesis of dynamic sounds. The research addresses the problem of the sound design limitations of synthesis techniques based on CA. These limitations fundamentally stem from the unpredictable and autonomous nature of these computational models.

Therefore, the aim of this thesis is to develop a sound synthesis technique based on CA capable of allowing a sound design process. A critical analysis of previous research in this area will be presented in order to justify that this problem has not been previously solved. Also, it will be discussed why this problem is worthwhile to solve.

In order to achieve such aim, a novel approach is proposed which considers the output of CA as digital signals and uses DSP procedures to analyse them. This approach opens a large variety of possibilities for better understanding the self-organization process of CA with a view to identifying not only mapping possibilities for making the synthesis of sounds possible, but also control possibilities which enable a sound design process.

As a result of this approach, this thesis presents a technique called Histogram Mapping Synthesis (HMS), which is based on the statistical analysis of CA

evolutions by histogram measurements. HMS will be studied with four different automatons, and a considerable number of control mechanisms will be presented. These will show that HMS enables a reasonable sound design process.

With these control mechanisms it is possible to design and produce in a predictable and controllable manner a variety of timbres. Some of these timbres are imitations of sounds produced by acoustic means and others are novel. All the sounds obtained present dynamic features and many of them, including some of those that are novel, retain important characteristics of sounds produced by acoustic means.

Contents Overview

Table of Contents

List of Figures

List of Tables

Acknowledgments

I would like to express my sincere appreciation to

my supervisor, Prof. Eduardo Reck Miranda for introducing me to the amazing world of cellular automata, for believing in my potential, guiding my research in the right direction and for his continuous support.

my family since this PhD has been largely pursued as a part time student living with them in Spain. To my parents, Jesus y Laura and my sister Maria for their patience and practical support in spite of the fact that they would have preferred me to focus my efforts on more mainstream endeavours. It has been my brother David, to whom I am indebted most, who has consistently encouraged me and given me crucial support to accomplish this degree. Without him it would not have been possible.

the HuMPA Research Centre of the University of Plymouth for the studentship I was awarded at the last stage of my PhD in order to finish it as a full time student.

all the colleagues I had the opportunity to meet at ICCMR, especially Hanns Holger Rutz, Noris Mohd Norowi, Alexis Kirke, Marcelo Gimenes, Adolfo Maia, Joao Martins, Qijun Zhang and Alicja Knast for all their help during my short stays and companionship during my residences in Plymouth.

the internal and external examiners, for their constructive criticism, stimulating questions and useful comments.

my young nieces, Paula, Adriana and Helena for their innocent source of childish amusement which gave me much needed respite from my studies and kept me sane.

Author's Declaration

At no time during the registration for the degree of Doctor of Philosophy has the author been registered for any other University award without prior agreement of the Graduate Committee.

This study was partially financed with the aid of a studentship from the Centre for Humanities and Performing Arts Research (HuMPA) of the University of Plymouth.

A course on Advance Object Oriented Programming was undertaken at the Polytechnic University of Valencia, and an Advance Level Certificate in English Language was awarded by the “Escuela Oficial de Idiomas” of Spain.

A record of activities can be found in Appendix C.

Word count: 36,291.

Plymouth, 18 July 2012

Signed_____

Chapter 1 Introduction

Sound synthesis was historically one of the first research topics in the field of computer music. Since sound is the raw material of music, this area of research has over the years been a major theme.

The fundamental motivations for sound synthesis research are on the one hand the imitation of sounds produced by acoustic instruments, and on the other hand the search for new sounds. Some applications drive the research towards the imitation of sounds, such as the reduction of costs in music production. But other goals are more relevant from an artistic point of view, such as the design of instruments that initially sound like an acoustic model, but which have further prospects of extending its timbral possibilities. Some compositional works in this area developed at the Interdisciplinary Centre for Computer Music Research (ICCMR) are “Sacra Conversazione” by Eduardo Reck Miranda (Miranda 2005) and “Ophidian” by David Bessell (Bessell 2011). Such extension of instrumental possibilities is relevant to the second motivation, the search for new sounds, where the goal is to widen the palette of timbres available to the composer. This research area, in turn, also has relevance to the first motivation, the imitation of sounds, because even though novelty is the goal, it can be desirable to include and simulate certain properties of acoustic instrument sounds. The reason for doing so is to provide some naturalness to the novel sound in order to be more acceptable to human hearing and brain. Figure 1 graphically illustrates this scenario. Related to this overlapping between new and imitated sounds is the following quotation by Jean-Claude Risset:

> ‘The point of instrumental imitation is not only instrument duplication, of course. In particular, it sheds light on properties that can endow sounds with naturalness, richness, and also give them a characteristic identity’ (Risset 1985 p.12).

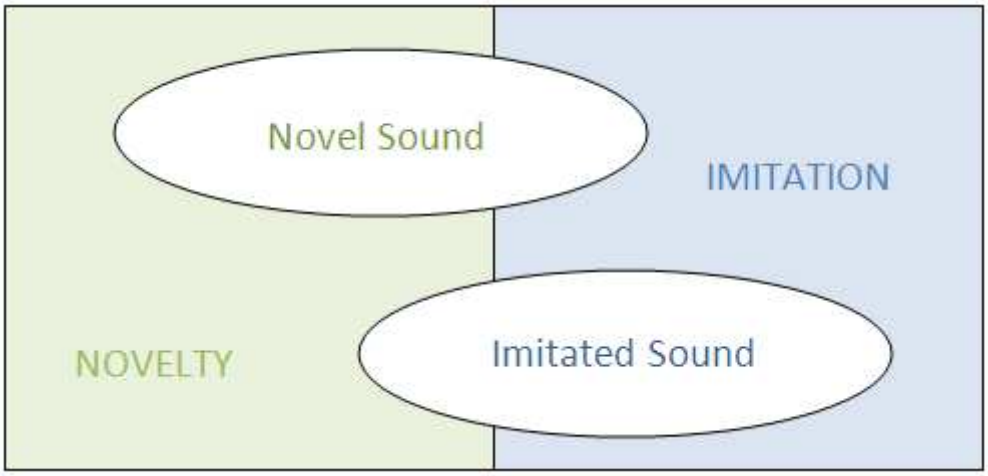

Figure 1: Overlapping between new and imitated sounds.

One of the main problems for digital sound synthesis is the amount of data that the synthesis instruments require in order to produce interesting results. This is, firstly, because usually, a synthesis algorithm is composed of many Unit Generators (UGs). As Curtis Roads explains:

> ‘UGs are signal processing modules like oscillators, filters, and amplifiers, which can be interconnected to form synthesis instruments or patches that generate sound signals’ (Roads 1996 p.89).

Moreover, The UGs usually have several input parameters, such as the frequency and amplitude of an oscillator. However, the real problem arises when the time dimension is considered. Note that the sounds produced by acoustic instruments present constant variations over time that make them pleasant to our hearing. Therefore, in sound synthesis we have to feed all the UG input parameters with streams of data over time in order to produce pleasant results. In many cases, such enormous amounts of data make the synthesis algorithms difficult to control manually and it becomes necessary to

investigate alternative methods for obtaining such control data. By way of illustration, Roads describes a number of approaches for obtaining control data for the particular case of the additive synthesis technique. These include data obtained from a sound analysis stage, which can then be used to resynthesise the original sound. A further possibility with this methodology is the transformation of sounds by modifying the original analysis data. Other sources of control data can be found in non-musical domains. For instance, the contour data of mountains has been used by some composers to feed their synthesis instruments (Roads 1996 p.143).

Cellular Automata (CA) also represent a solution to this problem of control because with few parameter specifications it is possible to obtain massive amounts of structured data. CA are of interest to computer musicians because of their emergent structures –patterns not created by a single rule but through the local interaction of multiple units with relatively simple rules. This dynamic process leading to some order allows the musician to explore new forms of organization. In sound synthesis, CA are normally used for controlling the parameters of a synthesis instrument over time. The goal is to transfer the structured evolution of CA onto the sound synthesis domain. This is normally done through a mapping: a set of correspondences between different domains. The origins of this line of research can be traced back to the 1980s when Barry Truax suggested the use of self-organised systems for the control of granular synthesis (Truax 1988 p.25). Since then, this has been an active area with recent research publications including (McLaughlin et al. 2010) and (Morimoto 2010). However, the main problem of synthesis techniques based on CA is the lack of control exercised over the automata, something that restricts the sound

design possibilities from the user's point of view. Thus, this thesis is in great part motivated by the challenge of providing such control over the CA. More details on the scientific motivations of this thesis are developed in the following section.

1.1 Motivations

At the beginning of my PhD degree I was primarily interested in the investigation of ways that new computational methods might provide novel approaches to the synthesis of dynamic sounds. To that end, I decided to investigate dynamic systems which show great promise in this role and are increasingly used nowadays because of the improvement of computers in terms of computational power. I decided to focus on CA which, as documented in Section 2.3, have many properties that are useful for sound synthesis. In addition, CA have seen far less study in the audio domain than in the music domain.

The literature makes constant references to the lack of control available in synthesis techniques based on CA. A given automaton, through a given mapping process, is capable of feeding over time a synthesis instrument with dynamically structured data (something which produces interesting results) but, and this is the rub, at its own will. I was motivated by the idea that the previously mentioned capabilities of CA could be boosted if the user's intentions could be also complementarily transferred to the synthesis process by means of control mechanisms.

Therefore, this thesis is new music technology research which, centred around the user in order to address their technical and artistic needs, investigates CA algorithms for computer-aided sound design.

1.2 Thesis Research Overview

It is appropriate at this point to briefly discuss the main problems, methods and results of this thesis, before they are developed in detail in subsequent chapters.

This thesis addresses the problem of the sound design limitations of synthesis techniques based on CA. These limitations fundamentally stem from the unpredictable and autonomous nature of these computational models which makes them difficult to be controlled. Therefore, the aim of this thesis is to develop a sound synthesis technique based on CA capable of allowing a sound design process. In order to achieve this goal, a novel approach is proposed based on the use of digital signal processing (DSP) methods that enable a better understanding of the self-organization process of CA. The main problem of developing such a synthesis technique can be divided into these three key sub-problems:

- To identify/develop DSP methods that offer useful information on the evolution of CA, and to select appropriate CA.

- To devise mapping processes between the CA domain and the sound synthesis domain to make the synthesis of sounds possible.

- To identify control possibilities and develop control mechanisms which enable a flexible sound design process.

The main problem has been addressed by the invention and development of a new spectral synthesis technique called Histogram Mapping Synthesis (HMS). HMS is discussed below and the following peer reviewed publications report on various milestones of the development of this technique: (Serquera et al. 2008), (Serquera et al. 2008), (Serquera et al. 2010), (Serquera et al. 2010), (Serquera et al. 2010), (Serquera et al. 2011).

Solving the first sub-problem has been largely a matter of experimentation. A DSP method has been identified/developed: the statistical analysis of CA evolutions by histogram measurements. It constitutes the basis of HMS and has provided useful information on the evolution of multi-state CA. On the basis of this DSP method, it has been possible to identify mapping possibilities for sound synthesis and control possibilities for sound design. Regarding specific automatons, this thesis explores the possibilities of three CA that offer interesting results with HMS: the multitype voter model, the hodge podge machine and the plurality vote rule. In addition, a number of modifications and extensions have been made to these CA, some of which could be considered as new automatons.

Regarding the second sub-problem, HMS is significant because it offers different mapping possibilities depending on the resemblances that can be found between the outcome of a CA histogram analysis and certain spectral components of sounds. These possible mappings based on such resemblances are distinctive because in most other synthesis techniques based on CA there is not an intuitive correspondence between the components of the automaton and the components of a sound.

As for the third sub-problem, a considerable number of control mechanisms have been identified and developed, which show that HMS enables a reasonable sound design process. With these control mechanisms it is possible to design and produce, in a predictable and controllable manner, a variety of timbres. Some of these timbres are imitations of sounds produced by acoustic means and others are novel. All the sounds obtained provide dynamic features and many of them, including some of those that are novel, retain important characteristics of sounds produced by acoustic means.

1.3 Thesis Structure Overview

The structure of this thesis has been largely influenced by (Chinneck 1988), a guide that has been useful to many postgraduate students, as demonstrated by the fact that it has been translated into several languages. The main reason for following this guide is that it fits with the type of research that I have pursued.

Following this introductory chapter, “Chapter 2 Background Information” gives information on the two main fields that this interdisciplinary research connects:

Cellular Automata and Digital Sound Synthesis. The chapter concludes with a section on how these two fields can be combined.

“Chapter 3 Previous Research on Sound Synthesis with CA” presents preceding work on the synthesis of sounds using CA. It is strictly a presentation since a critical analysis will come in the next chapter. This chapter is organised into two sections. The first one encompasses research which does not deal with the control of CA. The second section includes research that shows a concern, to a certain extent, on the CA control issue.

“Chapter 4 Problem Statement” is a key chapter: the main research problem is formulated. After a concise statement of the problem, a critical analysis on the basis of certain criteria of previous systems will be developed in order to justify that the problem has not been solved previously, and therefore, it is still an open research problem. Finally, it will be discussed why this problem is worthwhile to solve.

“Chapter 5 Histogram Mapping Synthesis” is the core of this thesis. It describes how the problem has been solved, at least to certain extent.

“Chapter 6 Conclusions” contains the following elements: the conclusions, a summary of contributions to knowledge, and suggestions for further research. All these three sections are developed in numbered concise paragraphs for the sake of organization and clarity.

Finally, three Appendices are presented. “Appendix A: Sound Examples” gives explanatory notes for each of the sound examples that can be found in the accompanying CD. “Appendix B: Video Examples” introduces each of the video examples that can be found in the same CD. And “Appendix C: Record of Activities” lists the publications and presentations carried out during the PhD.

Chapter 2 Background Information

This chapter gives background information on the two main fields that this interdisciplinary thesis connects. The first section will focus on Digital Sound Synthesis and the second on Cellular Automata. The chapter concludes with a section on how these two fields are, or can be connected.

2.1 Digital Sound Synthesis

2.1.1 Taxonomy

Over the course of time, a considerable number of digital synthesis techniques have been developed. In order to facilitate their study a number of taxonomies have been proposed that classify them into different categories. However, there is no universal agreement among experts on these classifications. As a starting point we can take as a reference the taxonomy devised by Eduardo Reck Miranda in (Miranda 2002 p.xvi). He assumes every technique is based on a model and establishes four categories according to their modelling approach: loose modelling, spectrum modelling, source modelling, and time-based approaches. It is important to note that a specific technique may be considered as belonging to more than one of these categories. For instance, Miranda discusses subtractive synthesis within two different categories, namely, spectrum and source modelling.

Miranda describes the first class as follows:

> ‘Loose modelling techniques […] tend to provide synthesis parameters that bear little relation to the acoustic world. They are usually based entirely upon conceptual mathematical formulae. It is often difficult to predict the outcome and to explore the potential of a loose model’ (Miranda 2002 p.xvi).

Examples of techniques within this category include Amplitude Modulation (AM), Frequency Modulation (FM), Waveshaping Synthesis, Walsh Synthesis and Wavetable Synthesis.

The next two categories, spectrum and source modelling, in contrast, provide synthesis parameters more related to the acoustic world. Miranda differentiates both these two classes as follows:

> ‘The fundamental difference between source and spectrum modelling techniques is that the former tends to model a sound at its source, whilst the latter tends to model a sound at the basilar membrane of the human ear’ (Miranda 2002 p.xvi).

A general description of the source modelling category is given as follows:

> ‘In general, source modelling techniques work by emulating the functioning of acoustic musical instruments […]. The key issue of source modelling is the emulation of acoustic sound generators rather than of the sounds themselves. For example, whilst some synthesis techniques (e.g. additive synthesis) attempt to produce a flute-like sound using methods that have little resemblance to the functioning of the flute, a source modelling technique would attempt to synthesise it by emulating a jet of air passing through a mouthpiece into a resonating pipe.
>
> The implementation of a source model is not straightforward. However, once the model is implemented, it is not complicated to interpret the role of their synthesis parameters. Take, for example, a singing voice-like instrument. A loose model using FM, for instance, would provide relatively complex synthesis parameters, such as modulation index and frequency ratio. Conversely, a source model using waveguide filtering, for instance, would provide more easily interpreted synthesis parameters, such as air pressure, vocal tract shape and throat radiation output’ (Miranda 2002 p.xvii).

The source modelling category embraces, among others, Modal Synthesis, Waveguide Synthesis, the Karplus-Strong algorithm and Physical Modelling techniques.

Regarding spectrum modelling approaches, Miranda introduces the category as follows:

> ‘Spectral modelling techniques are the legacy of the Fourier analysis theory. [...] Fourier analysis considers that a pitched sound is made up of various sinusoidal components, where the frequencies of higher components are integral multiples of the frequency of the lowest component. [...] timbre is the result of the presence of specific components and their relative amplitudes [...]. Despite the fact that not all interesting musical sounds have a clear pitch and the pitch of a sound may not necessarily correspond to the lower component of its spectrum, Fourier analysis still constitutes one of the pillars of acoustics and music’ (Miranda 2002 p.50).

Synthesis techniques within this category include Additive Synthesis and Formant Synthesis. This modelling approach will be studied in more depth later in this thesis.

Finally, the last category is described as follows:

> ‘time-based techniques approach synthesis from a time domain perspective. The parameters of time-based synthesis tend to describe sound evolution and transformation of time-related features; e.g. in terms of time lapses’ (Miranda 2002 p.xvii).

Examples of time-based techniques include Granular Synthesis, Pulsar Synthesis and Waveset Distortion.

2.1.2 Spectrum modelling

2.1.2.1 Introduction

Since Histogram Mapping Synthesis (the synthesis technique discussed in this thesis) can be considered as a spectral technique, we are going to study in more depth the spectrum modelling category. To that end, this introduction will start by explaining the concepts of Fourier analysis, spectrum and spectral components.

A Fourier analysis involves the decomposition of a periodic function into simple sinusoidal components with certain amplitudes frequencies and phases. For our purposes, the periodic function can be a periodic sound wave which can be decomposed into simple sine waves. However, most of the pitched sounds that we hear are not perfectly periodic but quasi-periodic, and many others, like noise, might be not periodic at all. In order to address this problem, Fourier theory includes a number of mathematical solutions which facilitate the decomposition of non-periodic functions (it is beyond the scope of this thesis to examine the mathematical details of this).

Perry R. Cook illustrates this sound decomposition process with the analogy of decomposing white light into coloured beams through a prism (Cook 2002 p.52). Subsequently, he gives a set of useful definitions:

> ‘The set of individual amplitudes and phases of the sines that make up a sound are called a *frequency spectrum*. The mathematical technique used to turn a time-domain waveform into a frequency-domain spectrum

> is called *transform* […] Using the frequency spectrum to inspect aspects of a sound is called *spectral analysis*' (Cook 2002 p.52).

In addition to the abovementioned mathematical solutions, a number of algorithms have been developed that efficiently compute a Fourier transform. One of the best known is the Fast Fourier Transform (FFT). In order to take into account the evolution of sound over time, another algorithm is used, the Short-Time Fourier Transform (STFT). This algorithm slices the sound signal into short pieces and applies the Fourier transform to each slice. As part of this process, a windowing operation is normally applied which consists in multiplying each slice of the sound by a windowing function. These windows adopt different shapes such as rectangular, triangular, Gaussian bell, etc. The windowing has important consequences in the resulting frequency domain representation because such multiplication in the time domain becomes a convolution in the frequency domain between the Fourier transforms of both signals (the sound and the windowing function). It should be noticed that even if none of these windowing functions are used, the slicing process is still equivalent to using a rectangular window. Another relevant parameter is the length of the slices, or the window size. The Fourier theory establishes that the longer the window, the more frequency resolution we will have. However, we can see that with long windows, we lose time resolution. This is known as the trade-off between frequency and temporal accuracy. For further reference on these and other STFT settings, see (Zolzer 2002 p.379). The STFT transform gives information on the temporal evolution of the frequency content of a sound. The graphical representation of this information is given by the so-called spectrogram. The spectrogram has three dimensions: frequency, time and amplitude or spectral magnitude.

We will now analyse different spectrograms computed with the STFT transform in order to examine some basic features. After this analysis it will be seen that some sounds can be better modelled by means other than sinusoidal components. The first example will help us identify the sinusoidal components of a sound within a spectrogram. Figure 2 shows the zoomed spectrogram of a saxophone tone. We can see a number of equidistant narrow bands all along the duration of the tone. These are the sinusoidal components, also known as partials. They are equidistant because they have frequencies which are integer multiples of the fundamental frequency (the left-most narrow band in the figure) and therefore, they follow the so-called harmonic series. Hence, they are also referred to as harmonics. Other sounds such as bell tones have non-harmonic partials, meaning that the partials do not follow the harmonic series and are not equidistant.

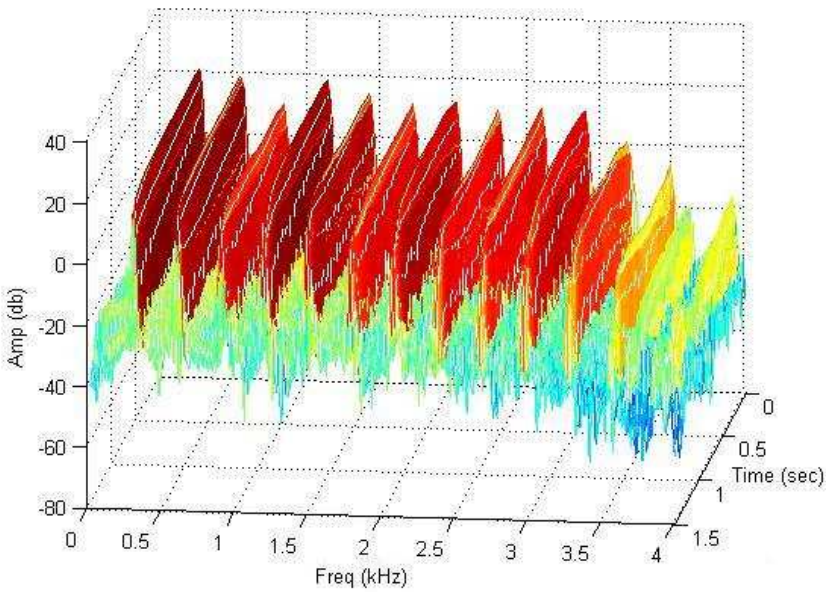

Figure 2: Spectrogram of a saxophone tone.

The next example is the spectrogram of a piano tone (Figure 3). Here we can also see a set of narrow bands. These fade out because the piano produces damped sounds.

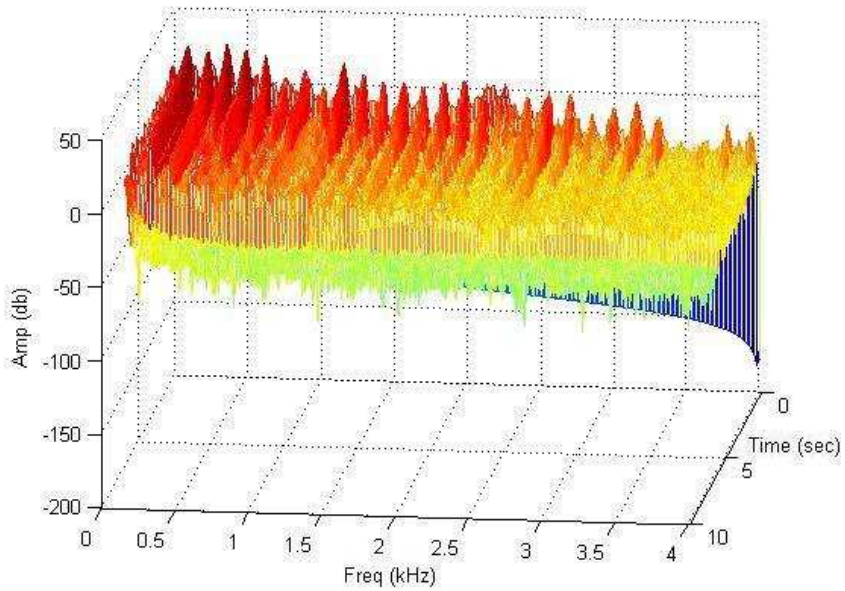

Figure 3: Spectrogram of a piano tone.

Let us note now that the sounds produced by acoustic instruments usually present more spectral complexity in the attack portion than in the rest of the sound. For instance, the piano has a strong attack with the corresponding noise of the hammer stroke. We can verify this fact by removing the sinusoidal components of the piano sound by means of a careful DSP process. As a result of this operation it remains a noisy structure called residual, corresponding to the hammer stroke, as shown in Figure 4 (see Sound Example 1 in Appendix A). In this figure we do not see prominent narrow bands, meaning that this structure would not be well modelled by sinusoidal components. It would require an extremely large number of sinusoidal components to model it. Instead, the hammer stroke would be better modelled by a combination of a filtered noise component and a transient component (this will be examined later).

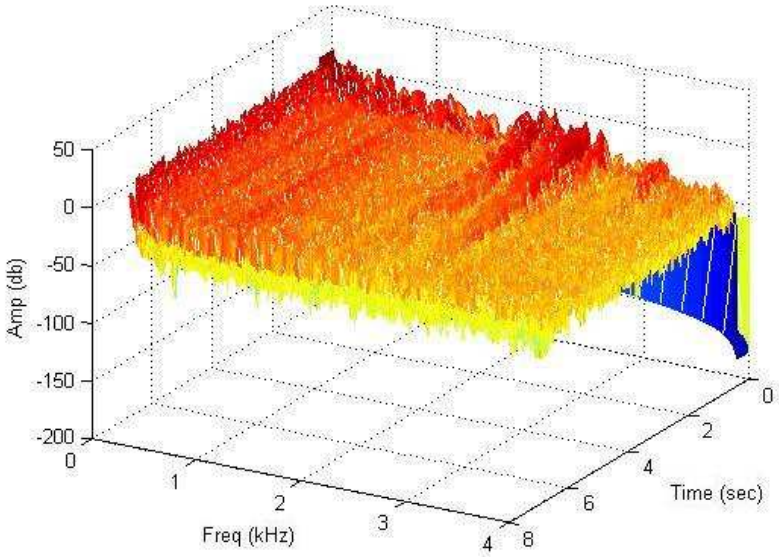

Figure 4: “Residual” of a piano tone.

Following this example we are in a better position to introduce the basis of spectrum modelling approaches. Spectral models model the spectrum of sounds by means of elementary or basic spectral components. In the piano example these were the sinusoidal components and the noise and transient components, but in other sound types other types of components can be identified such as spectral envelopes, formants, etc. In general, different spectral models make use of different types of spectral components and different techniques for processing them, as it will be studied in the following subsections.

These components are normally extracted in an analysis stage, and are able to represent the most characteristic features of the sound in question. Although modelling is a simplification of reality, some of these models can achieve excellent reconstructions of the original sound, in what is called the resynthesis stage. Sometimes the difference between the original and the synthesised

version is imperceptible to human hearing. In addition, it is normally possible to manipulate these components prior to resynthesis, and thus produce new sounds that are transformations of the original. This happens in the transformation stage. Note that it is also possible to do pure synthesis by creating these components from scratch, as is the case of HMS. In this way, the knowledge gained from an empirical practice in sound analysis can be used to obtain interesting results.

In the following subsections a representative selection of spectral models and techniques will be studied from the point of view of the spectral components they use. These models, techniques and spectral components will have a presence later in this thesis.

2.1.2.2 Sinusoidal Additive Synthesis

Sinusoidal additive synthesis is a spectral technique based on Fourier theory. As its name indicates, the spectral components utilized are sinusoidal components. The technique consists in creating a complex sound by summing up simple sinusoids. It can be expressed in mathematical form as follows:

$$y(t) = \sum_{i=1}^{I} A_i(t)\sin(2\pi f_i(t)t + \theta_i) \tag{1}$$

where I stands for the number of sinusoids and $A_i(t)$, $f_i(t)$ and θ_i are the instantaneous amplitude and frequency and the initial phase of the i-th sinusoid respectively.

Figure 5 illustrates the technique by means of a block diagram with oscillators – in which the phase information is normally omitted.

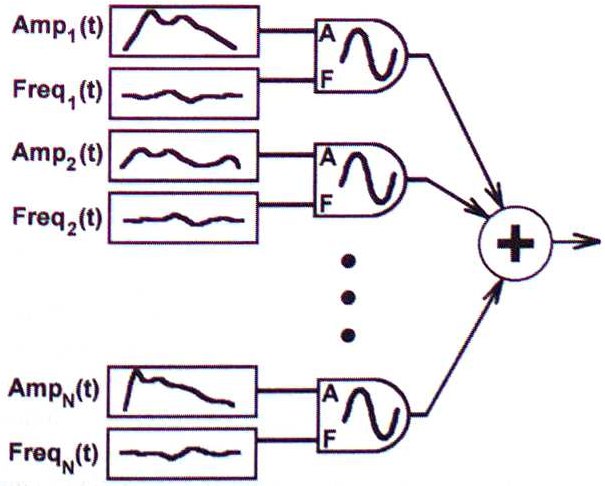

Figure 5: Sinusoidal additive synthesis. (Cook 2002).

Needless to say, the main problem of additive synthesis lies in the provision of control functions for the time-varying amplitudes and frequencies of every oscillator. At this point it is appropriate to introduce a widely established amplitude envelope model: the ADSR model. ADSR stands for attack, decay, sustain and release; each of which denoting a different stage in the production of a musical sound. This model will be used to refer to these stages during the course of this thesis. The model is illustrated in Figure 6.

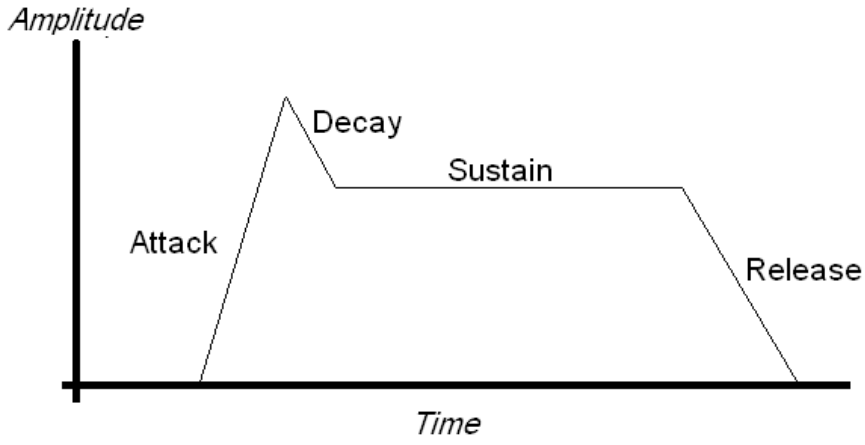

Figure 6: ADSR envelope model.

2.1.2.3 Subtractive Synthesis

In subtractive synthesis, a complex sound source is used to excite one or more filters (Figure 7). This technique is also known as source-filter synthesis.

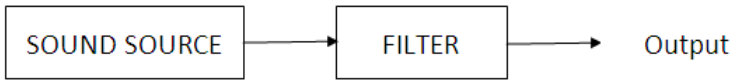

Figure 7: Subtractive synthesis block diagram.

Subtractive synthesis can be considered a spectral technique –the effect of the filter is better understood in the frequency domain. The filters are normally characterised in the frequency domain by the magnitude frequency response or filter shape. These are spectral envelopes which by means of multiplication (convolution in the time domain) reshape the spectrum of the input signal. The filtering process is illustrated in Figure 8. Typical sources of excitation signals are noise and impulse train generators, as well as triangular, square or

sawtooth oscillators. Typical standard filter shapes are low-pass, band-pass and high-pass.

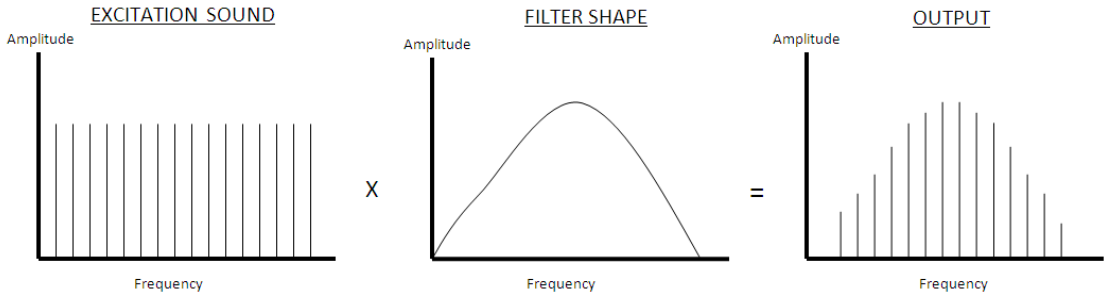

Figure 8: Filtering process in the frequency domain.

As with other synthesis techniques, in order to get interesting results the technique has to be supplied with time-varying control data. That can be achieved by the provision of time-varying filter shapes or magnitude responses.

2.1.2.4 Sines + Noise + Transients Models

A number of analysis/synthesis modelling approaches have been developed based on three types of spectral components, namely, sinusoidal, noise and transients. In these models, the sinusoidal and noise components are characterized by their slow varying spectral evolution, as opposed to the faster nature of transients. Since the “transient” concept is broader than the concept of “sine” or “noise”, the following clarifications are useful. Harvey Thornburg defines transients as follows:

> ‘events where the spectral content changes abruptly, or regions for which spectral content is best modeled as undergoing persistent change’ (Thornburg 2005 p.iv).

In addition he gives the following examples that can be considered as transients:

- ‘Abrupt changes in amplitudes, phases, or frequencies: in recordings of acoustic material, these changes are often due to energy inputs on the part of the performer; hence, abrupt change transients often associate with onsets of note events or other phenomena that may be notated in the score

- Rapid decays in amplitudes, usually associated with attack regions following onsets of percussive sources

- Fast transitions in frequencies and amplitudes: musical examples include expressive pitch variations (portamento, vibrato, etc.) and timbral transitions (such as a rapid shift in the vocal formant structure)

- Noise and chaotic regimes, primarily responsible for textural effects: environmental sounds, such as rain or crackling fire, exhibit persistent textures which are important to preserve in resynthesis; [...]’ (Thornburg 2005 p.3).

We will now examine the chronology of developments in these modelling approaches. The first models to appear were the sinusoidal models which model the spectrum of sounds using only time-varying sinusoids. These models consist of a number of algorithms that analyse the STFT of sounds in order to extract the amplitudes, frequencies and phases of the sinusoids. These parameters could then be used to control sinusoidal additive synthesis to resynthesise the original sound. Transformations of the original sound are possible through modifying these parameters prior to resynthesis. Examples of these models are the sinusoidal transformation system (STS) (McAulay et al. 1986) and PARSHL (Smith et al. 1987).

However, these models are not flexible and efficient in representing noisy sound sections such as the attack of musical sounds. In this respect, Serra and Smith explain:

> ‘Using sinusoids to simulate noise is extremely expensive because, in principle, noise consists of sinusoids at every frequency within the band limits’ (Serra et al. 1990 p.13).

Subsequently, spectral modeling synthesis (SMS) (Serra et al. 1990) improved the sinusoidal models by explicitly modelling noise. Serra describes his method as follows:

> ‘Our particular approach is based on modeling sounds as stable sinusoids (partials) plus noise (residual component), therefore analyzing sounds with this model and generating new sounds from the analyzed data. The analysis procedure detects partials by studying the time-varying spectral characteristics of a sound and represents them with time-varying sinusoids. These partials are then subtracted from the original sound and the remaining “residual” is represented as a time-varying filtered white noise component. The synthesis procedure is a combination of additive synthesis for the sinusoidal part, and subtractive synthesis for the noise part’ (Serra 1997).

However, as Verma and Meng explain, transients are not well represented in SMS:

> ‘Although this technique has been very successful, transients do not fit well into this model, because transients modeled as filtered noise lose sharpness in their attack and tend to sound dull. Because transients are poorly modeled by sinusoids or noise [...], transients must be handled separately’ (Verma et al. 2000 p.47).

As a result, these researchers developed a system called transient modeling synthesis (TMS) (Verma et al. 2000) which extends SMS by explicitly modelling transients. Figure 9 illustrates this three-part decomposition.

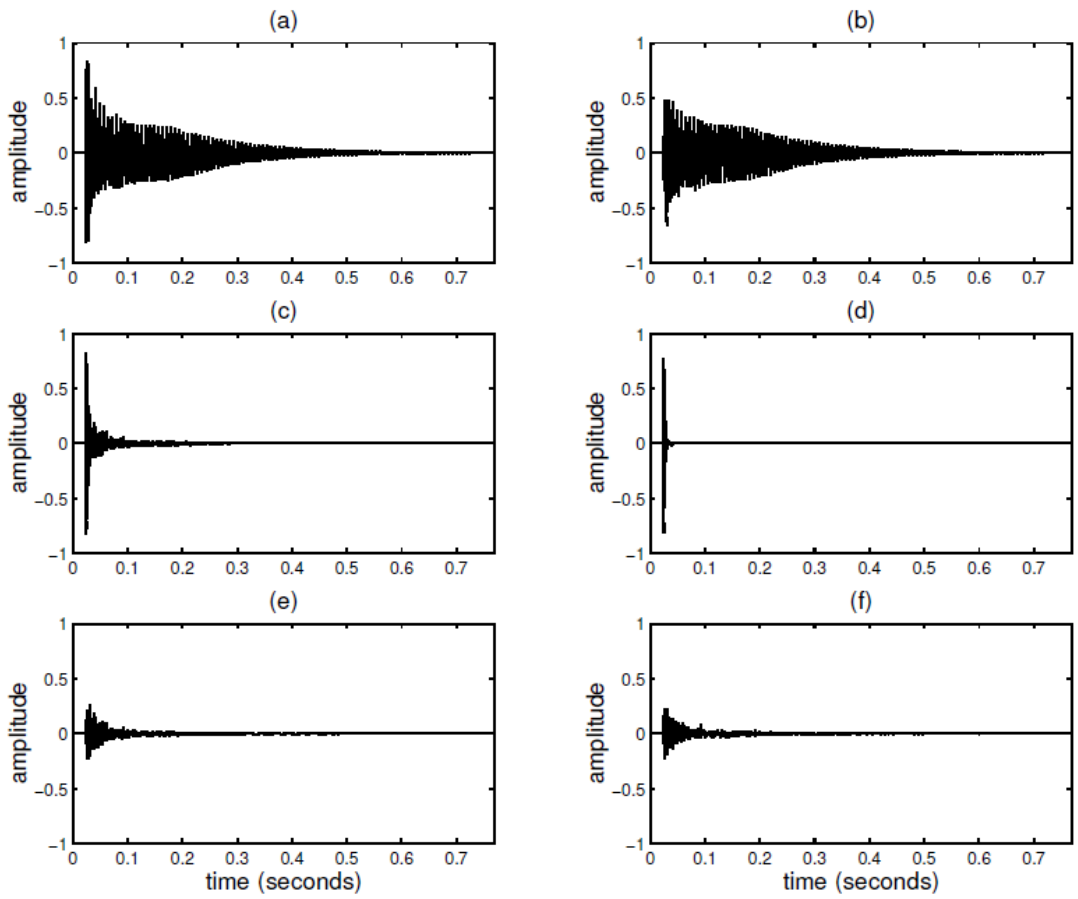

Figure 9: Spectral modelling of a xylophone tone. (Verma et al. 2000).
Original xylophone (a); synthesized sinusoids (b); first residual containing transients + noise (c); synthesized transients (d); second residual containing noise (e); and synthesized noise (f).

Note that the transformation stage of these models strongly benefits from sounds being separated into different components. For instance, for time-stretching (modification of the duration) or pitch-shifting (modification of the pitch) processes, it is usually better to leave the noise and/or transient components unprocessed.

2.1.2.5 Formant Synthesis

In this last subsection we are going to deal with two types of spectral components, namely, spectral envelopes and formants. A spectral envelope, already mentioned in Subsection 2.1.2.3, is a smooth curve that envelopes the magnitude of a spectrum. Formants are prominent lobes or peaks in a spectral envelope.

The spectral envelopes and formants represent an alternative to a detailed Fourier analysis for characterising sounds. Formants have been used mainly for the analysis and synthesis of voice sounds but can also be used for modelling instrument sounds. A type of sound can be characterised, independently of pitch, by the number of formants and their respective center frequencies, bandwidths and amplitudes. Figure 10 shows the spectrum of a synthesised [a] vowel with its spectral envelope representing three formants.

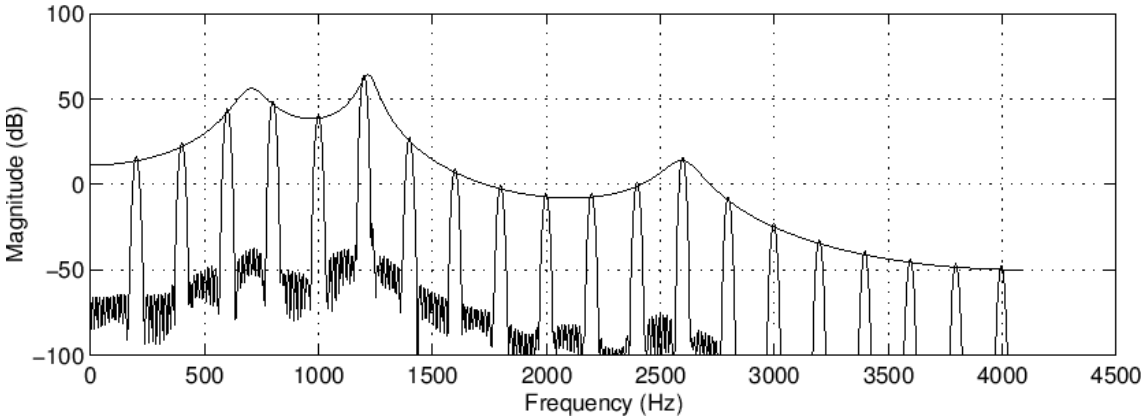

Figure 10: Spectrum of a synthetic [a] vowel with formant envelope. (Smith 2011).

Different techniques, such as additive synthesis, subtractive synthesis, FM, etc., can generate formants. However, there are techniques specifically developed for formant synthesis. Two of them are FOF (Fonctions d'Onde Formantique)

and VOSIM (voice simulator). Originally, these were attempts toward the simulation of voice and are described in (Miranda 2002 p.67) and (Roads 1996 p.299). WF (window-function) on the other hand, was developed to simulate the formants of traditional music instruments, and is described in (Roads 1996 p.311).

2.2 Cellular Automata

Let us now deal with CA. (Adamatzky 2001 p.11) presents a quotation from mathematician Stanislaw Ulam which constitutes the first evidence of CA theory:

> 'Suppose one has an infinite regular system of lattice points in E^n, each capable of existing in various states $S_1,\ldots,S_k$. Each lattice point has a well defined system of m neigbours, and it is assumed that the state of each point at time t + 1 is uniquely determined by the states of all its neighbours at time t. Assuming that at time t only a finite set of points are active, one wants to know how the activation will spread' (Ulam 1960).

In addition, as Adamatzky recalls, that quotation 'quite accurately defines a cellular automaton itself' (nevertheless, below an alternative description will be given adapted to the work pursued in this thesis). The term "Cellular Automata", however, was coined in the work of mathematician John von Neumann (Neumann et al. 1966) on self-reproducing machines (Sarkar 2000 p.82).

To give a description of the functioning of these models, the following will be stated embracing the CA used in this thesis: Cellular automata are normally implemented as a regular grid of cells in one or more dimensions. Each cell may assume any state from a finite set of values. CA evolve in successive generations at every time unit. For each generation, the values of all cells

change simultaneously according to a set of transition rules that takes into account the states of the neighbouring cells. The states of the cells may represent different colours, and therefore the functioning of a two-dimensional cellular automaton may be displayed on the computer screen as a sequence of images, like an animated film.

In this description, four CA features are addressed which, according to (Sarkar 2000 p.85) characterise an automaton, namely ‘the geometry of the underlying medium which contain the cells; the local transition rule; the states of the cell; and the neighborhood of a cell’. In the following paragraphs these features are studied, allowing more CA variants to be examined in addition to those utilised for this thesis.

Cell states: In common CA, the cells may assume any state from a finite set of discrete values. All the CA used in this thesis use values corresponding to a finite set of integers. But other options are possible. For instance, Rudy Rucker proposes the use of continuous-valued CA, that is, with cell states based on real numbers. These CA can be found implemented in the software called CAPOW, and (Ostrov et al. 1996) develops this idea. In addition, an important aspect related to the cell states is the initial configuration of cell values with which the automaton begins the run. The task of defining the initial configuration is referred to in this thesis as “cell seeding”. Most of the experiments that will be presented have been carried out with initial configurations consisting of uniform random distributions of cell values. Nevertheless, a number of different cell seeding processes will be also studied.

Geometry: This thesis uses CA implemented in two-dimensional grids. Other dimensions such as one and three are also common, and more dimensions are possible. In practical terms, when the grids are finite, another aspect must be considered –the boundary conditions. In this thesis periodic and fixed boundary conditions are considered. Periodic boundaries in 2D CA involve the connection of the left extreme with the right extreme and the top extreme with the bottom extreme. Therefore, the two dimensions are folded forming a torus shape. Fixed boundary conditions refer to the assignment of virtual adjacent cells to the extremes of the CA, whose values do not change during the run. A common type involves having these adjacent cells equal to zero, known as fixed zero or null boundary condition. Other boundary conditions not used in this thesis are, for instance, adiabatic and reflecting. Adiabatic refers to the duplication of the cell values in the virtual adjacent cells. In reflecting boundaries, the grid is reflected at the boundary so that the virtual adjacent cells adopt the values of the opposite neighbour. See Figure 11.

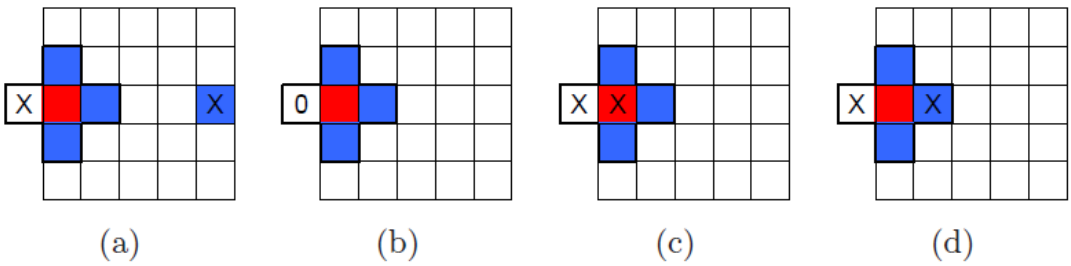

Figure 11: CA Boundaries. (Tovar et al. 2004).
(a) Periodic, (b) Fixed, (c) Adiabatic and (d) Reflecting.

Another type of cellular automata which is not used in this thesis is the structurally dynamic CA. Andrew Adamatzky defines this class as follows:

> 'The notion of structurally dynamic automata implies that the lattices evolve dynamically: the links connecting cells evolve together with the cells' (Adamatzky 2001 p.16).

Neighbourhood: In this thesis, two types of neighbourhoods (and some of their variants) are used: the “von Neumann” and the “Moore” neighbourhoods. The von Neumann neighbourhood defines neighbouring cells as the four orthogonally adjacent cells surrounding the central cell. The Moore neighbourhood includes not only the orthogonally but also the diagonally adjacent cells. In both these neighbourhoods the center cell is considered as a neighbour of itself. A variant of these neighbourhoods has been considered in this thesis, in which the center cell is not defined as a neighbour. Further variations of the Moore neighbourhood, defined with higher ranges, have been used in this thesis. That is, including cells which are further than one cell away from the central cell. A final example of another type of neighbourhood not used in this thesis is the MvonN, which is a combination of the von Neumann and Moore neighbourhoods. See Figure 12.

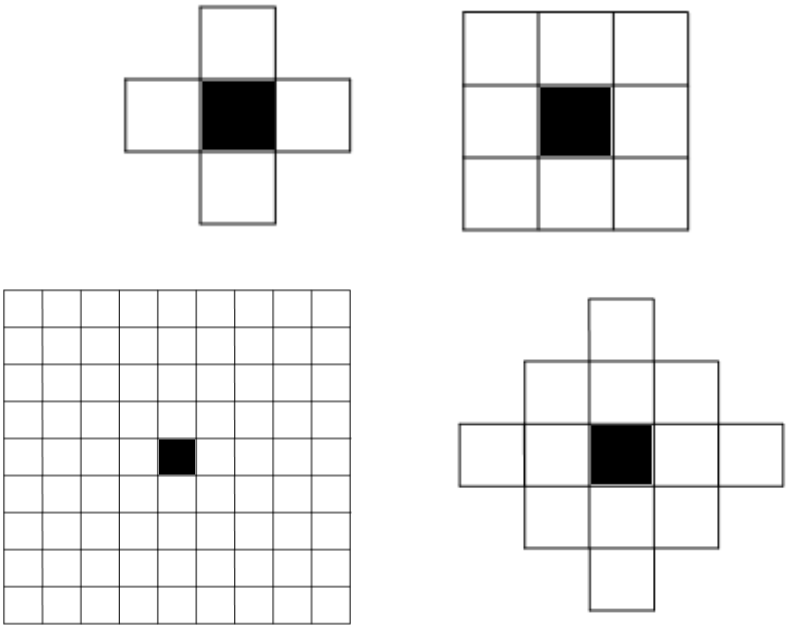

Figure 12: 2D neighbourhoods.
(Top-Left) von Neumann. (Top-Right) Moore. (Bottom-Left) Moore of range 4. (Bottom-Right) MvonN.

Rules: The transition rules that have been implemented in this thesis are of two types: deterministic and non-deterministic. With a deterministic rule, for a given configuration of cell states, the updated cell state is always the same. With a non-deterministic rule the next state is not only dependent on the neighbourhood, but also on some random input and/or probabilistic component. A probabilistic rule gives the probabilities that each cell will transition to the next possible state. In this case not all the cells update their state simultaneously and the automaton could be considered as “asynchronous”. (However, Adamatzky points out that ‘no proper asynchronous cellular automaton has ever existed’ (Adamatzky 2001 p.16)). Another interesting type is the “hybrid” CA in which different rules are applied to different cells. Furthermore, “programmable” CA are those that change the rule of a cell at each time step. These types of CA are implemented in Section 5.4.3.2 (Control by Coupled CA) where, although the same rule has been applied, it has been done so with different rule parameter values and therefore, in practical terms, that can be considered as different rules, which can change over time for each cell.

As we have seen, many different types of CA can be developed, and indeed, many have been developed due to the fact that CA-related research has experienced major growth over the recent decades. Such an active research rate makes the task of describing all the possible CA implementations difficult. Here it has been provided a non technical introduction aimed at non CA experts, such as musicians, and mainly centred on the CA implementations relevant to this thesis. For CA definitions and classifications using a mathematical approach, see the book (Adamatzky 1994). In an attempt to conceptually embrace all CA it is useful to cite a quotation by Richard J.

Gaylord and Kazume Nishidate, where they define the essential characteristic that a cellular automaton must possess –this quotation, in turn, quotes the book "Turtles, Termites and Traffic Jams" (Resnick 1994):

> 'The sine qua non of a CA model is that the behavior of the model is "determined not by some centralized authority but by local interactions among decentralized components"' (Gaylord et al. 1996 p.ix).

It is also beyond the scope of this section to give a description of the vast research activity on CA. However, it is worth mentioning a major theme of research, the characterisation of global CA dynamics, because it is relevant to the approach developed in this thesis, as we will see in Section 5.1, in which DSP analysis methods are proposed to characterise the CA behaviours. Research in this vein includes Wolfram's CA classification into four categories, the use of order/chaos measures such as entropy, and the use of Lyapunov exponents to characterise CA. For further reading see (Ganguly et al. 2003 p.5).

2.3 Interdisciplinary Connection between Cellular Automata and Sound Synthesis

To conclude this chapter, this section will present a number of CA properties that are useful for sound synthesis. This will highlight how these two fields are, or can be connected.

CA can be **Universal**. A system exhibits Universal Computation if it is capable of computing any computable operation. Over time, a number of CA have been proved to be universal, and this has encouraged researchers to use CA in a

variety of different application areas, including music technology. Note, however, that universality is by no means a necessary condition for obtaining interesting musical results. Christopher Ariza points out:

> 'the question of whether a CA functions as a proper machine, capable of universal computation or otherwise, has no significance on the extraction of artistically useful value streams' (Ariza 2007 p.29).

CA are **dynamical systems,** by the following definition:

> 'A system becomes a dynamical system when the current state of the variables comprising the system is at least in part dependent on previous states of the system. It is this dependence over time that defines the intrinsic dynamics of the system' (Boker 2007 p.132).

Also, space and time are discrete in CA, and quantities take on a finite set of discrete values. In this sense, CA are highly suitable for modelling sound and music (although this thesis focuses on sound), which are both fundamentally time-based and can be thought of as systems in which a finite set of discrete values (e.g., amplitudes, frequencies, musical notes, rhythms, etc.) evolve over time.

CA are **algorithmic**. An algorithm is defined as 'a sequence of unambiguous instructions for solving a problem, i.e., for obtaining a required output for any legitimate input in a finite amount of time.' (Levitin 2007). In sound synthesis, with algorithmic techniques it is possible to do pure synthesis, that is to say, to generate sounds from scratch as opposed to other types of synthesis that require pre-recorded material. This has some advantages in certain situations. For instance, in the video games industry there is a growing trend known as "procedural audio", which uses algorithmic techniques to generate sounds. The advantages of procedural audio are the reduction of disk access since it

constitutes a bottleneck in these systems. Also, procedural audio provides more variability of results as opposed to playing back always the same recorded sounds or music.

CA are **complex systems**. Typically, a complex system is composed of smaller interacting units which do not anticipate its emergent behaviour. Complex systems are relevant to sound synthesis because with few parameter specifications (which define the small units and their interconnections) we obtain complex evolutions. This is potentially useful for users of a synthesis instrument who do not have technical backgrounds, and therefore prefer not to deal with many technically complex parameters. In addition, the behaviour of complex systems is typically characterised by the emergence of self-organising patterns, and –as mentioned in the Introduction chapter– this allows the computer musician to explore new forms of organization. For a discussion and definitions of the previously mentioned key terms “emergence” and “self-organization”, see (Wolf et al. 2005).

Another key element is that CA **irreducibility** leads to **unpredictability**. CA irreducibility refers to the fact that it is not possible to ascertain in advance the value that a cell will hold after a number of generations. The only way to find out the value is to compute all the previous CA generations. In common situations, there is no shortcut in this process, and in this sense, the automaton is unpredictable. Such unpredictability has positive and negative implications in the sound synthesis domain. On the one hand, a level of unpredictability is accepted, and often desired, in systems for generating music and sound. An illustrative example is the following by Jacques Chareyron on the preference of

synthesis techniques offering a certain level of unpredictability, over other conventional techniques which are computationally reducible:

> ‘This characteristic of conventional synthesis techniques is a great advantage since it gives the composer complete control over the results. On the other hand, the sounds become easily recognizable and foreseeable, and it may be argued that they proclaim too loudly their mathematical origin and lack the degree of unpredictability the human mind expects from an artistic process’ (Chareyron 1990 p.37).

On the other hand, being under unpredictability conditions implies limited controllability, and a lack of a reasonable level of control restricts the music or sound design process.

Chapter 3 Previous Research on Sound Synthesis with CA

This chapter introduces previous work on the synthesis of sound using CA. The criteria adopted to select the research included in this review are as follows:

- The research is well documented in the literature of music technology.

- The research has had an impact on the field and is significantly cited in the literature.

- The research is globally accessible through the main channels of scientific distribution and not restricted to certain geographical zones. For instance, (Kreger 1997) seems to be relevant research presented in an Australasian conference, but unfortunately it is not accessible in Europe, even through the British Library.

Figure 13 shows the model that most of the synthesis techniques based on CA follow. Normally, a mapping process is used to connect the two main domains: the CA domain and the sound synthesis domain. In this way, CA are normally used for controlling the parameters of a synthesis engine over time. One of the few pieces of research that does not follow this model is (Toguchi et al. 2008) where the automaton itself is a physical model and so there is not an explicit mapping between two domains.

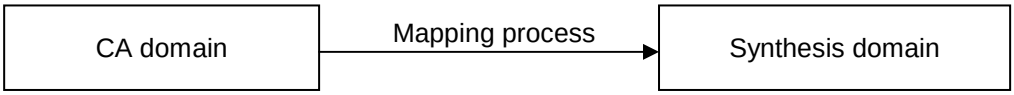

Figure 13: Generic model of CA-based synthesis techniques.

The following review, based on this model, focuses on the type of CA used, as well as on the mapping procedure for controlling the sound synthesis parameters. This chapter is organised into two sections. The first encompasses research which does not deal with the control of CA. The second section includes research that incorporates, to a greater or lesser extent, this issue.

For a review not restricted to sound synthesis, but also considering CA applications in the music domain, see (Burraston et al. 2005).

3.1 Research Not Involving CA Control

(Bowcott 1989) was one of the first attempts to synthesise sounds using CA. Bowcott was inspired by Xenakis' Markovian Stochastic Music theory (Xenakis 1992 p.43), in which a complex sound is described as a sequence of screens containing clouds of grains. Bowcott suggested that the 2D cellular automaton "Life" by John Horton Conway (Gardner 1970) could be a 'useful model to the automatic generation of such screens' (see Figure 14). His work also incorporated Truax's suggestion on the use of self-organised systems for the control of granular synthesis (Truax 1988 p.25).

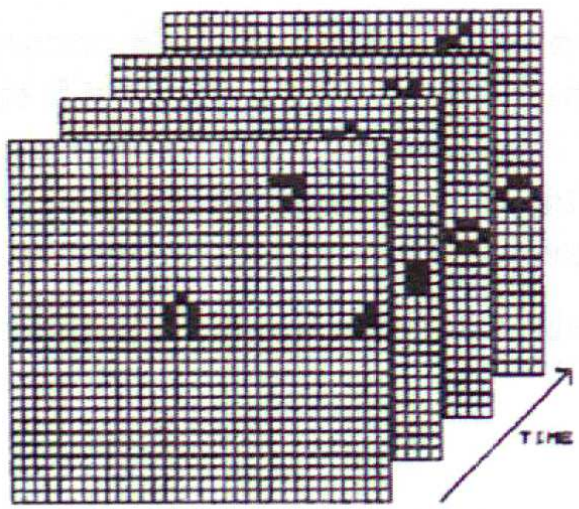

Figure 14: A number of generations of a “Life” evolution. (Bowcott 1989).

“Life” is a binary automaton, i.e. the cells can be either "dead" or "live". In Bowcott’s system ‘the evolution of these cells creates a list of time and x, y coordinates of the live cells’. Bowcott explored different mapping possibilities for these position parameters in the control of granular synthesis. He was concerned with certain mapping strategies, the results of which do not reflect their CA origins. According to him,

> ‘what should be sought after is a mapping which parallels the behaviour of the generating system in the musical domain, not a representation of the visual output’ (Bowcott 1989 p.57).

The mapping processes that he explored ranged from simple approaches where the x and y coordinates map onto frequency/amplitude pairs, to less obvious mappings in which the cellular space is seen ‘as a harmonic area of partials for one axis against predetermined frequencies in the other’ (Bowcott 1989 p.57).

Bowcott succeeded in synthesising ‘interesting timbres which varied in time’ but, as the author says, at the early stages of the project it was necessary more experimentation because of the ‘wide potential’ that the system provided.

In (Katrami et al. 1991) the use of digital images is investigated. These are obtained from either 1D binary cellular automata or fractal landscape algorithms. They are used as frequency filters of the analysis data of a Phase Vocoder (PV). As the authors explain in the paper,

> ‘The PV is an analysis-synthesis system that has an intermediate data, the time variant Discrete Fourier spectrum, of the input signal’.

The input signal for their system is a sound file. The explanation continues:

> ‘These intermediate data can then be used to resynthesise timbres at different pitches or different rates from the original. The analysis of sound is achieved by windowing the data. Within each time-window, the PV divides the spectrum into a number of equally spaced bands known as channels’ (Katrami et al. 1991 p.107).

In this system, the filtering process preceding the resynthesis is as follows: each time-window is filtered by an image line. Their PV has 512 channels and therefore ‘a fixed length of 512 pixels per image line is used’ (see Figure 15). In the case of binary CA image lines, a pixel with a value 0 ‘cuts entirely the corresponding channel’s amplitude while’ a pixel with a value 1 ‘leaves the channel’s amplitude unchanged’. The authors state that direct resynthesis can be achieved ‘by using a flat spectrum as the input to this process’, meaning that the system can perform a form of pure subtractive synthesis.

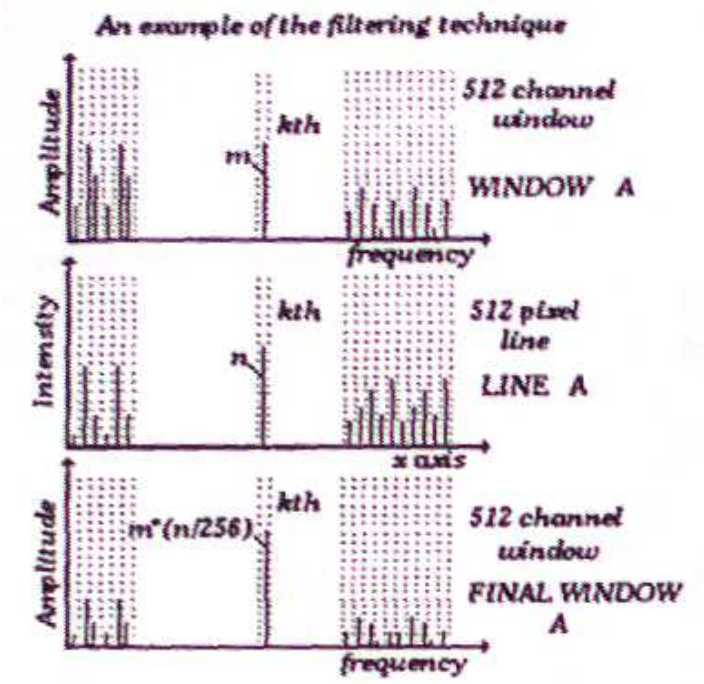

Figure 15: Filtering process in Katrami's system. (Katrami et al. 1991).

The "ca" system (Vaidhyanathan et al. 1999) uses a binary 1D cellular automaton. It considers only neighbourhoods of three adjacent cells and allows the use of 256 different rules to evolve the automaton. In this work,

> 'cellular automata are combined with granular processing. The sound sample is granulated and each grain is passed through a bank of filters' (Vaidhyanathan et al. 1999).

The user can adjust the bandwidths and center frequencies of the filters. An automaton of size 32 controls a set of 32 gain multipliers in order to modify the output of each filter over time. According to the authors:

> 'The main focus of ca is in transforming the harmonic structure of the grain. It deals with sampled sounds only and not synthesis of sound grains' (Vaidhyanathan et al. 1999).

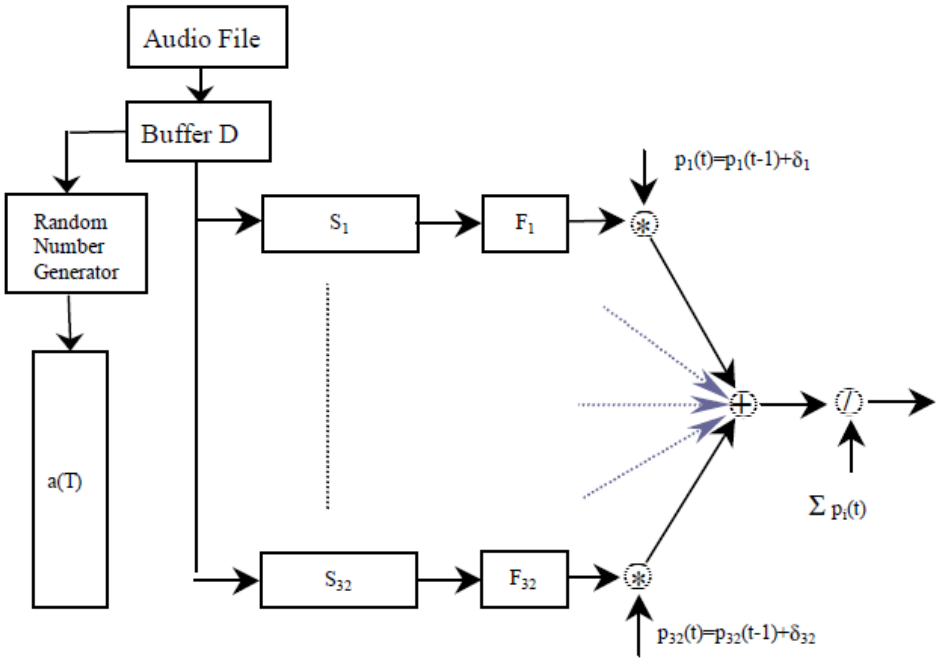

Figure 16: Block diagram of “ca”. (Vaidhyanathan et al. 1999).

Figure 16 shows a block diagram of “ca” providing a more detailed description of the system. The module denoted by **a(T)** represents the 32-cell automaton. Each of the cells controls the gain multipliers $\mathbf{p_i}$ which follow the filters $\mathbf{F_i}$, by determining the value of $\boldsymbol{\delta_i}$ in the expression $\mathbf{p_i(t) = p_i(t\text{-}1) + \delta_i}$, based on the rule shown in Figure 17.

$a_i^{(T)}$	$a_i^{(T+1)}$	δ_i
0	0	0
0	1	+Є
1	0	-Є
1	1	0

Figure 17: A table showing a rule of the “ca” system. (Vaidhyanathan et al. 1999).

T represents the fact that the automaton grows once in every **ω** samples, where **T = ω * t** and **Є = (1.0 – 0.0) / (ω/4)**.

(Comajuncosas 1998) describes a CA algorithm implemented in Csound. The synthesis system proposed utilises a 1D multi-state automaton with four cell

states. The automaton controls an additive synthesis through the following mapping procedure: the states of the automaton cells control the amplitudes of the oscillators (see Figure 18). The frequencies of the oscillators are left to be determined, and according to Comajuncosas:

> ‘A key feature of this design is the flexibility to map the frequencies of the oscillator bank’ (Comajuncosas 1998).

The user establishes the lowest frequency, corresponding to the first oscillator. The frequencies of the other oscillators are related to this fundamental and are calculated by means of mathematical functions. Using this design ‘a great variety of textures can be created with just a single’ automaton.

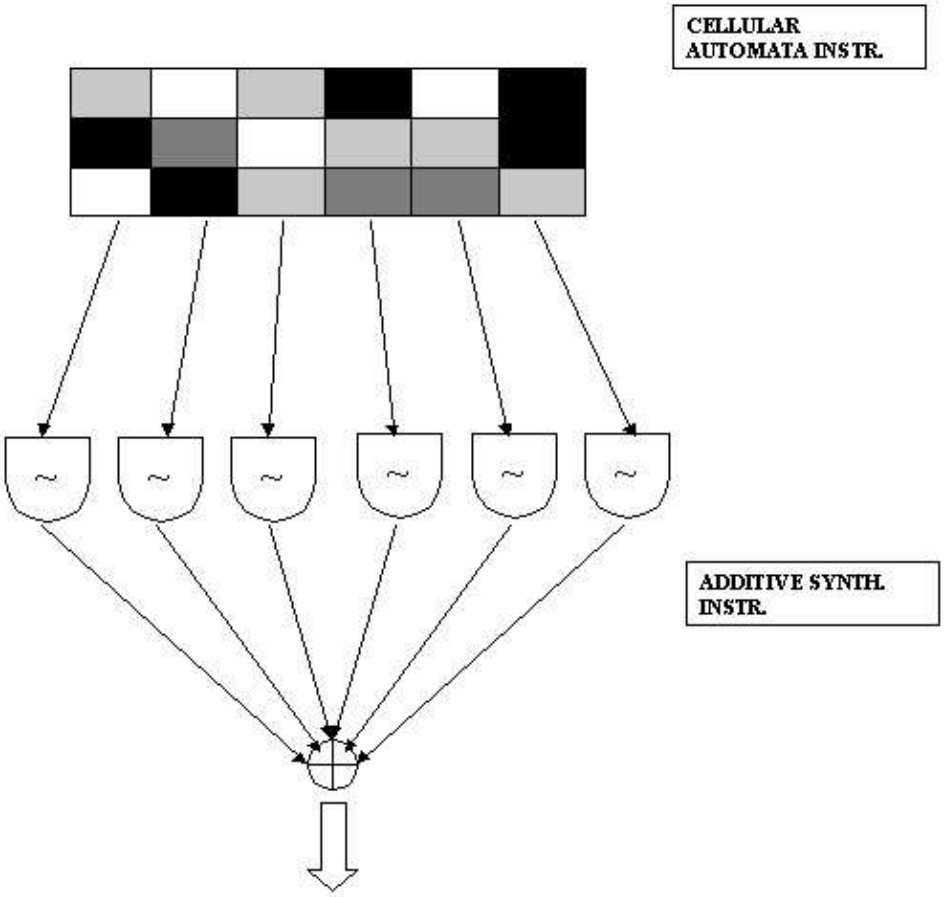

Figure 18: Mapping cell values onto oscillator amplitude values. (Comajuncosas 1998).

3.2 Research Dealing with CA Control

(Hunt et al. 1991) present the CA Workstation, a system developed to control music and sound parameters with 1D binary CA. The main focus is the exploration of the interrelationship between interesting visual patterns, and the sounds or music they can produce. The authors recognise that 'aesthetically pleasing visual patterns do not automatically lead to interesting sounds' (Hunt et al. 1991 p.167).

The paper describes different mapping methods. However, at the time of publication the system was in an early stage of developments regarding sound synthesis, with the authors providing only a few suggestions for the control of granular synthesis. Nevertheless, this work is of relevance to this thesis because it addresses the control of CA. Hunt et al. pointed out, referring to the problem of utilising CA algorithms:

> '... most of the time the composer is at the mercy of the algorithm –the music being automatically produced' (Hunt et al. 1991 p.165).

Thus, their CA Workstation 'aims to interrupt this automatic process to let the composer interact directly with the CA at a high-level and thus to make significant artistic decisions' (Hunt et al. 1991 p.165). Their system allows the composer to start and stop the automaton and change its parameters at any point. The composer is also able to zoom in and select areas containing visually interesting patterns which can then be used to control different music or sound parameters.

An example of how this system can work in the MIDI domain is now detailed. Figure 19 shows a screen containing a CA visual pattern. Then,

> 'the composer allocates a pitch to each horizontal cell position across the screen. The note corresponding to that pitch will be played for as long as the CA has an active (black) cell at the current generation. This process can be visualised to be a moving line which triggers notes (according to the position of the active cells) as it passes vertically over the pattern' (Hunt et al. 1991 p.166).

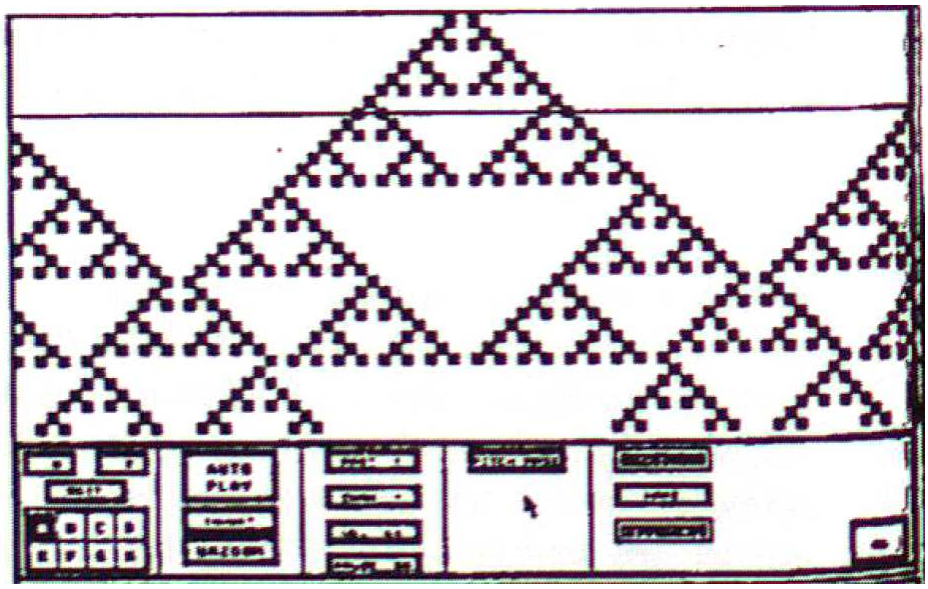

Figure 19: CA Workstation screenshot. (Hunt et al. 1991).

Chaosynth (Miranda 1995) is a granular synthesizer. It is based on ChaOs (Chemical Oscillator), a 2D multi-state automaton devised by Miranda as a different version of the hodge podge machine (Gerhardt et al. 1989). As the author explains:

> 'ChaOs […] is a metaphorical model of a neurophysiological phenomenon known as a neural reverberatory circuit […]. ChaOs can be thought of as an array of identical electronic circuits called nerve cells. At any moment, a nerve cell can be in a quiescent, depolarized or burned state. The array tends to evolve from an initial random distribution of these three states in the grid to an oscillatory cycle of wave patterns' (Miranda 1995 p.297).

Chaosynth dynamically controls the frequency content of the sound grains in the following manner: the states of the automaton cells are associated with frequency values (set by the user), and oscillators are associated with multiple cell groups as shown in Figure 20. Then at each iteration of the automaton, each oscillator produces a sinewave with a frequency determined by the arithmetic mean over the frequency values of the corresponding cells. Each grain is finally produced by additive synthesis of these sinewaves. The amplitudes of each oscillator are set by the user. With this mapping, Miranda managed to capture the behaviour of ChaOs which 'resembles the way in which most of the natural sounds produced by acoustic instruments evolve: sounds tend to converge from a wide distribution of their partials (for example, noise) to oscillatory patterns (for example, a sustained tone)' (Miranda 1995 p.297). As a result, Chaosynth produces a wide variety of novel granular sounds that sound pleasant.

Miranda's work is concerned with the sound design possibilities of Chaosynth from a control perspective, as the following quotations demonstrate:

> 'The user also specifies the dimension of the grid, the amount of oscillators, the allocation of nerve cells to oscillators and the parameters of ChaOs (that is, the number of states, the resistances of the potential divider, the capacitance of cells and the number of iterations)' (Miranda 1995 p.298).

> 'Variations in tone color are achieved by varying the frequency values and the amount of nerve cells per oscillator' (Miranda 1995 p.299).

Regarding the specific control of his automaton, Miranda established that 'different rates of transition from noise to oscillatory patterns' are obtainable by changing the values of the rule parameters.

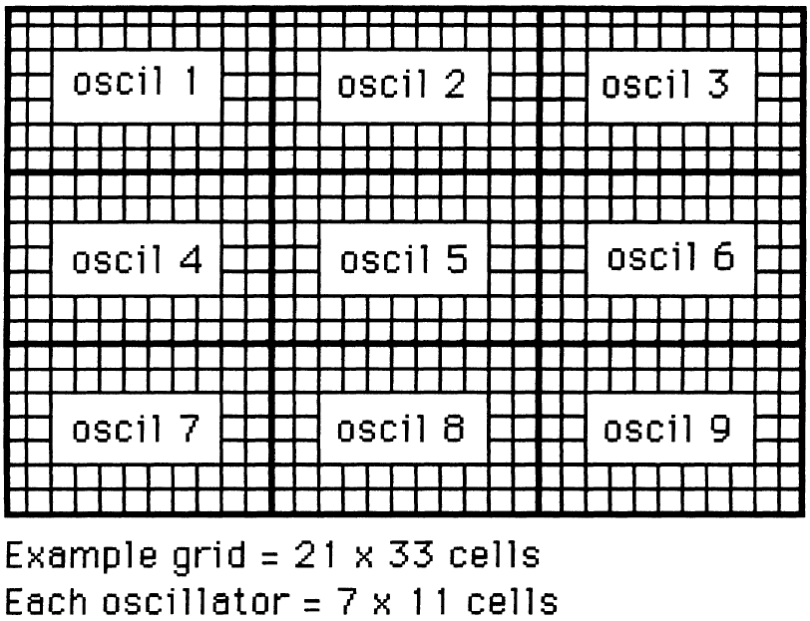

Figure 20: An example of a distribution of oscillators in Chaosynth. (Miranda 1995).

The last system reviewed here is LASy (Linear Automata Synthesis) (Chareyron 1990). The mapping procedure is described as follows:

> ‘A wavetable stored in memory is seen as a linear automaton: i.e., every element (sample) of the wavetable is a cell of the automaton. The value of an element of the waveform is the state of the corresponding cell of the automaton. Generation after generation, the state of every cell is computed according to the transition rule of the automaton, and these values are used to feed a Digital-to-Analog Converter (DAC) for sound output’ (Chareyron 1990 p.27).

Thus, LASy is a synthesis technique where CA are seen as Digital Signal Processors. A sound sample is received as input data and processed according to the transition rule. (As Chareyron says: ‘The LASy process may be seen as the operation of a digital filter over a time-limited signal’). The modified sound becomes the output of the system as well as the input for the next iteration. This process is repeated over and over and ‘presents an analogy to the Karplus-Strong algorithm’. Figure 21 shows a block diagram of the LASy algorithm.

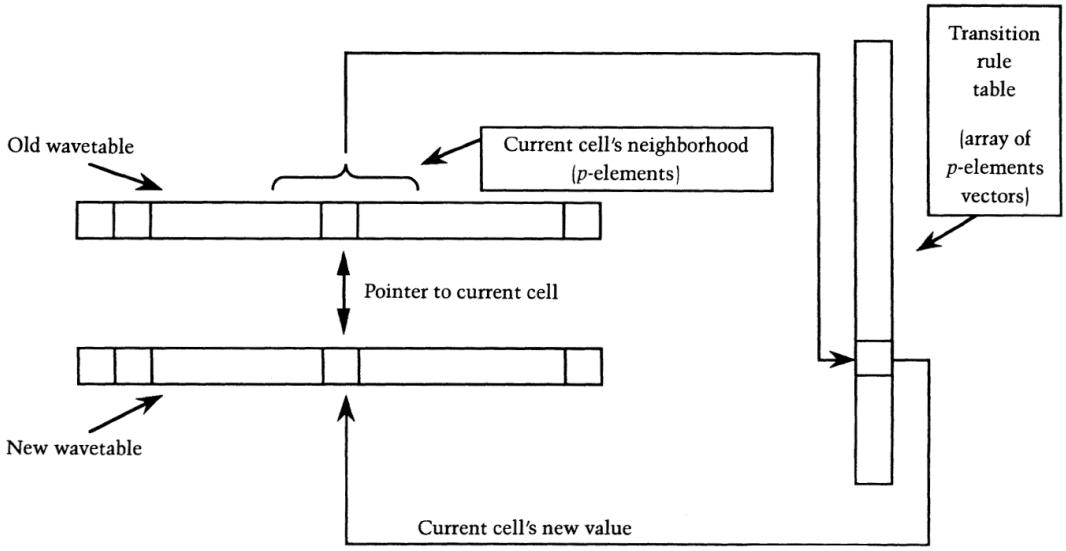

Figure 21: Implementation of the LASy algorithm. (Chareyron 1990).

As we have seen LASy works with 1D multi-state CA with the number of states determined by the bit depth of the DAC. Therefore, potentially interesting is the large number of possible transition rules that can be defined. However, on the downside, such a large number of possible states makes the definition of the transition rules a complex task. According to Chareyron 'it is difficult to extract the ones capable of producing useful results' (Chareyron 1990 p.28). Chareyron proposes the use of parameterised functions to generate the new cell values of the transition-rule table. With this approach, different rules can be created from each defined function by assigning different values to the parameters. According to Chareyron:

> 'Because theoretical results about the evolution of the sound are not available for most of the automata, this evolution may only be predicted using empirical methods. If the transition rule is characterized by a limited number of parameters, it will be possible to establish relationships between the variation of every parameter and the evolution of the harmonic spectrum of the sound output. The results of such a study will allow better control over the results of the synthesis' (Chareyron 1990 p.31).

Regarding this controllability, (Chareyron 1990) presents a case study of a simple parameterised linear function from which, depending on the values of the parameters, one can obtain:

- ‘sounds whose amplitude and spectral complexity grow in time’ and ‘the opposite behaviour’ (Chareyron 1990 p.32). ‘In such cases, we may use two (or more) transition rules in succession’ (Chareyron 1990 p.37). Such successions could be controlled by the user in real-time by means of note-on and note-off signals or messages.

- ‘transition rules whose actions may be seen as a linear filter operating upon the original table’ (Chareyron 1990 p.32). Such a filter was theoretically analysed.

- ‘a variation of the Karplus-Strong algorithm’ (Chareyron 1990 p.32).

According to the above paper, a great variety of sound evolutions, including more complex ones, can be achieved through: more complex functions (including those leading to non-linear filters), a combination of functions, by manually introducing modifications to the final table, and through a good choice of the initial waveform.

Chapter 4 Problem Statement

The aim of this thesis is to develop a sound synthesis technique based on CA capable of allowing a sound design process. By “sound design process” we shall understand a control over the synthesis process, on the part of the user, towards a satisfactory result.

The main criticism regarding sound synthesis techniques based on CA concerns the sound design limitations experienced by users. For instance, Rocha Iturbide points out, referring to CA:

> ‘... le danger d’utiliser ce genre de processus dynamiques automatiques complexes, pour le contrôle de la synthèse, est que nous tendons à limiter le contrôle sur eux, et qu’à la fin, ils deviennent plus importants que nous, qui devenons des êtres passifs se contentant d’observer le résultat’ (Rocha-Iturbide 1999).

Which translated means: the danger of using this kind of complex automatic dynamic processes for the control of synthesis is that we tend to limit the control over them and in the end, they become more important than ourselves, who become passive beings satisfied with the observation of the results.

We can also gather from this quotation that the control of CA is a key aspect to address to allow the widening of the sound design possibilities. As we saw, the research reviewed in Section 3.1 (Research Not Involving CA Control), allow the user, at most, to establish the values of certain synthesis parameters, and others are controlled over time by the automaton. With this approach such systems produce a wide variety of sounds. However, the sound design possibilities are limited by the fact that the control of the CA is not addressed.

All of this research certainly falls within the category of systems critiqued by Rocha Iturbide.

Only a handful of systems, those presented in Section 3.2 (Research Dealing with CA Control), have shown concern on the sound design problem taking into account the control of CA. However, there is room for more research along these lines as the following analysis will elucidate. For instance, in (Hunt et al. 1991) the issue of gaining control over the evolution of CA is addressed, but the work did not go into the synthesis of sounds in any depth. Whilst it provides specific mapping procedures in the MIDI domain, it does not provide any concrete mapping method in the synthesis domain. In the absence of further published work on sound synthesis, the potential of this approach for sound design remains to be seen.

Regarding Chaosynth, Miranda admitted that ‘more research is needed to gain a better understanding of the role of Chaosynth’s parameters’ (Miranda 1995 p.299). In order to alleviate the problem of exploring the wide range of new sounds that Chaosynth can produce, Miranda and other experienced users of his system carried out exhaustive work on the classification of the Chaosynth sounds (Miranda et al. 2000). They created a taxonomy with five general sound classes: Fixed Mass, Flow, Chaotic, Explosive and General Textures. In addition, different subclasses were defined for each of these types of sounds. The classification was made largely based on the spectral evolution of the sounds, which is the most powerful characteristic feature of Chaosynth. For instance, whilst the Fixed Mass class refers to sounds with ‘a stable and steady spectrum where the frequencies of the grains are kept within a fixed band’

(Miranda et al. 2000 p.99), the Flow class refers to sounds in which 'the frequencies of the granules tend to move collectively in one direction' (Miranda et al. 2000 p.99).

However, the values the cells of ChaOs adopt over time are unpredictable. Consequently the frequency values computed by Chaosynth, which create the spectral evolution, are unpredictable as well. This seems to be a handicap to determining the role of Chaosynth's parameters in general and of ChaOs' parameters in particular in the sound design process.

As was mentioned in the previous chapter, LASy (Chareyron 1990) is a synthesis technique that provides a certain level of control. As we saw, Chareyron presented a case study from which we could obtain with certain level of prediction three types of sound evolution. In addition, we saw that according to his paper more variety of sound evolution could be obtained. But these varieties of evolution were not studied in the paper in terms of prediction/control. Nevertheless, Chareyron stated the following:

> 'From personal experience, I would say that roughly 50 percent of the new automata I build generate sounds with the characteristics I expected from them, while about 10 percent of the automata give unexpected but interesting results […] (the other ones should only be forgotten)' (Chareyron 1990 p.38).

Finally, the paper does not provide information on the control of an important sound attribute, the pitch.

From this critical analysis of previous work one can identify that in order to develop a sound synthesis technique capable of allowing a sound design process, more research is needed regarding the control of CA. Furthermore,

apart from the control of CA, another capability that such desired synthesis technique should demonstrate is the possibility of producing a variety of sounds. In this respect, based on the generic model shown in Figure 13, the following criteria are established in order to estimate the potential of a system for producing a variety of sounds:

- Variety of CA used: Does the technique use only one automaton or can many types be used? We would expect that the more CA rules available the more variability there can be in the resulting sounds.

- Variety of mappings used: Does the technique work with only one mapping process or can it use many? We would expect that the more mappings that can be used, the more variability there will be in the resulting sounds.

- Variety of synthesis engines used: Does it use only one synthesis engine or can it use many? Again, we can expect that the more engines that can be used, the more variability in the sounds produced.

- Type of synthesis: Is it pure synthesis or/and audio processing? By incorporating both these types of synthesis, the results should provide more variability.

Through reviewing previous work on the basis of these criteria one can identify that none of them addresses the use of a variety of synthesis engines. A lack of a variety of engines has, therefore, as a consequence, a lack of a variety in the type of synthesis used. Nonetheless, those systems that are based on audio processing (usually filtering) can always perform pure synthesis, provided that they can generate excitation signals such as noise, or other complex waveforms. However, only Katrami's research explicitly reported on the implementation of this possibility, to certain extent, as we saw in Chapter 3. In other respects, only a few systems –the CA Workstation and Bowcott's research– consider or suggest more than one mapping procedure. Miranda expresses the desirability of mapping variety as follows:

> 'We are, however, aware that instead of providing a system that uses only one mapping possibility, we should provide the means for user-specification of other mapping possibilities' (Miranda 1995 p.299).

Finally, most of the reviewed systems make use of a variety of CA rules. Note that in the case of Chaosynth, even though only one cellular automaton is used, it can produce different behaviours depending on the value of its rule parameters. Therefore, in this sense, it can be considered equivalent to having access to a variety of CA.

The following table summarises the critical analysis undertaken in this chapter by presenting the different systems in relation to the adopted criteria. In "Chapter 6: Conclusions" these criteria will be also used to evaluate Histogram Mapping Synthesis, the technique to be discussed in this thesis.

	Variety of CA	Variety of mappings	Variety of engines	Type of synthesis	Control / prediction
Bowcott's system	No	Yes	No	Pure synthesis	Limited
Katrami's system	Yes	No	No	Both	Limited
ca	Yes	No	No	Audio processing	Limited
Comajun-cosas' system	Yes	No	No	Pure synthesis	Limited
CA Workstation	Yes	Yes	No	Pure synthesis *	In MIDI domain
Chaosynth	No **	No	No	Pure synthesis	Low
LASy	Yes	No	No	Audio processing	Certain

* Only suggested
** Variety achieved through different behaviours.

Table 1: Previous systems in relation to defined criteria.

In conclusion, as it has been shown in this chapter, existing sound synthesis techniques based on CA have involved rather limited possibilities for sound design. Although CA can allow the production of a rich variety of dynamic sounds, due largely to these sound-design limitations its practical use demands significant experimentation, and hence relatively few musicians, usually those within a research environment, use them. Dealing with this problem in order to allow a reasonable level of control over the synthesis will enable professionals (composers, sound designers, live performers, etc.) to utilise CA as a tool in their work so they may benefit from the rich potential. Full control over generative processes like CA is not possible, but is –conceptually speaking– not desired. Also note that "control" and "prediction" are two concepts that go hand in hand. For instance, we have seen with LASy that the prediction of the CA

outcome is an important aspect for having the synthesis process under control. In general terms, the more unpredictable a generative process is, the more difficult it will be to control it. A research activity seeking a good balance between the predictable and the unpredictable seems to be a sensible direction. The achievement of such goals, combined with technological improvements in terms of computational capacity, could then attract industry interest so that a new generation of “natural” synthesizers can be developed –Nature is never under total control.

Chapter 5 Histogram Mapping Synthesis

5.1 Problem Analysis and DSP Approach

At this point, the problem that this thesis tackles is twofold. Firstly, in order to develop a sound synthesis technique based on CA, it is necessary to devise a mapping procedure. Secondly, in order to allow a sound design process it is necessary to provide control mechanisms for the synthesis technique.

Regarding the design of control mechanisms we have seen that the control of CA is a complex but crucial aspect of research to be addressed. It is difficult to make generalizations because different CA offer different control possibilities. Also, for the same automaton, the devised control mechanisms may differ depending on the mapping used. Nevertheless, I identify two general problems to be faced, common to virtually all the automata, which stem from the nature of CA:

- The autonomous evolution of CA: An automaton is defined as a device operating under its own hidden power. The term is derived from the Greek "automatos" which means acting of one’s own will or independently, self-acting, self-moved.

- Irreducibility and unpredictability (discussed in Section 2.3): Under unpredictability conditions the design of control mechanisms is hindered.

The approach adopted to deal with these issues will be explained below. Let us now examine the mapping problem. Designing a mapping procedure is a trial-and-error endeavour open to the most arbitrary solutions. However, mappings producing interesting musical results are not straightforward to obtain. This thesis avoided a critical analysis of the existing mapping methods because it is generally agreed the difficulty of evaluation of the quality of a given mapping procedure. Although in the music domain, (Kirke et al. 2007) and (Beyls 2004) are studies dealing with this issue.

As a first step towards a mapping design, noteworthy is the practice of some authors who manage to identify analogies between CA behaviours and sonic phenomena or musical processes. Then, they try to design a suitable mapping process which captures or reflects these analogies. For instance, Bowcott established the musical aims of his work on the basis of this kind of analogies as follows:

> ‘When one observes “Life” within a graphical environment, one can become very familiar with a number of configurations that appear: if these configurations and their evolution can be perceived as sound objects, the process could be a useful asset to the composer’ (Bowcott 1989 p.56).

However, the achievement of his goal has not been demonstrated in the literature. Miranda’s work is a successful example of this methodology. As we saw in Chapter 3, he managed to capture the behaviour of ChaOs which resembles the way in which most sounds of acoustic instruments are produced. He also adopted this strategy, in the music domain, for his system CAMUS. For this he ‘devised a number of experiments in order to study whether cellular

automata that exhibit pattern propagation behaviour could be used or adapted to model the propagation of musical patterns' (Miranda 2001).

In order to identify these analogies between CA behaviours and sonic phenomena an important question arises: what type of information is available to us in order to characterise the CA behaviours? Bowcott and Miranda analysed them visually, which is a very limited type of immediate analysis. It can also be unreliable or misleading; note that the visual display of CA is the result of an elementary mapping: cell values are normally assigned to grey levels or different colours and, for instance, with different palettes of colours the same output can be displayed in many different ways, especially in the case of multi-state CA.

A similar situation to the abovementioned limitation is found in the analysis of sounds. By displaying the waveform we can perceive some basic information such as its degree of timbral complexity, but it is through Fourier analysis that we obtain exhaustive information about its frequency content. Analogously, a DSP analysis of CA evolutions could reveal information beyond what can be perceived from their visual display.

Therefore, this thesis proposes to consider the output of CA as digital signals and to use DSP procedures to analyse them. This approach opens a great variety of possibilities to better understand the self-organization process of CA, with the view to identifying:

A) Analogies from the sound synthesis domain: if certain analogies are found we would be more likely to design "natural" mappings, in contrast with more artificial or arbitrary mappings designed when there are not clear correspondences between CA behaviours and sonic phenomena. Even though no clear analogies are found, the output of such a DSP analysis may provide useful information about the CA evolution that could be the inspiration for other musical applications.

B) Possibilities of CA control for sound design with respect to the two main problems stated above:

- With a DSP analysis it may be possible to monitor the CA evolution, being therefore in a better position to make decisions on when (and how, as stated in the following paragraph) to alter the autonomous evolution.

- With a DSP analysis it may also be possible to better characterise not only the CA behaviours, but the behavioural effects caused to the CA by the control mechanisms that may be devised. That would help to establish cause-effect relationships which could bring a degree of predictability regarding the type of sounds that we could obtain.

These are merely points of departure based on initial intuitions for dealing with main general problems. This thesis will demonstrate that this approach is suitable not only for dealing with these, but also with other specific research problems, depending on the automaton used.

To the best of my knowledge this approach has never been pursued before in order to develop methods for sound synthesis based on cellular automata. Histogram Mapping Synthesis is a result of such an approach.

5.2 Histogram Mapping Synthesis (HMS)

Histogram Mapping Synthesis is a new synthesis technique based on a statistical analysis of CA evolution. The functioning of a two-dimensional multi-state automaton is considered as a sequence of digital images, each of which is analysed by means of histogram measurements. Such a DSP analysis gives a histogram sequence that can be displayed as a 3D plot.

The histogram of a grey-level digital image is a graphical representation of the number of occurrences of each grey level[1] in the image. By dividing the number of occurrences by the total number of pixels of the image, the histogram is normalised and expressed in probabilistic terms giving an estimate of the probability of occurrence of each grey level in the image, i.e., it is estimated the probability distribution of the image values (Pratt 2007 p.160). Expressed in mathematical form, according to (Gonzalez et al. 1993):

> 'The histogram of a digital image with grey levels in the range [0, *L* - 1] is a discrete function

$$p(r_k) = n_k / n \tag{2}$$

[1] Apart from the definitions stated in this section, this thesis may refer to colours instead of grey levels because CA are usually displayed using a palette of different colours.

where r_k is the kth grey level, n_k is the number of pixels in the image with that grey level, n is the total numbers of pixels in the image, and k = 0, 1, 2, …, L – 1' (Gonzalez et al. 1993 p.171).

It is important to note that the sum of all the histogram bins will be equal to one:

$$\sum p(r_k) = 1 \tag{3}$$

It is also important to have in consideration the weight with which every pixel contributes to the normalised histogram. It inversely depends on the image size as follows:

$$w = 1 / n \tag{4}$$

meaning that a change of colour of a pixel in a small image will cause a larger variation or a more abrupt change between the corresponding pair of histograms than in the case of a larger image.

Let us go back to the histogram analysis of CA evolution leading to 3D histogram sequence plots. Initially, this kind of analysis was chosen for two main reasons. Firstly, in the histogram sequences the values adopted over time by the cells of multi-state CA are registered and displayed in a meaningful way. This is especially important when a large number of cell states are brought into play, whose values escape the visual perception. This kind of analysis opens up a new dimension of the self-organization process of CA. Delving into this new dimension provides unique insights that complement the spatial information gathered from the CA visual output.

Secondly, one can think of a certain analogy between histogram sequences and spectrograms. Firstly, their z-axes both correspond to the time domain. Secondly the x-axis of the histogram sequence, which represents the automaton's cell values, parallels the frequency domain represented by the x-axis of the spectrogram (in other words, there is a parallelism between the bins of the histogram sequence and the spectrogram's bins). Finally, the probability scale on the y-axis of the histogram sequence parallels the spectral magnitude represented by the spectrogram's y-axis. Note that this analogy is formulated here strictly between two types of graphical representations. Any analogy between the concepts behind these graphs would be more difficult to state.

Continuing the theme of the abovementioned analogy, this thesis documents that from an appropriate automaton, in the histogram sequences, which can be seen as temporal structures, it is possible to identify time-varying structural elements resembling spectral components of a sound such as sinusoidal components, noise components, transients, spectral envelopes and formants. With these structural elements we can design the time-varying frequency content of sounds; we can build spectrograms. These spectrograms can be rendered into sounds using different synthesis techniques.

Seen from another perspective, these structural elements can be used as control data to drive different synthesis techniques. Depending on the resemblance of these structural elements, different mappings onto appropriate synthesis parameters can be established. For instance, structural elements resembling partials can control additive synthesis through controlling the time-varying amplitudes and frequencies of sinusoidal oscillators. This process can

be seen as one that creates spectrograms out of histograms. In other cases, histogram sequences resembling sequences of spectral envelopes can control subtractive synthesis by controlling the time-varying magnitude response, or filter shape, of a time-varying filter. If the sound source is, for instance, a white noise generator, this process can still be seen as building spectrograms out of histograms. However, this is not the case if the sound source is, for instance, a sawtooth wave oscillator, because the harmonics introduced by such an oscillator are not initially represented in the histogram sequence.

Such possible mappings based on such resemblances make HMS distinctive; in most other techniques there is not an intuitive correspondence between the components of an automaton and the components of a sound.

5.2.1 Control

This subsection presents an important feature of HMS, namely, a control possibility for sound design, which can be generally applied regardless of the automaton used. Nevertheless, it should be noted that most of the control possibilities that HMS can offer over the synthesis depend to a significant extent on the automaton used. Hence, a number of control mechanisms will be presented during the course of this thesis. In this subsection it will also be helpful to re-examine, from the HMS perspective, the two main CA control problems highlighted in Section 5.1: the autonomous evolution and the unpredictability of CA.

Regarding the unpredictability of CA, unlike most other techniques, HMS does not map directly from the CA cells. These cells are the elements that introduce the unpredictability. Instead, the mapping processes are created from histogram sequences (Figure 22). Therefore, with those CA for which the histogram sequences are a means of characterising their behaviours and, whose behaviours are predictable by their parameter values, then by designing mappings after a histogram analysis it is largely avoided the unpredictability of cell values in favour of the predictability of behaviours.

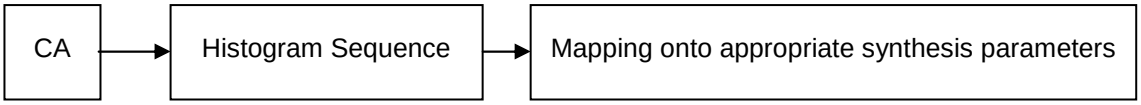

Figure 22: Mapping process from histograms instead of directly from CA.

Regarding the autonomous evolution of CA, it is only when a specific automaton is under study when potential mechanisms for controlling its evolution can be identified. Nevertheless, closely related to the autonomous evolution of CA is the automatic character that CA imprint on most synthesis techniques, something that reduces the possibilities of an offline sound design process.

In this respect, a key control possibility of HMS is that we can make decisions before rendering the sounds; we can develop an offline sound design process from the structural elements of the histogram sequences. Virtually infinite sounds can be designed from a single histogram sequence, by manipulating its structural elements and controlling the way they are mapped onto the sound synthesis domain. The fact of being operating in the frequency domain facilitates the control over all the sound attributes: intensity, duration, pitch and timbre. This is accomplished by means of spectral and mapping processes

applied to the structural elements of the histogram sequences; for example the assignment of frequencies, amplitude modifications, pitch shifting, time stretching, etc.

The possibility of designing a timbre and being able to reproduce it with different pitches, intensities and durations, constitutes a special case of sound design, namely, instrument design. Instrument design can be boosted with specific CA in which it is possible to predict the location and characteristics of the structural elements within histogram sequences resulting from different CA runs. With this level of prediction, the automatic characteristic introduced by CA is no longer a problem, but an advantage. Once a type of sound has been designed offline, it will be possible to automatically reproduce multiple instances of the original sound which will be similar but not identical; very much like the sounds produced by acoustic instruments. Of course, for each of these sound instances it will be possible to change some of the sound attributes, such as the pitch, duration and intensity, whilst largely preserving the originally designed timbre.

5.2.2 Modes of Use

Having introduced various ways of synthesising sound with HMS, let us study in detail the different modes of use of HMS.

Manual vs. Automatic Modes: these refer to the mapping process between the CA domain and the synthesis domain. In Manual Mode there is user involvement at the point between the CA run and the synthesis. After the CA run, an offline sound design process can be developed by manipulating the

histogram sequence, establishing which synthesis parameters will be controlled by which structural elements of the histogram sequence, and setting the values of the rest of the synthesis parameters. In Automatic Mode there is not user involvement at the point between the CA run and the sound synthesis. Usability depends on the level of predictability. With predictable histogram sequences, in the best case scenario, any process previously developed in Manual Mode can be automatically reproduced. Hence, HMS can automatically produce different sound instances of a type of sound previously designed, behaving like an instrument. Some synthesis parameters can be established beforehand for each instance (for example, to control the pitch). On the other hand, with a low level of predictability, even in the worst case scenario it is always possible to render sounds with random procedures involved in the mapping process.

Real-Time vs. Non-Real-Time Modes: these refer to whether the sound is rendered while the automaton is evolving or after the run is completed. In Real-Time Mode the sound is synthesised as the automaton is evolving. It has not been implemented yet, although it is potentially possible. In Non-Real-Time Mode the sound is heard after the CA run is completed.

Finally, **Interactive vs. Non-Interactive Modes** refer to whether there is user control over some parameters (of the automaton or the synthesis engine) on-the-fly, i.e. while the sound is being rendered. The Interactive mode has not been implemented yet, although it is potentially possible.

The relationships between these modes are as follows: Manual Mode is Non-Real-Time. Automatic Mode can be in Real-Time or Non-Real-Time Modes

(regarding the latter, think for instance of Csound type of applications). Real-Time Mode is Automatic. Non-Real-Time Mode embraces Automatic and Manual. Interactive and Non-Interactive Modes can be in Real-Time and Non-Real-Time, as well as in Manual and Automatic Modes.

5.2.3 The Search for Automata

The search for CA that are suitable for Histogram Mapping Synthesis is largely a matter of experimentation. Traditionally, an automaton was considered potentially interesting for sound synthesis after a visual analysis of its behaviour. Within the context of HMS, an automaton can be considered interesting depending on the results obtained from a histogram analysis of its evolutions.

The results of a histogram analysis can be more or less expectable depending on the automaton in question. For instance, as we will see, in the case of the multitype voter model (studied in Section 5.3) the main characteristics of its histogram sequences are, in general terms, as one might expect provided that one understands the model's behaviour. On the other hand, the histogram sequence structural elements of some behaviours of the hodge podge automaton (studied in Section 5.4) are by no means possible to expect in advance from its visual perception. It is the DSP analysis which reveals such useful information. It is worth mentioning that, although the first experiments with HMS resulted in essentially flat histogram sequences, the first successful results were these kinds of unexpected structures from the hodge podge machine. It was a discovery in its own right.

There is a requirement that every automaton must satisfy: it must be multi-state so that its corresponding histogram's x-axis is large enough so as to accommodate enough structural elements for the synthesis of complex sounds.

Regarding controllability, those CA which are defined by rules that include a number of parameters and whose behaviours are predictable from their parameter values are more likely to allow controllability. For instance, if different parameter settings lead to different behaviours, then it is more likely that the different types of sounds that can be obtained from them are predictable, and therefore, under control. Also, in general, by modifying the CA parameters it may be possible to modify the behaviours, as far as possible, as appropriate to our needs. In other respects, those CA evolving from a random distribution of cell values are of special interest for instrument design; the fact of being able to start from different random configurations may favour the generation of similar but not identical evolutions, leading to similar but not identical sounds.

This thesis focuses on two-dimensional CA, however, HMS can be produced using CA grids of different dimensions, and also with other computational models. Generally speaking, it is possible to obtain histogram sequences from virtually any stream of numerical information. However, it is my belief that self-organised systems are the most fruitful for producing musically interesting histogram sequences. Self-organised systems can be modelled using means other than CA, for instance, differential equations and Monte Carlo simulations (Camazine et al. 2003 pp.72-73).

In the following sections HMS will be studied with four different CA, specifically, the multitype voter model (Section 5.3), the hodge podge machine (Section 5.4), the plurality vote rule (Subsection 5.6.1) and the rug rule (Subsection 5.6.2). All these CA have been implemented in Matlab on the basis of the existing model descriptions. However, some original modifications have been explored as explained below for each case. These CA will be studied in relation to the mapping possibilities they can offer. Also, a number of control mechanisms will be presented. All the mappings and control mechanisms are original contributions of this thesis.

5.3 HMS with the Multitype Voter Model

5.3.1 The Multitype Voter Model

In 1953, geneticist Kimura introduced the stepping stone model (Kimura 1953). This process was broadly studied by other geneticists over twenty years before being rediscovered by probability theorists Clifford and Sudbury in 1973 (Clifford et al. 1973) being named the invasion process, and by Holley and Liggett in 1975 (Holley et al. 1975) under the name of the voter model (Cox et al. 2002 p.1348).

Nowadays, the voter model is considered one of the standard models of interacting particle systems (Liggett 1985). It can be interpreted as a model of opinion formation. A collection of individuals is defined,

> ‘each of which has one of two possible positions on a political issue. These possible positions are denoted by 0 and 1. Periodically, […] an individual reassesses his view in a rather simple way: he chooses a “friend” at random with certain probabilities and then adopts his position’ (Liggett 1985 p.226).

It can be also viewed as a model of competition. This interpretation is clear from the point of view of the invasion process:

> ‘Two species compete for territory along their mutual boundary […] and the result of conflict is the invasion by one of the species of territory held by the other’ (Clifford et al. 1973 p.581).

When the voter model is generalized to more than two opinions it is known as the multitype voter model (MVM).

The MVM can be simulated by means of a two-dimensional probabilistic cellular automaton. I have implemented the following transition rule on the basis of the stepping stone implementation found in (Fisch et al. 1996): at every time step, a coin with prescribed Update Probability of tails is tossed for each cell in the grid. If the coin comes up heads, the state of the given cell will be replaced by the state of one of its neighbours, selected uniformly at random from the specified neighbourhood. The coin-tossing process can be implemented by defining such a prescribed Update Probability as a CA input parameter. It will be a number chosen by the user, between 0 and 1. Then, the coin will come up heads if a random number between 0 and 1 generated for each cell and time step is higher than the Update Probability. Regarding the neighbourhood type, most of the experiments have been carried out considering the four orthogonally adjacent cells, namely, north, east, south and west. (It is also possible to consider all the eight adjacent cells by including the north-east, north-west, south-east and south-west). The MVM has been implemented with periodic boundary conditions, i.e., with a torus shape.

Starting from a uniform random distribution of cell values, or colours, as the initial configuration, this automaton self-organizes into clusters or areas of a single colour. See Figure 23 and Video Example 1 in Appendix B. As the rule is iterated, some areas will increase in size while others will decrease to the extent that they can disappear. In the end, one colour will prevail over the others when, according to the voter interpretation, consensus occurs (see Video Example 2). Note that the random inputs and probabilities in the rule make that different runs with the same settings will result in different evolutions.

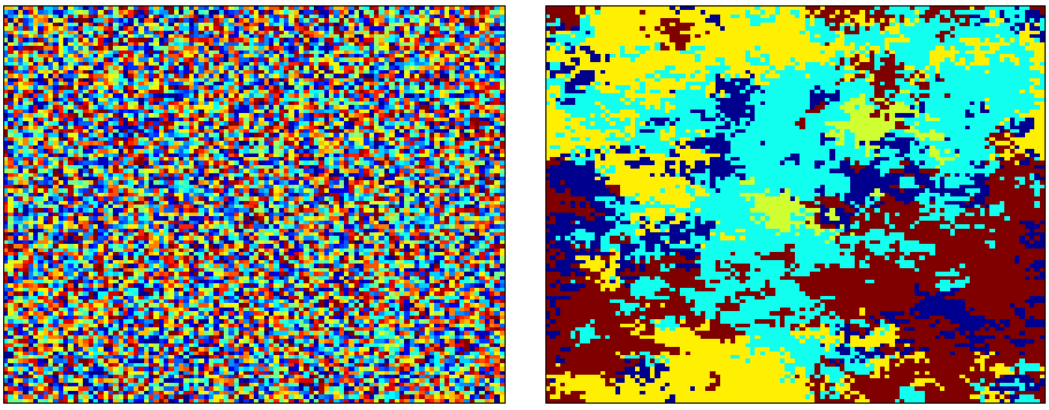

Figure 23: Two configurations of an MVM evolution.
From a random input (left) it self-organizes in coloured areas (right).

5.3.2 Features of the Multitype Voter Model Histogram Sequences

The histogram sequences of the MVM have interesting features that make them suitable for sound synthesis. Figure 24[2] shows the histogram sequence of an automaton of size 70x70, with 20 colours, and Update Probability = 0.5 over 4000 iterations.

[2] The x-axis of the histogram sequences are expressed as indexes of the palette of colours. The relationship between cell values and colour indexes is as follows: CellValue = ColourIndex -1.

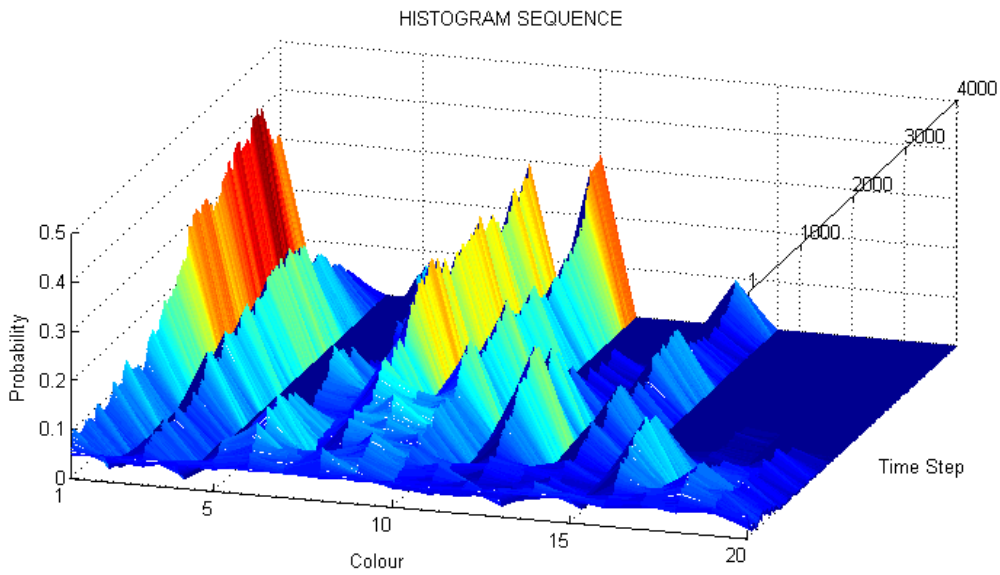

Figure 24: A histogram sequence of an MVM evolution.

The first important characteristic is that the bins of the histogram sequence may represent the time-varying amplitudes of sound partials[3]. This resemblance will be used as a starting point and to establish a narrative thread, although, however, other types of spectral components can be identified, as we will see later. In general, although this text will refer to additive synthesis of sinusoidal components, other synthesis techniques can be controlled with these structures. For instance, an extensive practice has been pursued with subtractive synthesis.

Returning to the analysis of the MVM histogram sequences, it should be noted that the MVM allows us to work with as many colours (i.e., partials) as we wish. In other respects, the disappearance of colours during the run of the automaton is an attractive feature because it parallels the behaviour of sounds produced by acoustic instruments; they usually produce more partials in the attack than in

[3] Assuming this premise, for the rest of the thesis it may be used the term “partials” for referring to MVM histogram bins.

the rest of the sound. Finally, note that since the automaton has a finite size and all the cells are occupied, when the total area covered by one colour increases then it means that the areas of other colours have decreased. In the histogram sequence this means that when some partials grow, other partials decrease. This is appealing for the synthesis of sound textures with interesting internal opposite movements (Figure 25). In this respect, a large interpolation of the histogram sequences prior to rendering the sounds may be desirable in order to synthesise the amplitude movements slowly enough so as to be perceptually noticeable.

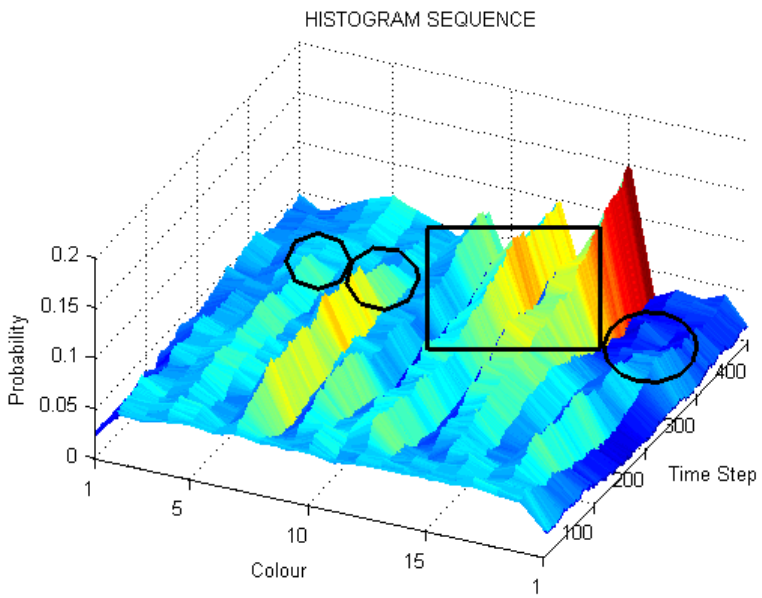

Figure 25: Opposite movements.
The circles show decreasing partials and the rectangle shows increasing partials at the same moment.

In summary, the MVM initially shows great promise for producing dynamic sounds. However, it is necessary to investigate control mechanisms in order to exploit this potential and to solve a number of problems that will be seen to arise in this investigation.

5.3.3 Control

5.3.3.1 Control with the Cell Seeding

The first problem to deal with concerns the lack of attack and release patterns in the histogram sequences. Seeding the automaton with a uniform random distribution of values means that the resulting histogram sequences do not start at zero. Instead, because the sum of all the histogram bins is equal to one at every time step, it can be seen that all the bins of the first histogram in the sequence will hold values close to 1 divided by the number of colours (Figure 26). Although these initial values are close to zero, most of the complex sounds synthesized by adding a considerable number of these partials, will have a high initial amplitude i.e., they will not have an attack pattern. This type of structures does not provide a release pattern either. Generally speaking the MVM will never generate release patterns for all the partials. At least one partial will always remain: the one that covers the whole automaton when consensus is reached.

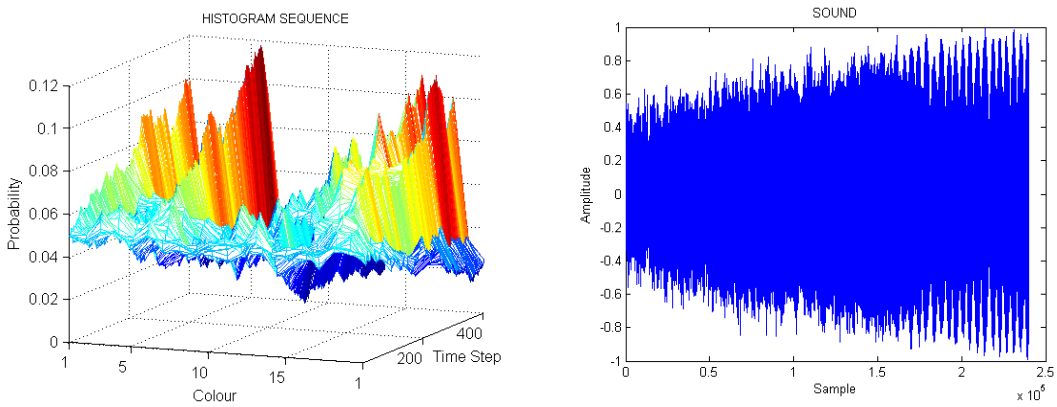

Figure 26: Lack of attack and release patterns.
(Left) A typical MVM histogram sequence.
(Right) A typical sound rendered from an MVM histogram sequence.

Driven by the motivation to resolve this problem a number of control mechanisms for sound design have been developed.

Attack Solution 1: As a first solution to the attack problem, the automaton will be seeded with just one cell per colour placed at random locations. The rest of the cells will be defined as “empty”. With this, it is propitiated a beginning of the histogram sequences closer to zero. The MVM rule is slightly modified to ensure a growth of the occupied cells, by not allowing occupied cells to become empty. Thus, the occupied cells will grow (it could be a growth model) to cover the whole automaton (Figure 27). At that point the attack phase is complete and the automaton will continue with its original rule because the absence of “empty” cells makes the imposition introduced in the new rule no longer effective.

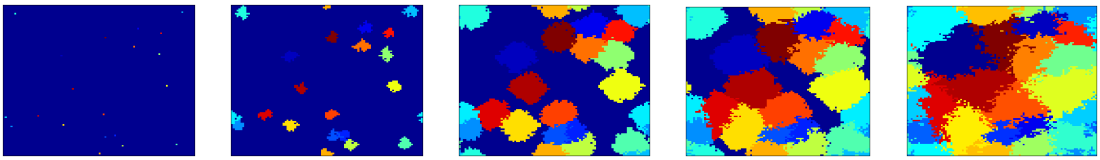

Figure 27: CA evolution for creating attack envelopes by Solution 1.

Figure 28 illustrates the type of result obtained from this solution. The partials start with an amplitude value equal to 1/*n*, which can be virtually zero in a large automaton. In addition, the partials present a growing attack pattern that corresponds to the growth of the seeds –note that in order to make this growth possible the automaton has to be relatively large with respect to the number of colours; otherwise, in extreme opposite cases the automaton could be largely

filled just by the initial seeding. The resulting synthesized sounds have a sigmoidal-like attack.

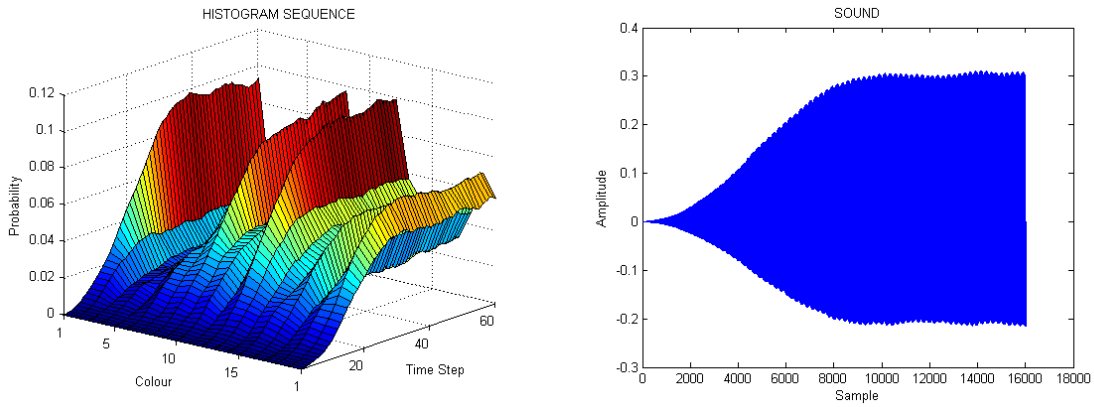

Figure 28: Attack patterns with Solution 1.

One characteristic of this solution is that each partial reaches a different amplitude value at the end of the attack. This is because the seeds are introduced at random locations, thus they start to compete for the territory at different times. This competition for the territory occurs when different coloured areas collide and leads to a decrease in their growth rate. Below we will see how to keep the relative amplitudes under control.

Attack Solution 2: An alternative solution has been developed which covers the automaton with growing areas but not with specific colours but with random colours. As a result, the automaton is filled with a random distribution of colours (Figure 29).

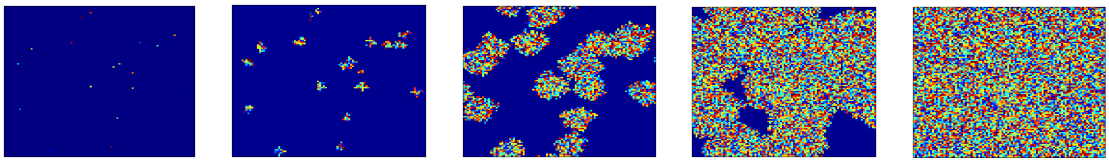

Figure 29: CA evolution for creating attack envelopes by Solution 2.

Figure 30 shows the corresponding histogram sequence which has, once again, sigmoidal-like attack patterns. Also, notice the similar amplitude values that the partials reach at the end of the attack.

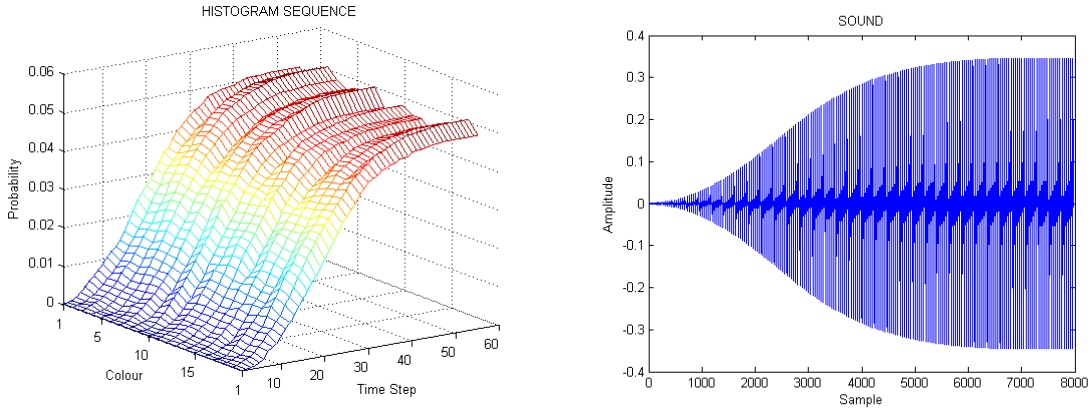

Figure 30: Attack patterns with Solution 2.

In addition, with this solution the automaton can be seeded with as few as just one cell. Therefore, this solution is stronger than Solution 1 in the abovementioned cases in which the automaton size had to be, for some reason, relatively small with respect to the number of colours.

However, it is Solution 1 which has to offer further control possibilities for sound design. For instance, with a different number of seeds for each colour we can control the relative amplitude of the partials, though to a certain extent due to the random processes involved. In Figure 31 both two variables, the number of seeds for each colour and the colour value, have an inverse relationship. This level of prediction is especially useful for working in Automatic and Real-Time modes because we can pre-design the relative amplitudes of an instrument, i.e. a type of sound. Also, a general property of all these attack solutions is the fact that the partials “know” how much they can increase; they will increase until the sum of their amplitude envelopes is equal to one, independently of the number

of partials. Hence they will not exceed ever a maximum amplitude value. This is important for additive type of synthesis, especially in real time, in order to prevent digital clipping.

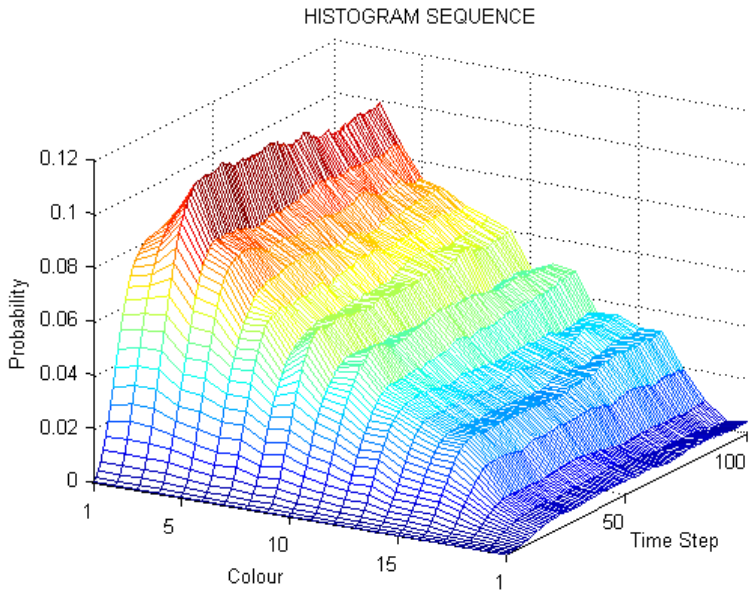

Figure 31: Controlling relative amplitudes.

In other respects, regarding other types of resemblances with spectral components, note that histogram sequences like Figure 31, in which the histograms are considerably smooth, can be viewed as a sequence of spectral envelopes. Finally, Figure 31 also serves as an example to illustrate the role of the weight with which every cell contributes to the histogram in the degree of abruptness of the fluctuations of a histogram sequence. As we saw, this weight inversely depends on the automaton size. In the above example the size of the automaton was made relatively large, 200x200, to keep the amplitudes stable. On the other hand, for Figure 32 the size is reduced to 50x50, whilst preserving the other parameter values, and it can be seen that the amplitude fluctuations are more abrupt.

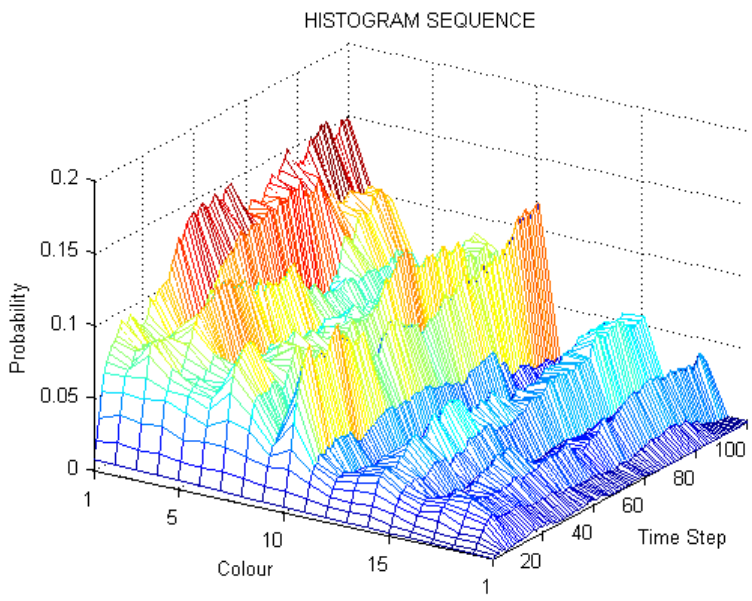

Figure 32: More abrupt fluctuations with smaller CA.

Continuing with further control possibilities, with Solution 1 we can also design entry delays for the partials by seeding the different colours at different CA generations. For Figure 33 the successive colours appeared with a delay of 5 CA generations. See Video Example 3 in Appendix B and Sound Example 2 in Appendix A.

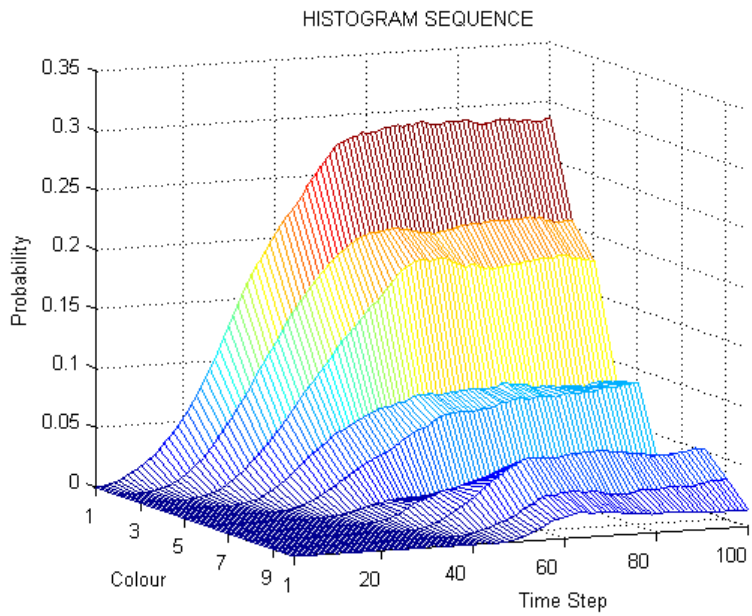

Figure 33: Controlling spectral complexity in the attack.

As well as its applications in Automatic and Real-Time Modes due to the level of prediction achieved, the above control mechanism is remarkable because with it we begin to exploit a powerful capability of HMS with the MVM, the control over time of the spectral complexity of sounds. This type of attack allows us to model a key feature in brass-like instrument attacks in which the spectral complexity increases, in terms of high-frequency components, as loudness increases (Risset 1985 p.12).

In order to create releases a method based on the opposite idea has been devised. Sources of “epidemics” are seeded at random locations, which will expand “killing” all the cells. The curves obtained look sigmoidal and produce different release times for each partial (Figure 34). This is interesting because the sounds of acoustic instruments normally have different release times for each partial. See Sound Example 3 in Appendix A.

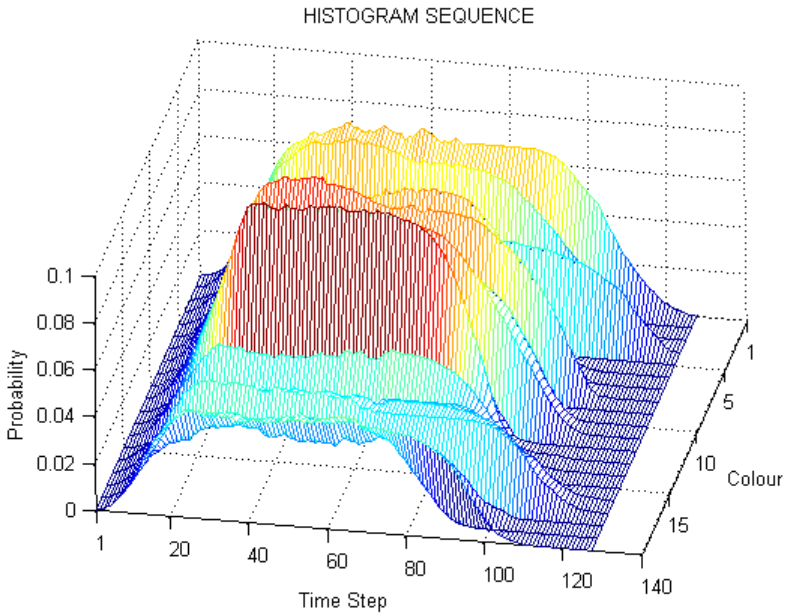

Figure 34: Attack and Release patterns.

The relative release times obtained with this solution are unpredictable, meaning that, for instance, the strongest partial is not necessarily the last one to fade out. This is due to the random locations of the “epidemics” and the unpredictable configuration of the automaton at the beginning of the release. Rather than seeding the sources of “epidemics” in random CA grid locations, they can be placed in the same cells where the Attack Solution 1 placed the initial seeds; then, for relatively short CA evolutions in which the coloured areas have not spread extensively, we are more likely to get release times more proportional to the amplitude values (Figure 35). See Sound Example 4 in Appendix A.

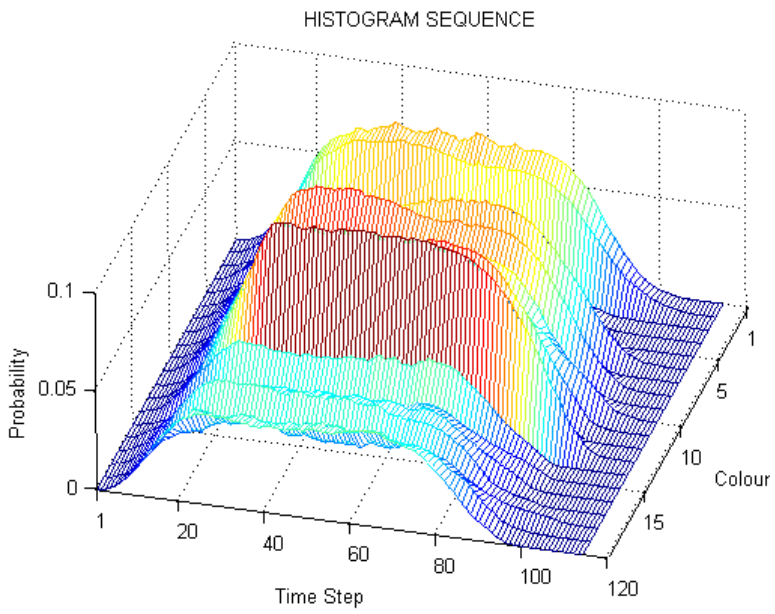

Figure 35: Release times more proportional to the amplitudes values.

In terms of implementation, the “empty” cells of the attacks and the “dead” cells of the releases may correspond to negative cell values; values that are not considered in the histogram computation.

In conclusion, different CA seeding processes along with certain control mechanisms slightly modify the MVM so as to produce, in different ways, attack and release patterns for each partial independently of the number of partials. Nevertheless, other types of attacks and releases different than sigmoidals can be designed through applying external envelopes to each partial or to the final sound produced by the original MVM completely seeded with a uniform random distribution of colours. This approach will be utilised for the rest of this research.

5.3.3.2 Control with the Relationship between the Size and the Number of Colours

As we have seen, the disappearance of colours during the MVM evolution parallels the behaviour of sounds produced by acoustic instruments in regards to the fact that in acoustic instrument sounds, the spectral complexity in the attack is normally higher than in the rest of the sound. In Figure 24 the automaton evolves from displaying 20 colours to 5 in around 4000 time steps. We can favour this dynamics of extinction by working with a smaller automaton. According to the invasion interpretation of the MVM, the twenty species will be competing for less territory, a fact that would provoke a faster extinction of many species. From another point of view, in such a smaller automaton there will initially be less exemplars per species (i.e., cells per colour), a fact that propitiates their extinction.

The problem with considerably small sizes with respect to the number of colours is that the automaton may achieve consensus too soon after a rapid disappearance of most partials. That is, there can be too much uncontrolled

extinction. This often leads to an inconsistent structure with too few partials for the sustain part of the sound (Figure 36).

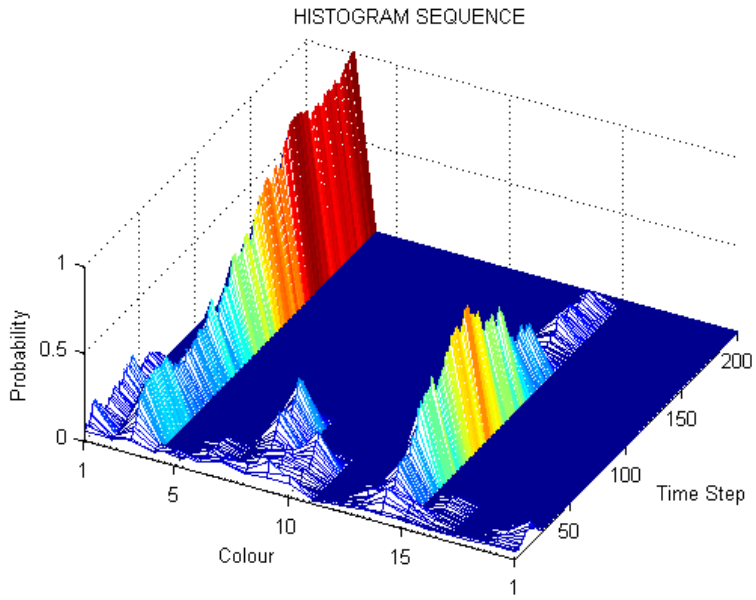

Figure 36: Extinction of colours propitiated but not under control.
Histogram sequence from an automaton with the same settings as for Figure 24 except for the size which has been reduced to 10x10.

Therefore, having propitiated the extinction of colours, the next step is to enable the coexistence of partials for the sustain of the sounds. A control mechanism has been devised based on the same principle of the relationship between the grid size and the number of colours. During the evolution of the automaton every histogram will be monitored. Then, when there remain a determined number of partials (determined in advance for the sustain), the automaton will be replicated a number of times. These replicas will be joined together in order to construct a larger automaton. Figure 37 illustrates this Replication/Enlargement process.

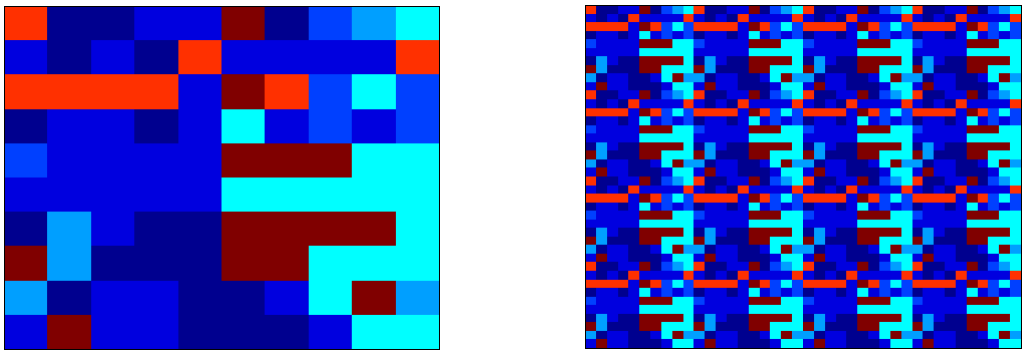

Figure 37: Replication/Enlargement process.
From size 10x10 (Left) to 40x40 (Right).

The normalised histograms of these two CA configurations are identical, and therefore there will not be any discontinuity in the histogram sequence. Having provided more space and more exemplars per species through the Replication/Enlargement process, the remaining partials will be able to coexist in the sustain for longer. Figure 38 illustrates the result of this control mechanism showing four histogram sequences, corresponding to four CA runs. For all these runs, the automaton has initially the same settings as for Figure 36, and the Replication/Enlargement process (from 10x10 to 40x40) is automatically activated for the sustain when five partials remain.

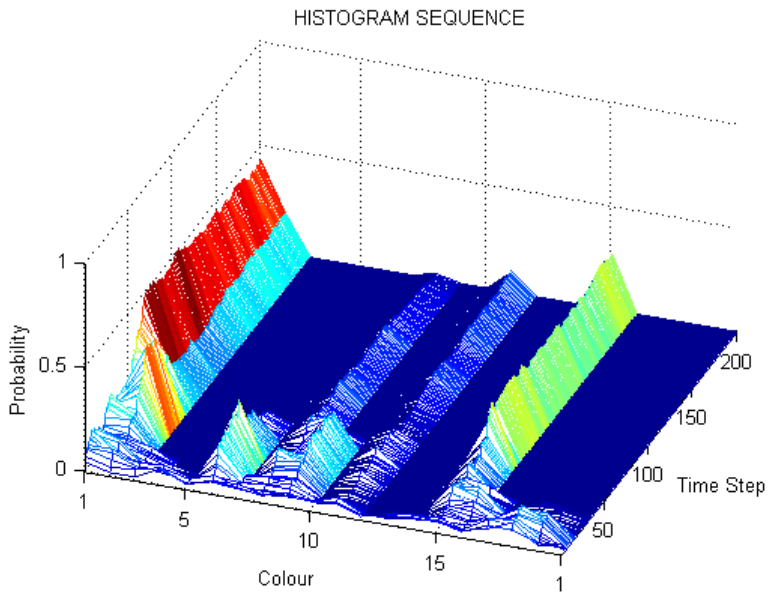

Figure 38: Controlling extinction and coexistence.
(First run)

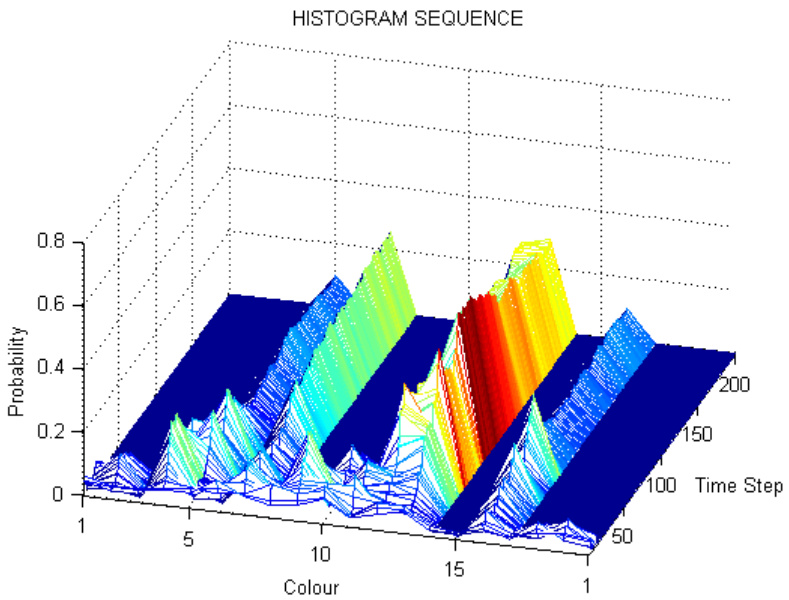

Figure 38: Controlling extinction and coexistence.
(Second run)

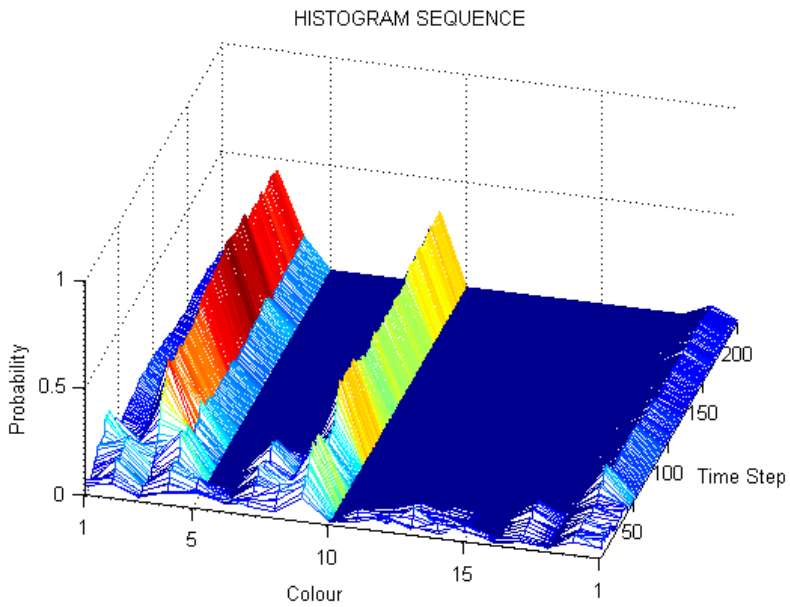

Figure 38: Controlling extinction and coexistence.
(Third run)

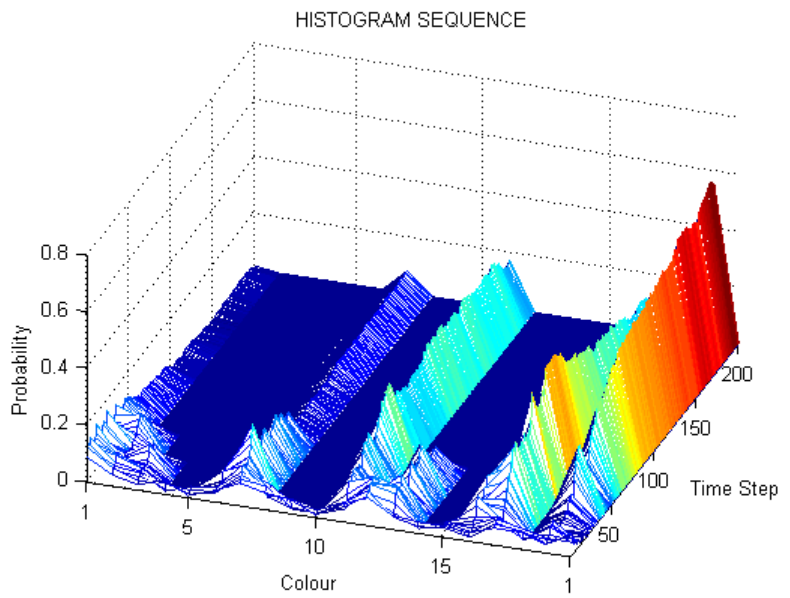

Figure 38: Controlling extinction and coexistence.
(Fourth run)

An interesting characteristic of these histogram sequences is the notable difference between the attack and sustain portions in terms of amplitude fluctuations. This is, again, due to the weight with which every pixel contributes to the histogram. At the beginning of the run, the automaton is small and this leads to interesting abrupt changes in the amplitudes that enhance the transient

character of the attack. Subsequently, there is a contrast with the steady character of the sustain due to the enlargement of the automaton.

In other respects, with this control mechanism we cannot predict which partials will remain in the sustain and what their amplitudes will be. Therefore, regarding modes of use, we will make the most of it for sound design in Non-Real-Time Mode. In this mode, instrument design will be possible because the level of prediction achieved through this control mechanism is considerable; the number of partials that are present in the attack portion and in the sustain portion is predictable. Thus, by assigning certain frequencies in Manual Mode to the partials of one of these histogram sequences, the frequency content of a timbre can be designed. Then, we can establish a relationship between those frequencies and the relative energies that the histogram bins will contribute to the sound. After that, in Automatic Mode, this relationship can be used to automatically assign those same frequencies to the appropriate partials (according to their energy contribution) in subsequent CA runs. This allows us to obtain similar but not identical sounds automatically. Note that the first part of the runs, prior to the Replication / Enlargement process, is of interest because the time-varying amplitudes of the partials evolve very differently in each run. This leads to a set of sounds which have interesting spectro-temporal differences between their attack portions. It is assumed that other essentials necessary for timbre definition, such as external amplitude envelopes, can be applied automatically as well. Finally, an indicator of the energies that the histogram bins will contribute to the sound can be determined by accumulating all their histogram values over time. All of the above is exemplified in the melody in Sound Example 5. Video Example 4 shows a practical usage of the

possibility of producing similar but not identical sounds in video games, with this control mechanism. Finally, another sound example using the Replication/Enlargement process can be found in Sound Example 6.

In conclusion, a control mechanism for sound design has been developed based on the relationship between the CA size and the number of colours. Such a relationship has proved to be a means of controlling the MVM dynamics in terms of extinction and coexistence. A DSP monitoring of the MVM informs on when to alter its autonomous evolution; that is, when to activate the Replication/Enlargement process. In the synthesis domain this all translates into control of the spectral complexity of sounds over time. The number of partials in the attack portion and the number of partials in the sustain portion are nearly under total control. It is not full control however, because the coexistence of partials in the sustain, although propitiated, is not totally guaranteed. For instance, if there is a partial that enters the sustain portion with a very low amplitude, probably it will not be able to survive for longer. Thus, it is not guaranteed that the initial relative amplitudes of the partials in the sustain will be preserved in the long run. Later, an alternative control mechanism will be provided which improves on these limitations.

5.3.3.3 Control with the Update Probability

From the point of view of the voter interpretation of the MVM, the Update Probability (UP) controls the number of individuals that reassess their opinion in each generation. The lower the UP, the more individuals will reassess their opinion simultaneously. In the context of HMS, for sound design purposes, this

UP role is only reliable in the attack/release solutions discussed above. The lower the UP value, the faster the seeds will grow and therefore the shorter the attack/release pattern will be. However, when the MVM evolves normally with all the cells occupied, the competition for the territory is random so it could not be said that, for instance, the lower the UP the faster the automaton will reach consensus. Tests in this vein have shown that an automaton with 20 colours, size 10x10 and UP=0.1 in many runs can take longer to reach consensus than the same automaton with UP=0.5. Nevertheless, some control mechanisms have been developed on the basis of the UP parameter.

5.3.3.3.1 Control with Colour-Dependent UPs

Originally, in the MVM the UP is equal for all the cells. An extended version has been developed here based on the assignment of colour-dependent UPs. From the point of view of the invasion interpretation this means having species with different degrees of vulnerability. A colour with a low UP has a high probability of “being eaten” by a neighbour, while a colour with UP=1 represents an invincible species; a cell with that colour never changes its colour. In this adapted model, therefore, the UP values represent the probability of survival of each species.

The model extension provides an extra level of control over colour extinction and coexistence. In the resulting histogram sequences, partials corresponding to colours with low UPs will decrease at the expense of those partials with higher UPs. As a result of this, it will be possible to predict which partials will become extinct more quickly and which will coexist for longer. This level of

prediction will be useful for instrument design because Automatic and Real-Time modes of use will be possible under predictable conditions. This extended model, hence, represents a control mechanism and with different UP distributions we will be able to design different types of sounds. We will firstly study some UP distributions for producing damped sound structures, and then some for producing sustained sound structures.

One way to ensure obtaining damped sound structures is to reserve just one colour which is assigned the maximum relative UP value. With UP=1 it is ensured that such invincible colour will prevail over the others. It has been observed, however, that it is not just a matter of being an invincible species, but a matter of holding the maximum relative UP value. The issue is relative UP values. Then, if we exclude this partial from the histogram sequence and we consider all the remaining partials that have faded out, we obtain a damped structure. In other words, if we sum the amplitude envelopes of all the partials but the one that has been excluded, the result is an always-decreasing curve. This is because at every time step the sum of all the histogram bins is equal to one, and we are excluding an always-increasing bin. Setting UP values inversely proportional to the colour values gives an example of such a UP distribution (Figure 39).

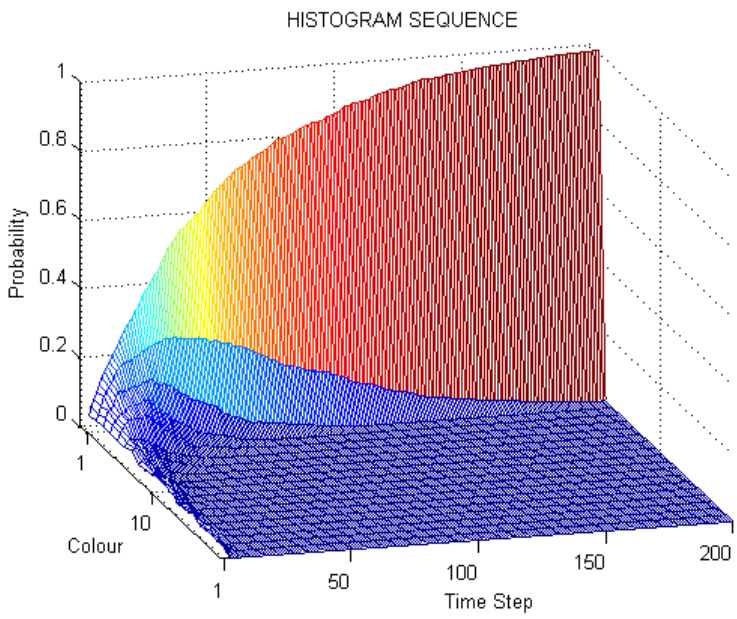

Figure 39: UPs inversely proportional to the colour values.
From a 30x30 automaton with 20 colours.

The next example serves to illustrate the potential for control with the UPs in the design of damped structures. With a different UP distribution the previous histogram sequence can be improved so as to provide a more consistent spectral structure. Three regions are defined: one region corresponding to low partials with UP=0.8, another region in between with UP=0.7 and, the highest region with UP=0.4. The first colour will have UP=1. Figure 40 shows the result of an automaton with 20 colours and size 100x100, with the first partial excluded.

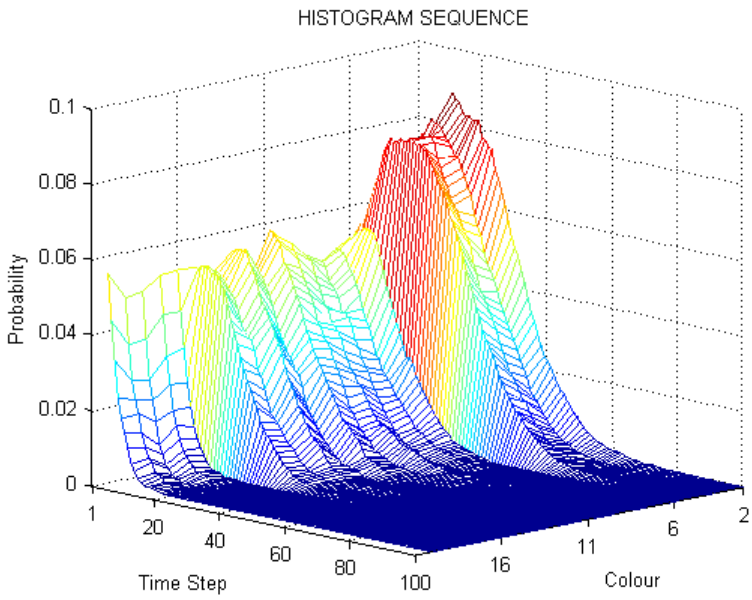

Figure 40: Three UP regions.

It can be seen that the three expected regions are clearly differentiated. Moreover, the level of prediction achieved in subsequent runs is very high. Figure 41 shows the result of a different run with the same settings. It can be seen that these two histogram sequences are very similar.

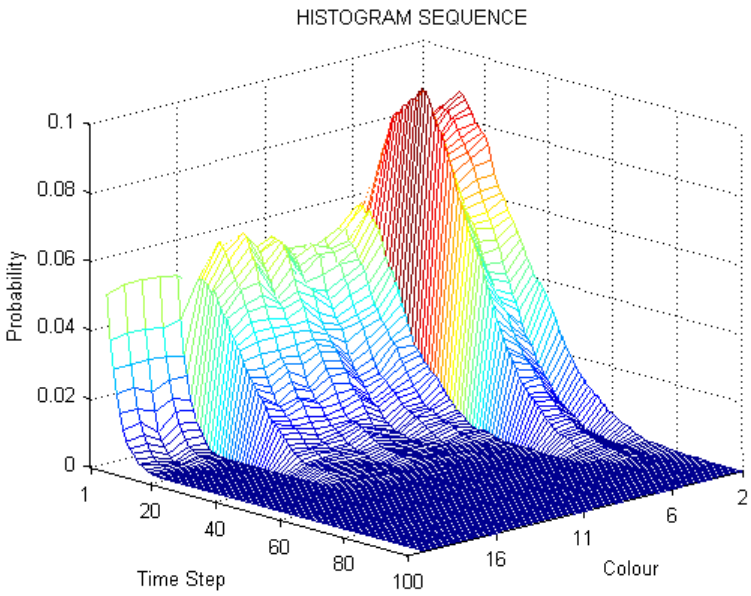

Figure 41: A different run with the same settings as for the previous figure.

Figure 42 shows the result from a smaller automaton (50x50) so as to evaluate the effect of the weight with which every cell contributes to the histogram. The larger automaton produces smoother curves.

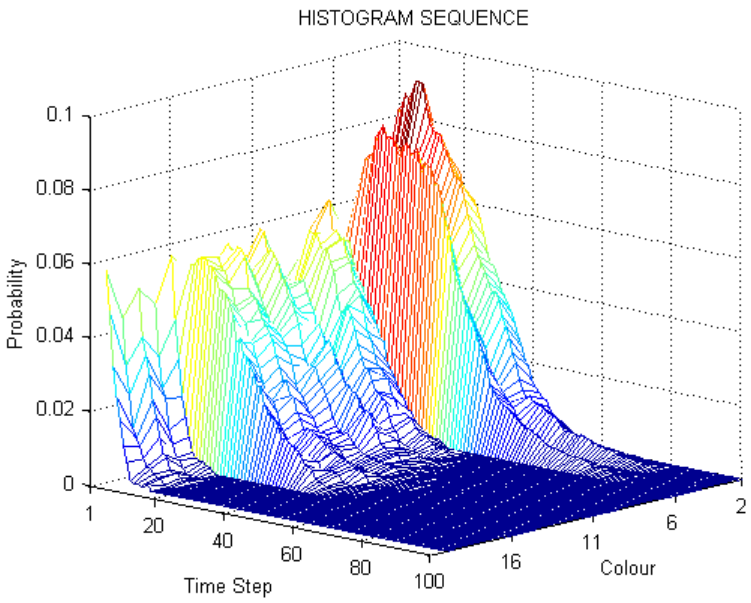

Figure 42: Same settings except for the size.

Another characteristic of larger automatons is that the decay times of the partials within the same region are very similar. In order to differentiate partials, the UP values can be generated with random numbers bounded around the three main regions (Figure 43).

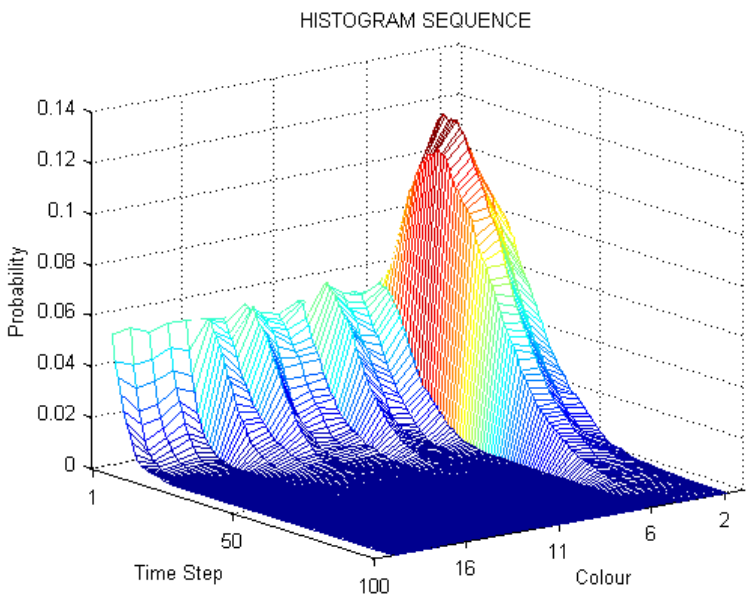

Figure 43: Differentiating partials within the same region.

In this way, different runs of the automaton give predictable structures but with differences in the decay times of the partials and in the maximum amplitudes that the partials reach. These differences, that can be made larger or smaller with the bounded random generator, are perceived as timbral variations. Thus, an instrument has been designed which can produce similar but not identical sounds. See Sound Example 7 and Sound Example 8.

With these structures it has been synthesized convincing imitations of plucked strings and bell tones depending on the frequency ratios between the partials. See Sound Example 9 for bell tone examples. Note that the strongest partials increase in amplitude before fading out. This is at the expense of other partials with lower UPs which decrease faster. Such dynamics are particularly useful in the synthesis of bell tones because they parallel the process observed in real bell sounds in which ‘energy passes between the various modes of vibration’ (Aldoshina et al. 2003) (Hibbert 2000). As a result some partials increase in amplitude while others decrease. This process is illustrated in Figure 44.

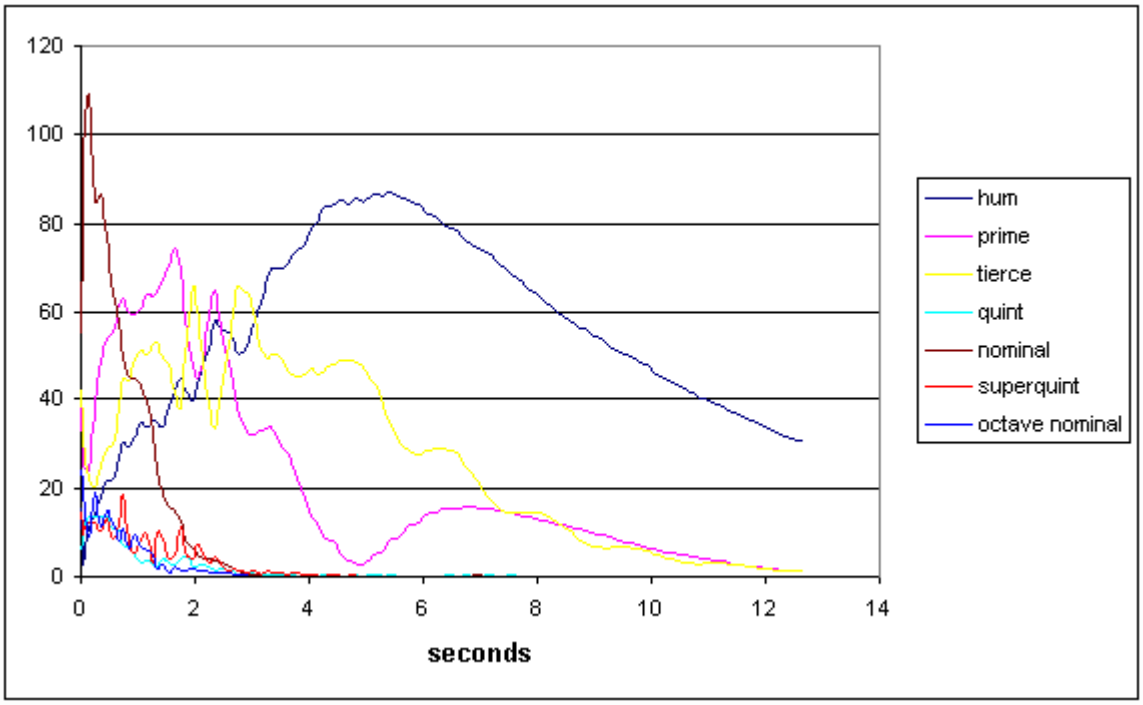

Figure 44: Variation in amplitude of the partials of a bell tone. (Hibbert 2000).

One way to obtain sustained structures is to assign the same maximum relative UP value to the colours (or partials) that we want to coexist in the sustain. The rest of the partials will have a lower UP value and therefore will fade out relatively quickly in the attack portion. In the sustain, the remaining colours will compete for “territory” in equal conditions and therefore, will coexist for longer (according to the relationship between the CA size and the number of colours that we studied in subsection 5.3.3.2).

The following UP distribution is an example of such an arrangement. We will establish three regions. One region corresponding to the partials that we want for the sustain with UP=0.6, another region with UP=0.5 and the last one with UP=0.3. Figure 45 illustrates the result from an automaton of size 50x50 with 20 colours.

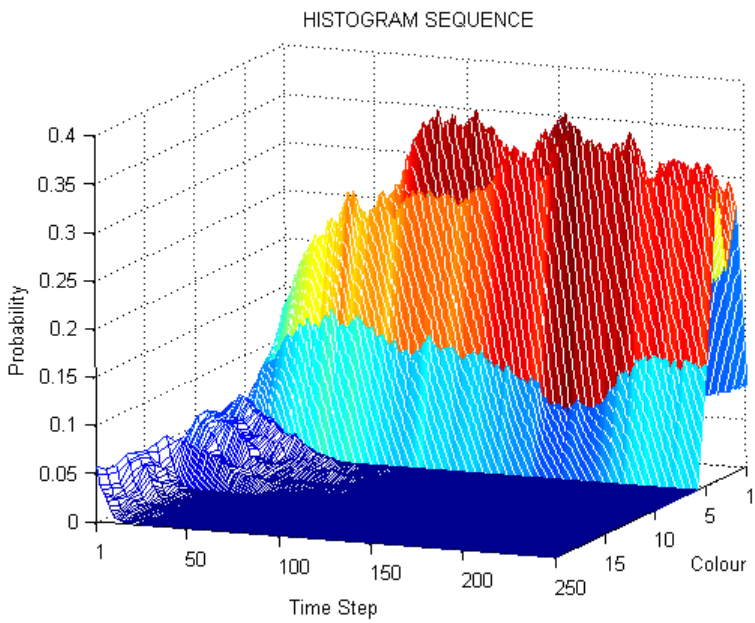

Figure 45: Control with UPs for obtaining sustained sounds.

The controllability achieved is of significance because it allows prediction of which partials will remain in the sustain. This is a step forward over the control mechanism studied in subsection 5.3.3.2. It has significant implications such as, for instance, the possibility of assigning a desired set of frequencies to the partials by the user beforehand, thus allowing to work in Real-Time Mode with a degree of predictability. The resulting sounds have more complexity in the attack portion due to the partials that fade out. These partials are damped and this imprints a specific character to the sounds; they sound like plucked strings or bell tones, but with the peculiarity of being sustained. When listening to these sounds, it is remarkable the fusion between the steady and transient partials, probably because of the energy transfer involved. Applying external envelopes enables new timbres to be designed. By applying an amplitude envelope in the form of a linear fade-in over those sounds which are similar to sustained plucked strings, violin-like sounds have been obtained. See

Sound Example 10.

In conclusion, the UP parameter provides new ways to control the extinction and coexistence of partials and hence, the spectral complexity of sounds over time. An extended version of the MVM has been developed based on the assignment of colour-dependent UPs. Some UP distributions have been examined that produce damped sound structures, along with others that produce sustained sound structures.

5.3.3.3.2 Control with Time-Varying Colour-Dependent UPs

This subsection presents a control mechanism that guarantees the coexistence of partials in the long run. This will be useful for the design of both, note type sounds and sound textures. Regarding the latter, a new textural concept will be presented below based on complex and dynamic sound beats.

The control over species vulnerability studied in the previous subsection can be used to prevent extinction. When a species were in danger of extinction, we would protect it by assigning a high UP. Therefore, in order to ensure coexistence the automaton must be monitored over time and the UP colour dependence must be variable in time. There are different ways to implement this. One possibility is to set the UP of each colour equal to 1 minus the value of its histogram bin, at every time step. With this approach, the new model self-regulates in such a way that the more a partial decreases, the less vulnerable it becomes, while at the same time, the more any other partial grows, the more vulnerable it becomes.

This self-regulation process can be activated when needed, for instance, in the sustain section so as to guarantee the permanence of a number of partials. To illustrate this possibility we will work with the same automaton settings as for Figure 36 and Figure 38. Then, when there remain a determined number of partials (determined in advance for the sustain), the automaton will automatically enter self-regulation mode (Figure 46).

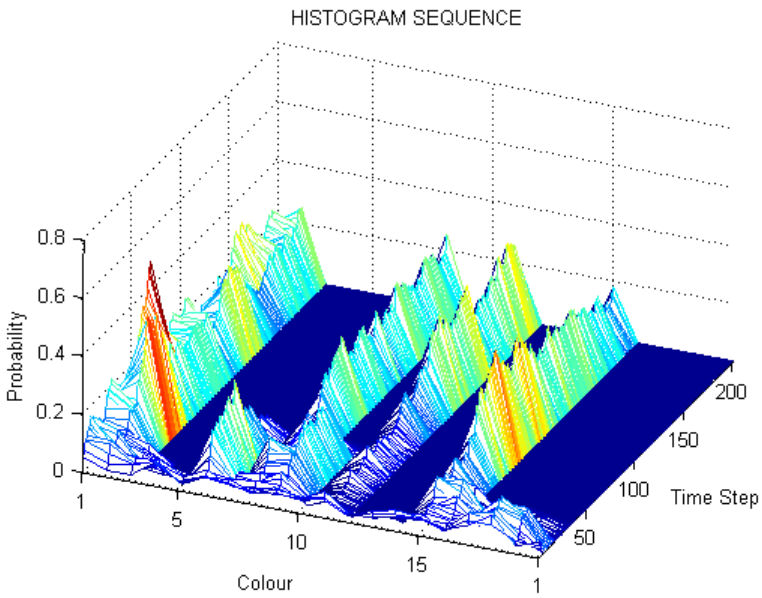

Figure 46: Self-regulation automatically activated when remaining five partials for the sustain.

As a result a transient attack is obtained in which many partials fade out. After this, a sustain is generated in which the permanence of the remaining partials is guaranteed. See Sound Example 11. It is important to note that those partials that enter the sustain portion with low amplitude, would not in normal circumstances be likely to survive for long, but with this self-regulation process they will rapidly “recover”, and increase their amplitude. In general, it can be seen that during self-regulation, the permanent partials undergo a process of equalization in which their amplitudes, although fluctuating, converge to values

close to 1 divided by the number of remaining partials. This is also interesting because the overall relative amplitudes of the partials become predictable. Then, if desired, we can always modify them at will by scaling any partial before rendering the sound. What is not predictable is knowing which partials will remain in the sustain. However, the self-regulation process may be used in combination with other control mechanisms which do offer this kind of prediction. For example, the one studied in the previous subsection on sustained sounds. Similarly, this control mechanism could be used in combination with a Replication/Enlargement process in order to differentiate the attack portion and the sustain portion in terms of abruptness of the amplitude fluctuations, a differentiation not apparent in Figure 46.

This self-regulating model has also been used for the synthesis of dynamic sound textures. In general, we have seen that the MVM produces histogram sequences with interesting opposite amplitude movements between the partials. This dynamic is appealing in the synthesis of polyphonic sound textures because it parallels opposite movements typically found in polyphonic music. This research has focused on opposite movements of just amplitudes, although such movements are commonly found in other musical elements such as pitch (within harmony). Regarding amplitude, opposite movements in polyphonic music occur when voices in the background become important and come to the foreground, increasing their intensity; while at the same time, previous foreground voices become less important and go to the background, decreasing in intensity. By working with relatively large CA sizes, with respect to the number of colours, we can obtain long histogram sequences which provide interesting movements, suitable for this kind of textures with appealing internal

evolutions. Figure 47 is an example from an automaton of size 100x100 with 20 colours over 12000 iterations. In addition, see Sound Example 12.

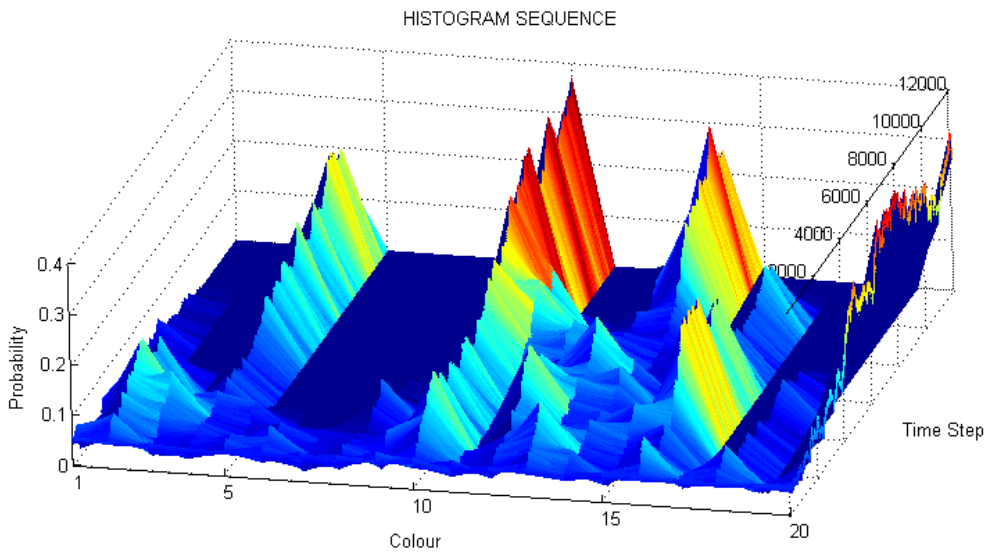

Figure 47: A long histogram sequence of the MVM.

Also, it is worth mentioning that with some processing in Non-Real-Time modes of use (both Manual and Automatic) these histogram sequences can be manipulated so as to exploit their dynamicity for the synthesis of more elaborate musical gestures. Figure 48 is the result of the following sound design process. Firstly, a histogram sequence is sorted according to the energy that each partial will contribute to the sound. Secondly, this histogram sequence is duplicated and flipped in two of its dimensions. Finally, the two modified histogram sequences are linked together. See Sound Example 13.

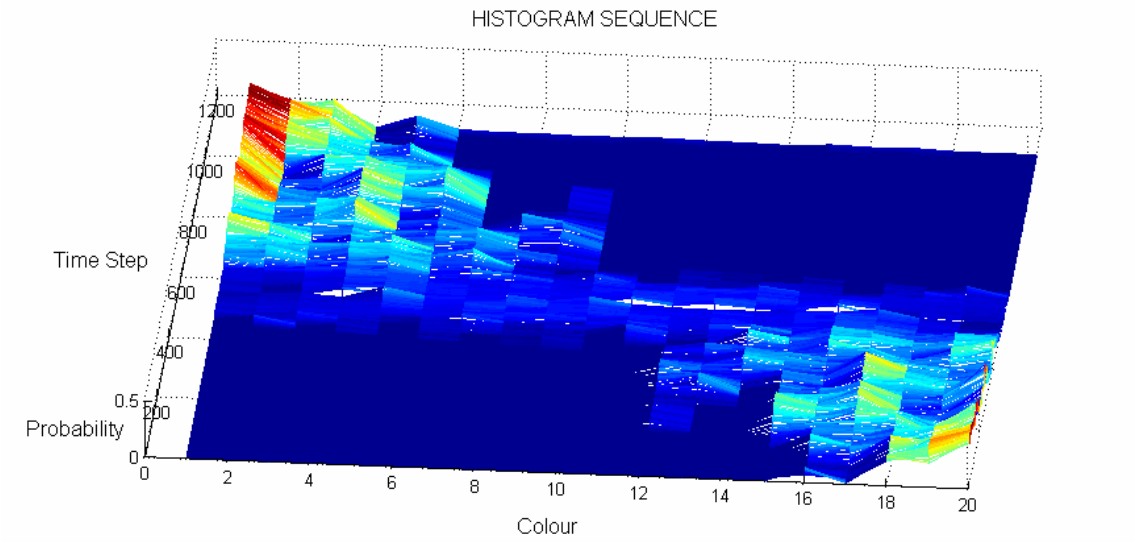

Figure 48: A manipulated histogram sequence.

A key disadvantage of these textures is that they are finite. The MVM cannot provide endless histogram sequences because consensus is always reached. Moreover, close to achieving consensus, when few partials remain, the spectral complexity of the textures may be less interesting. A possible improvement along these lines would be the inclusion of a mutation process. If new species entered the system through genetic mutation, the textures could be endless and more dynamic in terms of their frequency content. Different mutation processes are possible and an in-depth examination is beyond the scope of this thesis. Nevertheless, some experiments have been carried out in which the control with time-varying colour-dependent UPs has been used to protect a single mutant cell, the only one of its species, which otherwise would not be likely to survive for long.

In other respects, a research effort has focused on the exploration of the textural effects of sound beats. Initial experiments along these lines were carried out with long histogram sequences similar to the one in Figure 47. An interesting type of texture was synthesised by extracting a number of the

permanent partials, typically four, and assigning to them slightly different frequencies. The constant changes in the relative amplitudes of the four beating partials made the interference patterns change over time continuously, gradually and non-deterministically. It was obtained dynamic complex beats. The term “complex” refers to the use of more than two beat frequencies. These beats are perceived as a “bumpy” texture, in contrast to the unvarying rhythmic sequence found in common synthetic beats. This process provides a further option to existing techniques for producing similar kind of “bumpy” textures.

There were, however, a series of problems that hindered its practical use. Firstly, the overall amplitude was not constant (Figure 49). This is because although the sum of all the histogram bins is equal to one at each moment, only four out of many other partials were extracted. Also, as we have seen before, due to the disappearance of colours the texture could not be endless. Finally, Real-Time Mode was not potentially feasible because it was not possible to know in advance which would be the partials that would remain.

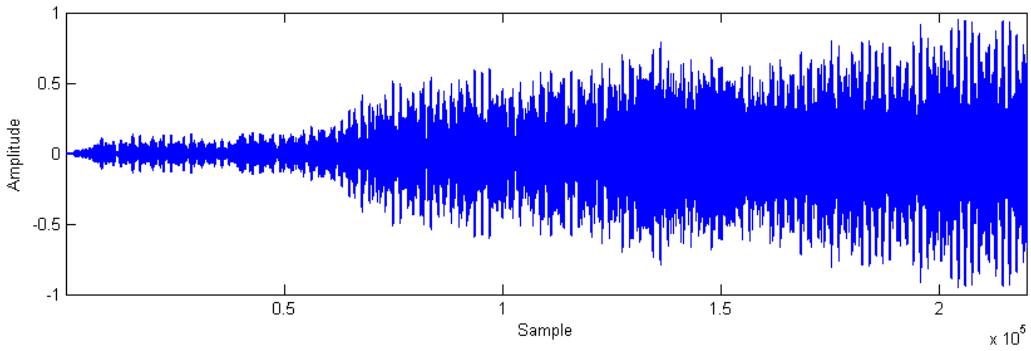

Figure 49: Dynamic complex beats with non-constant overall amplitude.

Working with an automaton with just four colours and the self-regulation algorithm activated from the beginning, so as to ensure that all the partials can coexist, it has been possible to produce this kind of textures with a constant

overall amplitude equal to one (Figure 50). Using this, the user can modify at will the amplitude of the sound; a dynamic beating oscillator has been created.

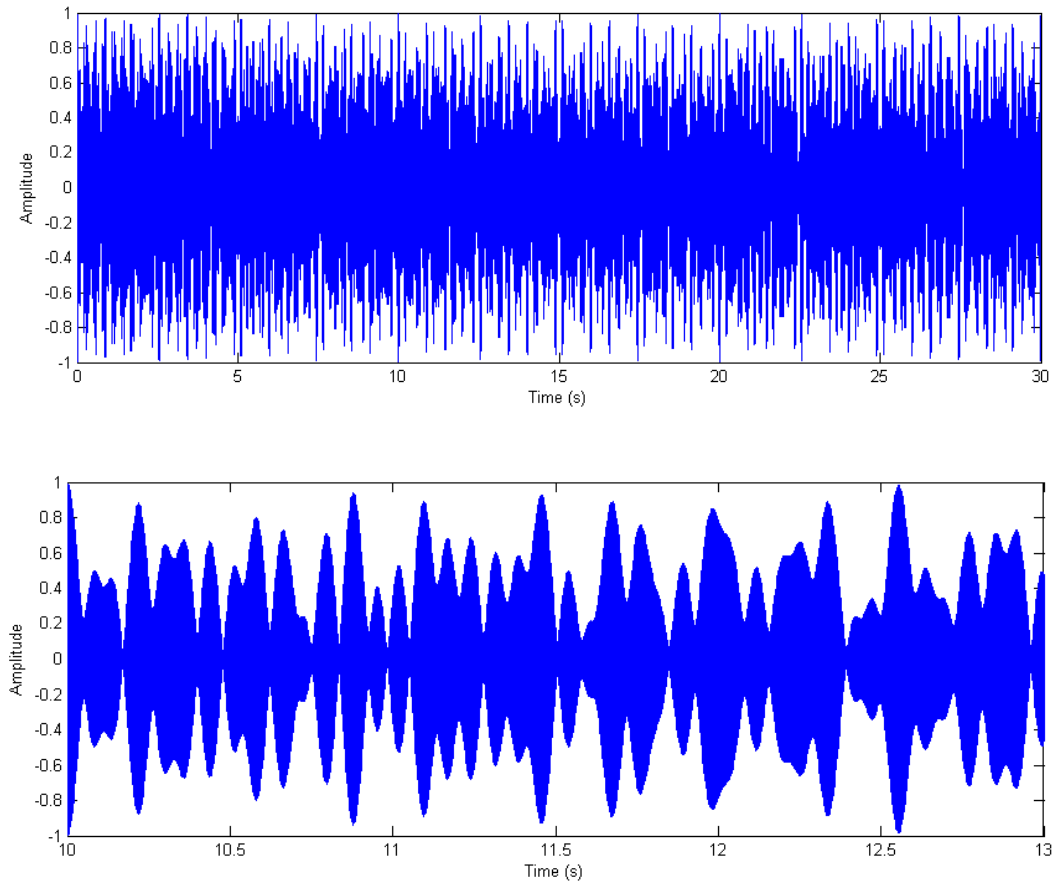

Figure 50: Dynamic complex beats with constant overall amplitude.
(Top) 30 s. sound with beat frequencies 1185.3, 1188.8, 1194.3 and 1199.0 Hz.
(Bottom) Zoom of 3 s. in order to highlight the irregular bumpiness of the waveform.

Endless duration and Real-Time Mode now become potentially possible. In order to produce notable amplitude fluctuations for each beating partial, and thus making the complex beats more dynamic, very small CA grids have been used (Figure 51) –remember that with small sizes each cell contributes a high weight to the histograms which, consequently, vary more from one generation to the next.

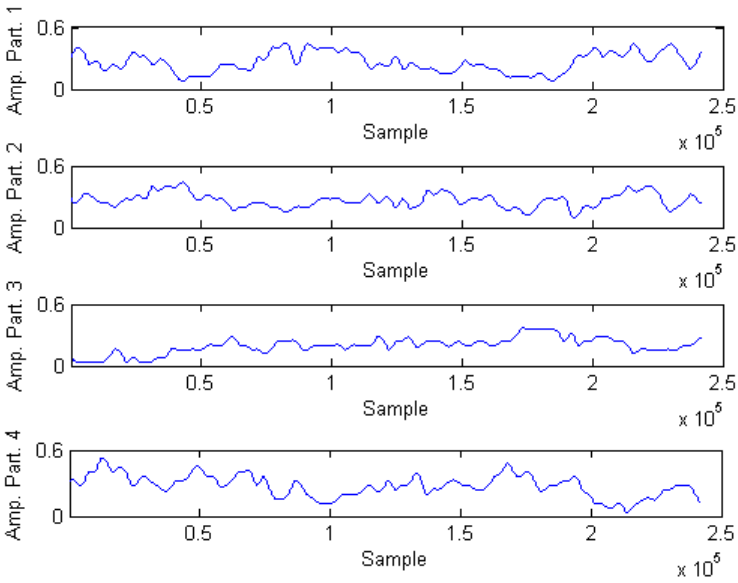

Figure 51: Time-varying amplitudes for four beating partials.
These are obtained from a 5x5 automaton. Note that the sum of the four amplitude envelopes is constant and equal to one (these curves are compensated). They are the result of a large interpolation of 100 time steps so as to produce a 30 s. sound at Fs=8000Hz.

With such small grid sizes the self-regulated model can be pushed to the limit and partials can exceptionally become extinct. In order to solve this problem, a mutation process has been incorporated into the system. At every time step it is checked whether a partial has become extinct and if so, one cell at random would change its colour to that of the extinct species. This mutation process would not normally work in the original MVM where just one species exemplar would be unlikely to survive. But in the self-regulated model such an exemplar will have a very high UP, so it will be virtually invincible. See the self-explanatory Video Example 5, which has sound examples of dynamic complex beats.

Last but not least, a series of successful experiments has been carried out on applying this textural concept to existing sounds, by means of audio processing. Technically speaking it is only needed to extract the amplitude envelope of the

dynamic complex beats. Then, any sound can be multiplied by this envelope. The result obtained is a tremolo-like effect (often used with electric guitars) with the novel characteristic of being complex and dynamic, and hence, more textural. In addition to the common controllers that a tremolo effect could have, including the depth, this new effect could offer control over the degree of complexity and dynamism. See Sound Example 14 and Sound Example 15.

To conclude, a new control mechanism has been developed which aims at ensuring the coexistence of partials by controlling the automaton over time. The basis of this method lies in a DSP monitoring of the automaton, which informs us on when and how to alter its autonomous evolution. The method has been applied to design note type sounds, as well as sound textures.

5.4 HMS with the Hodge Podge Machine

5.4.1 The Hodge Podge Machine

The hodge podge machine (HPM) is a mathematical model of the ‘oxidation of carbon monoxide catalyzed by palladium crystallites which are incorporated into a zeolite matrix’ (Gerhardt et al. 1989 p.209). Its creators ‘were mainly interested in the question whether the complex oscillation patterns of the aforementioned CO oxidation might be correlated to the formation of certain spatially ordered structures on the catalyst, which are caused by local interactions between different "catalytic units"’ (Gerhardt et al. 1989 p.209). These local interactions were determinant in the decision to use a CA modelling

approach. Once the 2D CA model was devised and its behaviours were studied, 'it turned out that this automaton not only describes the typical behaviour of the CO oxidation, but that it leads to a self-sustained organization of fascinating spatial patterns, such as circular and spiral waves. These are very similar […] to those observed in excitable media, e.g. the Belousov-Zhabotinskii reaction' (Gerhardt et al. 1989 p.210).

Let us now describe the model. In the HPM the states of a cell can be interpreted metaphorically as follows: the state characterized by a minimum value 0 is called "healthy". The state given by a maximum value *V*-1 is called "ill". All other states in between are called "infected". The model comprises three rules. Rule 1 is applied to "healthy" cells, Rule 2 to "infected" cells and Rule 3 to "ill" cells. Note that to broaden musical possibilities, this thesis includes experimentation with modifications of the original rules, and the use of different types of neighbourhoods. The original rules can be found in (Gerhardt et al. 1989 p.213). The following modified rules have been implemented leading to interesting results in a number of experiments:

$$m_{x,y}[t+1] = \begin{cases} round(A/r_1) + round(B/r_2) & for \;\; m_{x,y}[t] = 0 & Rule\,1 \\ \min\{round(S/A) + K, \quad V-1\} & for \;\; 0 < m_{x,y}[t] < V-1 & Rule\,2 \\ 0 & for \;\; m_{x,y}[t] = V-1 & Rule\,3 \end{cases} \tag{5}$$

where the state of a cell at a time step *t* is denoted by $m_{x,y}[t]$; *x* and *y* are the horizontal and vertical coordinates of the location of the cell in the automaton; *A* and *B* represent respectively the number of "infected" and "ill" cells in the neighbourhood; *S* stands for the sum of the states of all cells in the

neighbourhood; K, r_1 and r_2 are constants $\in \mathrm{N}$; *round* stands for the nearest integer; and V is the number of possible states that a cell can adopt.

In these modified rules, the rounding processes to the nearest integer differ from the original model, in which rounding-down processes are used (i.e. rounding towards minus infinity). With this small difference the automaton exhibits the same type of behaviours studied below. However, this modified version is preferred in this thesis because its Spiral behaviour is better characterised by the histogram analysis. Later, when this Spiral behaviour is introduced, some illustrative examples will be given.

The different neighbourhoods that have been considered will be specified below when the behaviours are presented. The automaton has been implemented with periodic boundary conditions, i.e., with a torus shape.

A desirable property of this CA model is its cyclic nature, which allows us to work with different ranges of cell values. The cyclic nature of this automaton is characterized by two end states (the minimum and the maximum cell values), by two positive feedback mechanisms which make the cell values increase (Rule 1 and Rule 2) and by a negative feedback mechanism which makes the cell values decrease (Rule 3).

Regarding behaviours, from an initial uniform random distribution of cell values, the HPM can exhibit various behaviours depending on the values of its rule parameters (K, r_1, r_2 and V), as well as on other specifications such as the type of neighbourhoods, and so forth. For sound synthesis with HMS, three suitable

behaviours have been studied. One is the Spiral behaviour, for which this automaton is best known. This thesis will refer to the other two as the Quasi-Synchronic behaviour and the Fast-Infection behaviour. In the next subsection we will study each of these behaviours and the histogram sequences they produce.

5.4.2 Features of the Hodge Podge Machine Histogram Sequences

In the following subsections three behaviours of the HPM will be studied with the help of a histogram analysis. Each of these behaviours can be obtained from a variety of different combinations of rule-parameter values and neighbourhood types. Therefore, some guidance will be provided on how to obtain them. The structural elements of the histogram sequences useful for sound synthesis will also be presented.

5.4.2.1 Quasi-Synchronic Behaviour

This behaviour can be obtained using the Moore neighbourhood. Regarding the values of the rule parameters, interesting results have been achieved with large values of V (number of colours), from hundreds and usually thousands. Based on experimental practice, the rest of the parameter values could be set as follows: K equal to a value near 20% of V, and r_1 and r_2 both equal to 2. By expressing the value of K proportionally to V one can explore many different sets of parameter values leading, hopefully, to this behaviour. However, this is an intuitive generalisation based on experimental practice and provided for

guidance only. Other parameter values can also lead to this behaviour. The behaviour can be seen in Video Example 6.

Let us now describe the behaviour with the help of a histogram analysis. The behaviour is labelled “Quasi-Synchronic”, because the values of all the cells increase and reach their maximum state fairly simultaneously. After the cells reach their maximum value and a new CA cycle begins, patterns with distorted circumferences emerge. These patterns create narrow bands, or peaks, in the corresponding histogram. From here, the cell values grow towards the maximum state during certain CA generations and the boundaries of the distorted circumferences become “blurred”. These patterns create wide bands in the histogram (Figure 52).

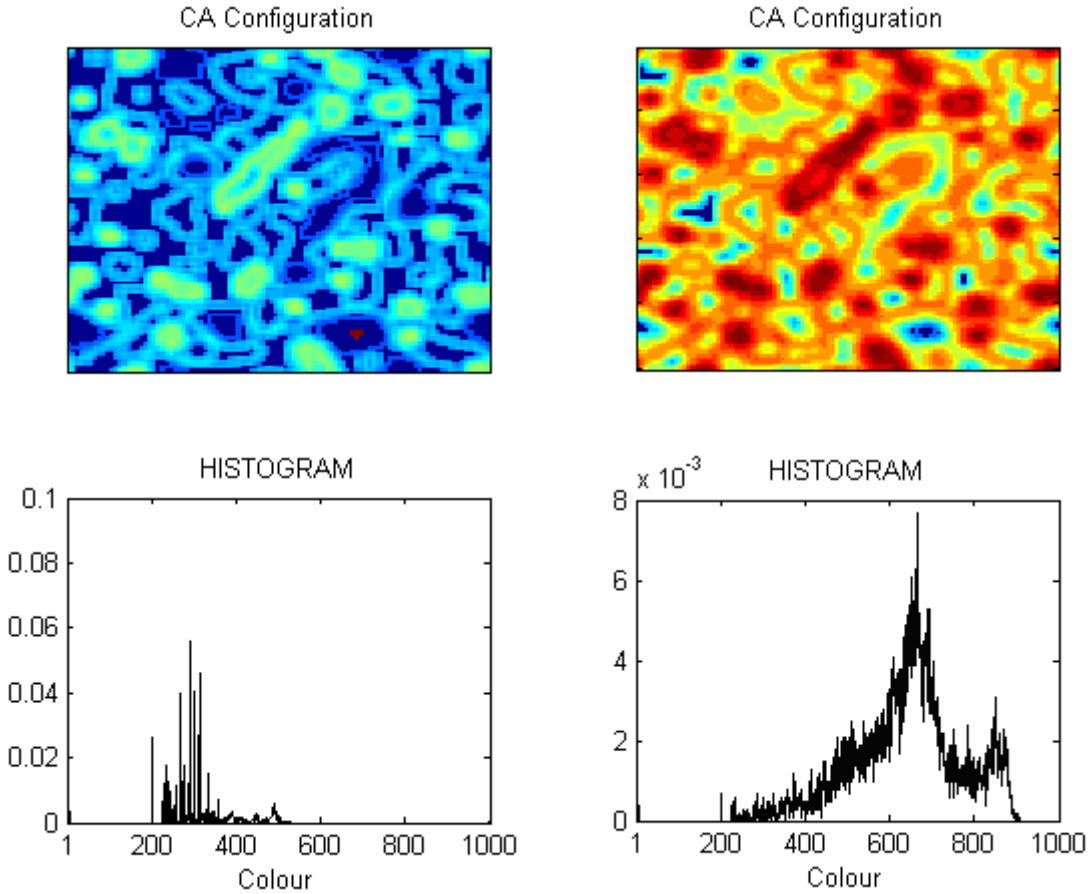

Figure 52: Distorted circumferences (Top) and the histograms they produce (Bottom). (Left) narrow bands. (Right) wide bands.
CA size of 100x100 cells.

At each cycle of the automaton this process is repeated and the histogram analysis reveals that the automaton self-organises through the same set of predominant colours –this set of colours depends on the rule parameter values. In addition, the distorted circumferences will emerge with different shapes in every cycle (Figure 53). These two facts will propitiate the formation of structural elements with time-varying “amplitudes” in the histogram sequences.

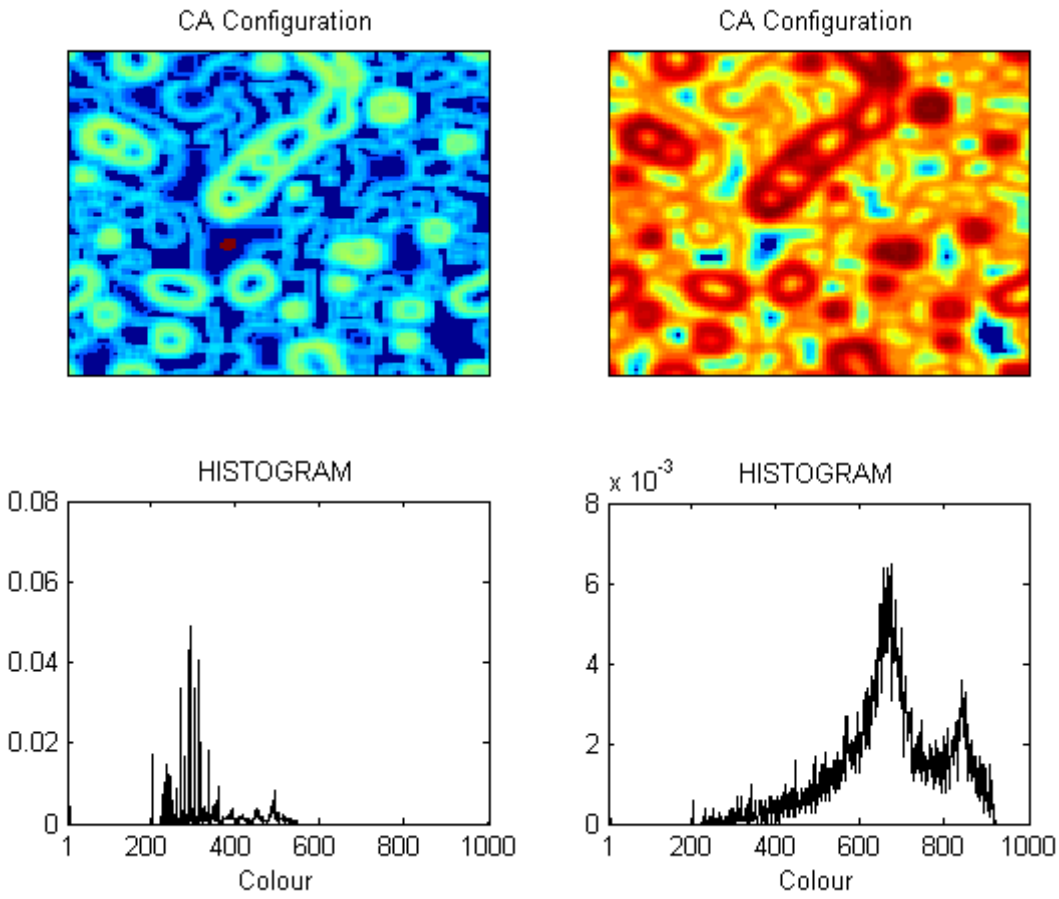

Figure 53: Patterns of another cycle of the automaton.
The distorted circumferences are slightly different.
The histograms register the same predominant colours with slightly different "amplitudes".

Figure 54 shows a typical histogram sequence of this behaviour from a frontal view allowing us to see all histograms superimposed. Three different zones have been established according to the type of structural elements they accommodate. NB1 and NB2 are zones with narrow bands and WB is a zone with wide bands. Notice that the left-most narrow band of NB2 is the highest in that zone, and is considerably separated from the other lower narrow bands. This feature will be used later to characterise histogram sequences of Quasi-Synchronic behaviour.

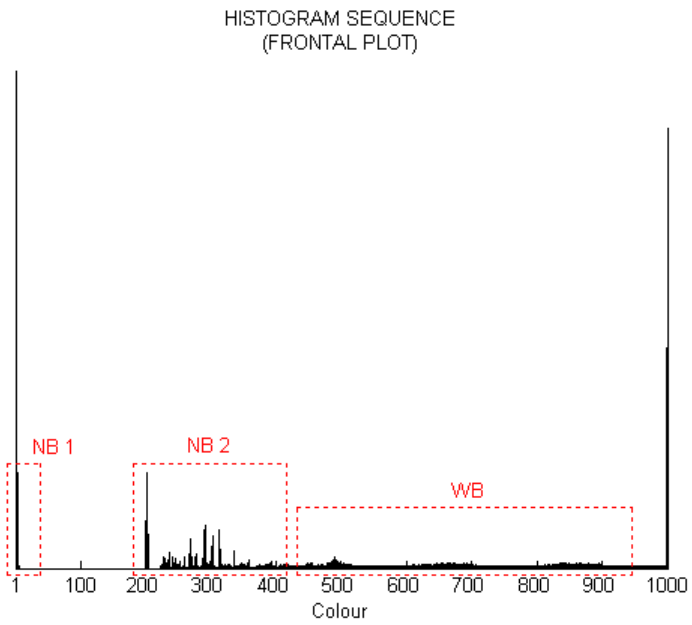

Figure 54: Frontal view of the histogram sequence of a Quasi-Synchronic behaviour. Three zones have been differentiated in red colour.

At this point, by the way of a digression, we can see how a histogram analysis can help us to better understand the functioning of the HPM automaton. In Figure 55 we can see, in blue colour, the effect of the rules and the *K* parameter. Rule 1 creates the NB1 zone. Between this zone and NB2 there is a gap, due to the addition of constant *K* in Rule 2. This means that the cells of the automaton never adopt such colours. Rule 2 creates the zones NB2 and WB as well as the peak in the last bin. Furthermore, Rule 3 creates the peak in the first bin.

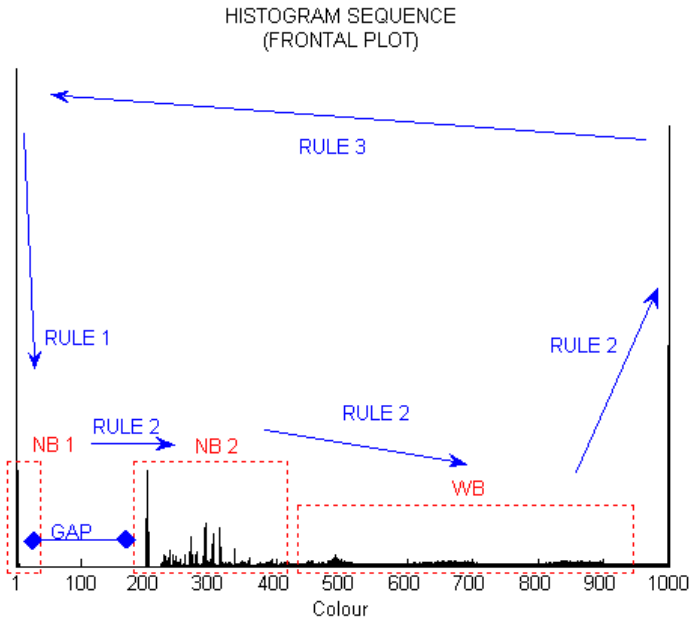

Figure 55: The role of the CA rules (in blue colour).

Let us go back to the histogram sequence structural elements. Due to the quasi-synchronic nature of this behaviour, the bins of the histogram sequences display an oscillatory behaviour in time. It has been a common practice to compute the envelope of each bin in order to obtain smooth time-varying structural elements that can finally resemble spectral components of sounds. For instance, Figure 56 shows the evolution of an NB2 zone bin (red stem plot) together with its envelope (discontinuous blue line).

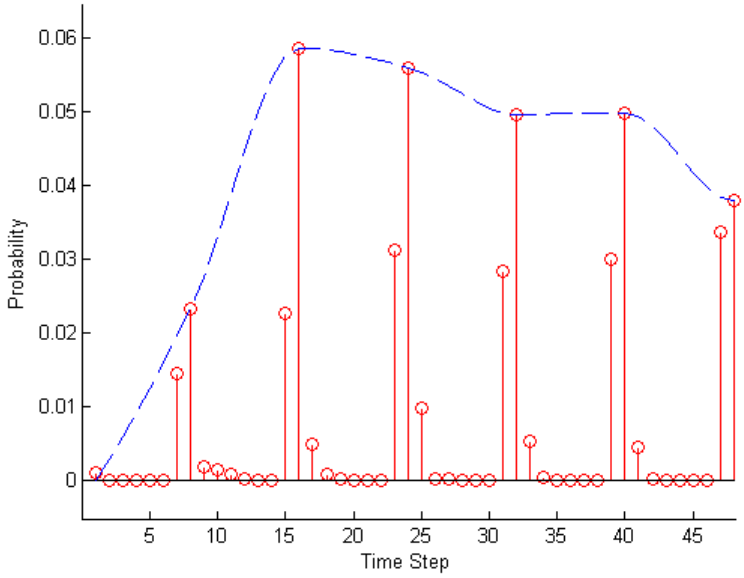

Figure 56: Computing the envelope of a histogram sequence's bin in the NB2 zone.

As has already been discussed, the first time step, which corresponds to the histogram of the initial seed with a uniform random distribution of colours, will be approximately equal to 1 divided by the number of colours. When working with hundreds or thousands of colours, this value can be approximated to zero. Moreover, as the first significant peaks in NB2 will appear after a few generations, or time steps, we can set the first value of the computed envelope exactly to zero.

As a result of this, we obtain an envelope with an attack pattern and a sustain which normally has interesting fluctuations; it finally resembles the time-varying amplitude of a sinusoidal component. Figure 57 shows several bins of a typical NB2 zone having estimated their envelopes. See Sound Example 16. Another interesting characteristic of these "sinusoidal components" is that they are quite correlated in amplitude. In other respects, it should also be noted that once the entire histogram sequence envelope is extracted, the sum of all the bins at any given time step is no longer one. Therefore, once the envelope of a histogram

sequence has been computed, the y-axis of its 3D plot no longer informs on a probability measure, and hence, will not be displayed in the figures.

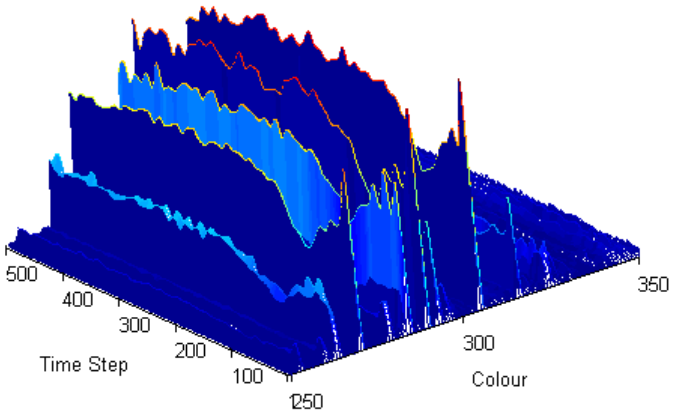

Figure 57: Zoom of an NB2 zone after envelope estimation.

In NB2 zone we can also find other elements resembling sound components. Some runs of the HPM exhibiting Quasi-Synchronic behaviour produce interesting peaks at the beginning of the histogram sequence that could be considered as transients (Figure 58).

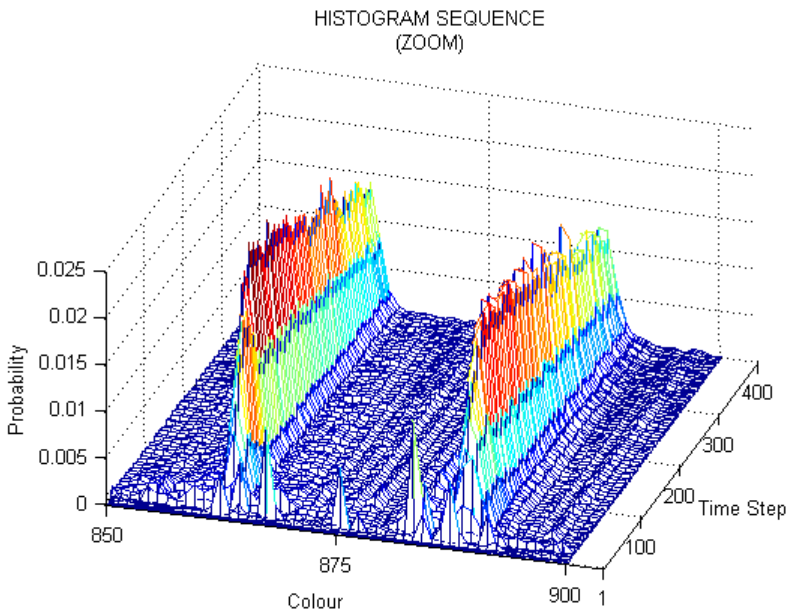

Figure 58: Zoom of an NB2 zone showing a noisy transient structure in the attack.

To conclude with the identification of structural elements resembling spectral components in the Quasi-Synchronic histogram sequences, some remarks will be given about the wide bands of WB zone. Figure 59 shows two of these structural elements.

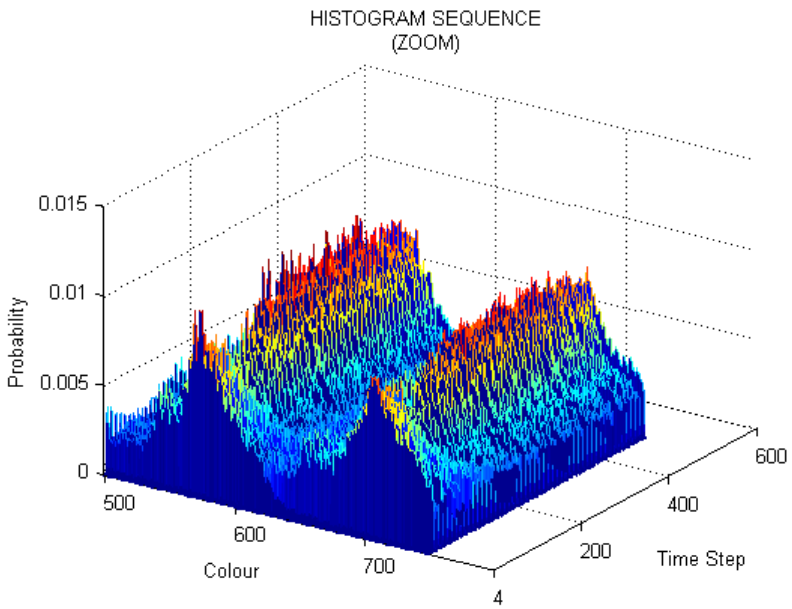

Figure 59: Zoom of two wide bands obtained from a Quasi-Synchronic behaviour.

An interesting characteristic of these bands is their amplitude correlation with the narrow bands in the NB2 zone. Regarding the number of them, note that certain HPM settings and rule modifications can lead to the production of many wide bands, for example, as many as ten of them. These wide bands can be considered as steady-state noise components, as formants, or in general terms, as a sequence of spectral envelopes. A number of experiments have been carried out with subtractive synthesis using these bands as the time-varying magnitude response of a time-varying filter. Different sound sources can be considered for this. For producing steady-state filtered noise, a white noise oscillator can be used. These bands have also been considered as spectral

envelopes for filtering sawtooth, triangle and square wave oscillators. Regarding formant synthesis, more work would be needed to model vocal tract or instrument resonances. It will be necessary to take into consideration the separation of each band and their subsequent location at specific center frequencies with appropriate relative amplitudes and bandwidths (by means of interpolation).

5.4.2.2 Fast-Infection Behaviour

Fast-Infection behaviour can be obtained using neighbourhoods that are smaller than the Moore neighbourhood. Although the Neumann neighbourhood could be an example, most of the experiments have been carried out by using an even smaller neighbourhood, in which the central cell is not viewed as a neighbour (by inheritance from Chaosynth –studied in Chapter 3, Section 3.2). Therefore, only the four orthogonally adjacent cells, namely, north, east, south and west are considered as neighbours. If the central cell is not considered as a neighbour of itself, a rule modification is necessary to prevent a division by zero in Rule 2 when the *A* parameter (number of infected cells in the neighbourhood) is zero. Note that this will never be the case when the central cell is considered as a neighbour, because the condition for applying Rule 2 is that the central cell is infected, and thus the *A* parameter would always be greater than or equal to one. The solution adopted is to check when *A* is equal to zero and to set it equal to 1 when applying Rule 2. This behaviour can be seen in Video Example 7.

The Fast-Infection behaviour is interesting because it generates non-sustained structures in the histogram sequences, as opposed to the (supposedly) endless

sustained structures of the Quasi-Synchronic behaviour. A histogram analysis reveals that the automaton reaches a long-term behaviour in which the cell values are registered only in the first bin, then in the NB1 zone, and finally in the last bin of the histograms. Before reaching this long-term behaviour the rest of the histogram's bins fade out to zero, creating non-sustained structures.

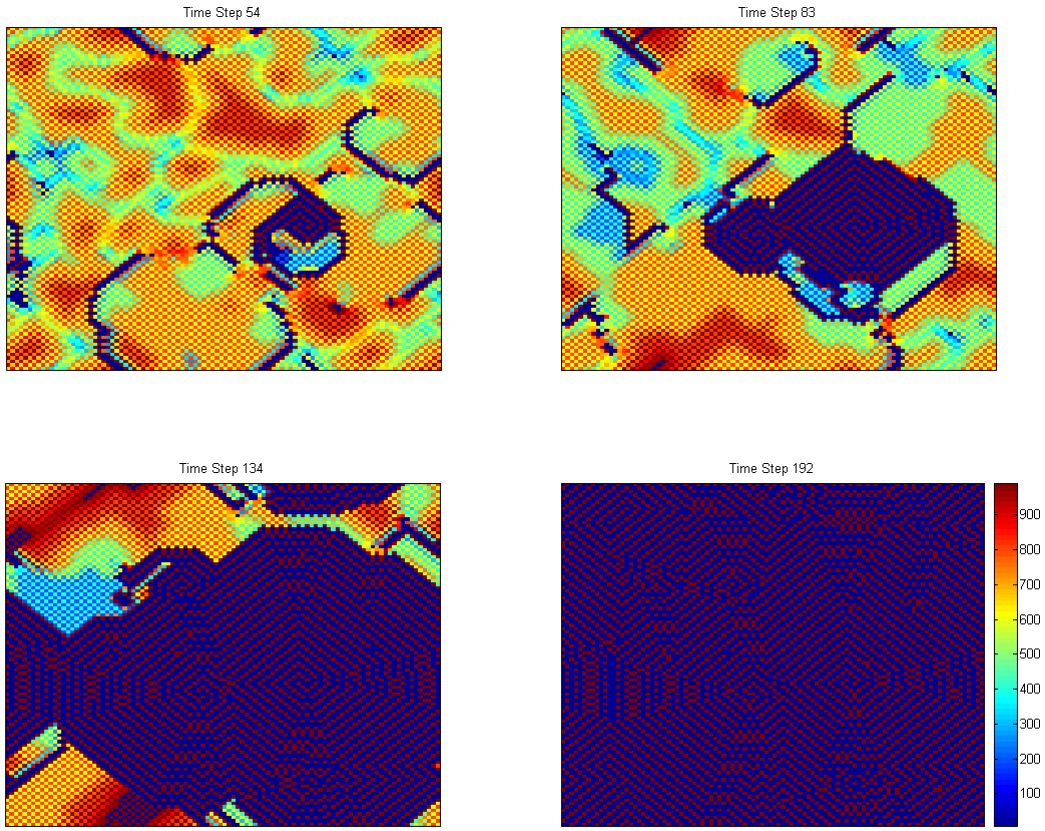

Figure 60: Fast-Infection behaviour.

Figure 60 shows four configurations of an evolution of the Fast-Infection behaviour. Even with a colour bar beside the last configuration, it is hard to have an idea of what values the cells attain. It is the histogram analysis which can give answer to this, as explained above. Also, after such a histogram analysis one can find an explanation for this behaviour with respect to the size of the neighbourhood. By considering fewer neighbours, the divisor of Rule 2 is lower, and thus, those cells that have been infected by Rule 1, which are

registered in the zone NB1 of the histogram sequence, will have great chances of getting immediately ill (this is why the name Fast-Infection was chosen) next time Rule 2 is applied. There are many combinations in the neighbourhoods leading to this effect. For instance, ill cells in the neighbourhood make the Rule 2 numerator very large, whilst not affecting the denominator. Thus, for instance, just by having one ill neighbour and only one infected neighbour, Rule 2 makes this cell automatically ill.

Figure 61 illustrates two types of non-sustained structures that can be obtained, after envelope estimation, from the histogram sequences of this behaviour. The top image shows a number of narrow bands, which can be considered as time-varying amplitudes of sinusoidal components. The bottom image shows a more compact structure, which can be considered as a sequence of spectral envelopes. See Sound Example 17.

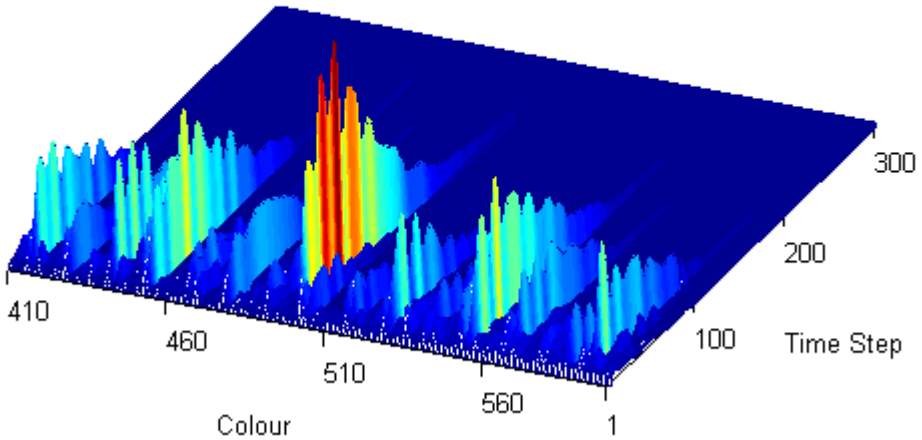

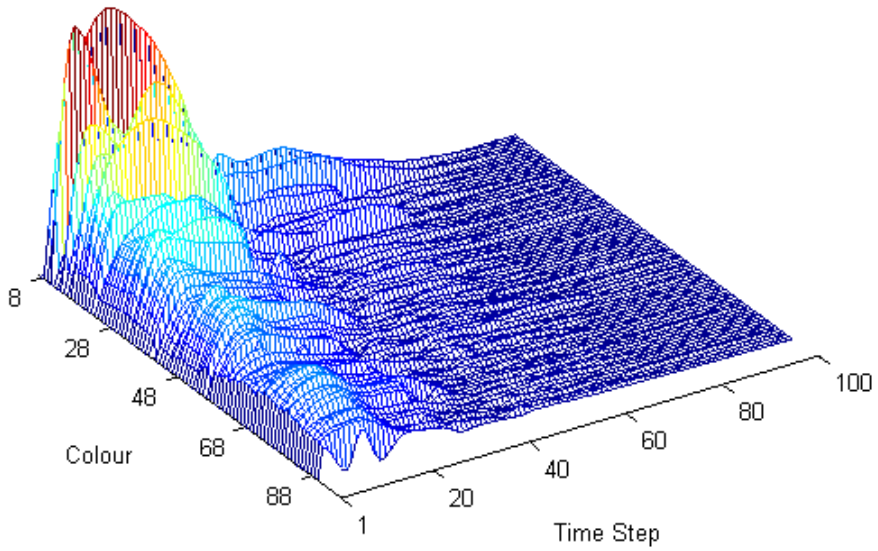

Figure 61: Non-sustained structures obtained from the Fast-Infection behaviour. Two zooms of two histogram sequences after envelope estimation.

5.4.2.3 Spiral Behaviour

The Spiral behaviour is interesting for sound synthesis because it parallels the way that acoustic instruments produce most tonal sounds: they start with a more complex vibration, often noisy, and then evolve into a (quasi)periodic

vibration, characterized in the spectrum by a harmonic structure. The Spiral behaviour can be described as follows. From a uniform random distribution of colours, the HPM initially exhibits a behaviour resembling the Quasi-Synchronic behaviour mentioned earlier. Then, at some point, spiral-like patterns begin to emerge and propagate in the form of “waves”. These spiral waves often evolve to cover the whole CA grid. See Figure 62 and Video Example 8.

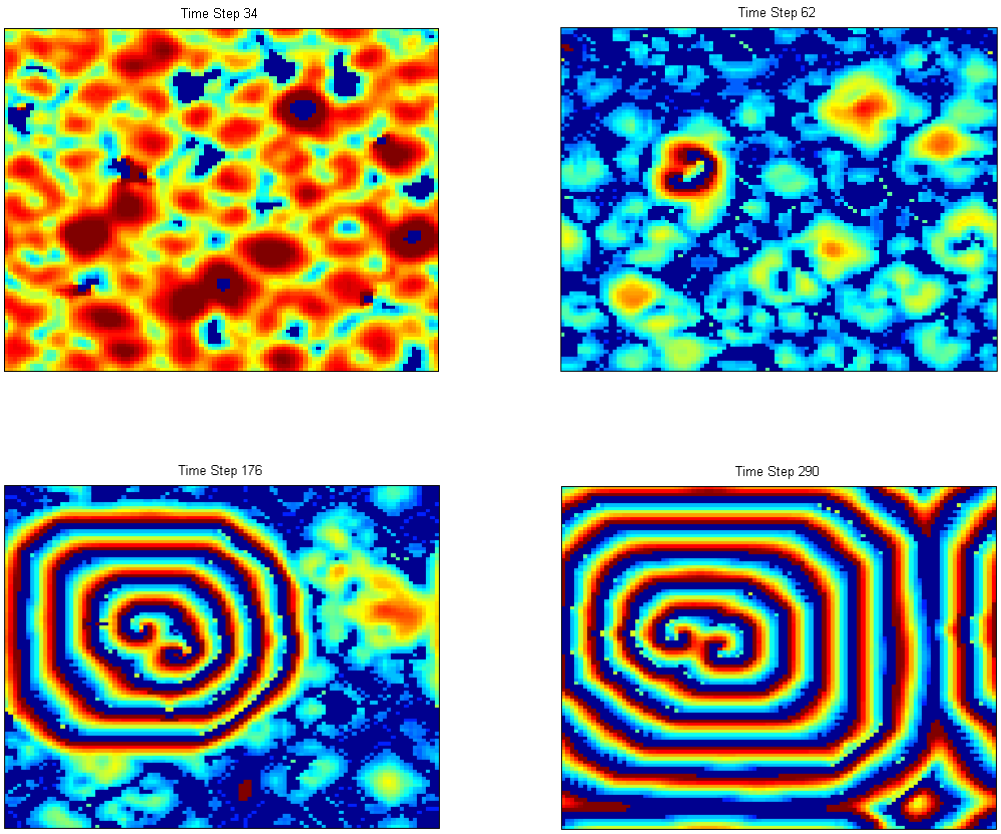

Figure 62: Spiral Behaviour.
100x100 automaton with Moore neighbourhood. V = 200. $r_1 = r_2$ = 6. K = 30.

The histogram analysis captures this evolution. The analysis produces a type of histogram sequences which remain essentially flat (resembling a noise spectrum) until the spirals emerge. At this point, some narrow bands will appear in the histograms (resembling a harmonic structure). These will increase until the spiral waves are covering the whole CA grid. At this stage, the narrow

bands stop increasing and settle into relatively stable values. This description ignores the first and the last histogram's bins which reach very high values and have always been discarded for the sound synthesis. See Figure 63 and Sound Example 18.

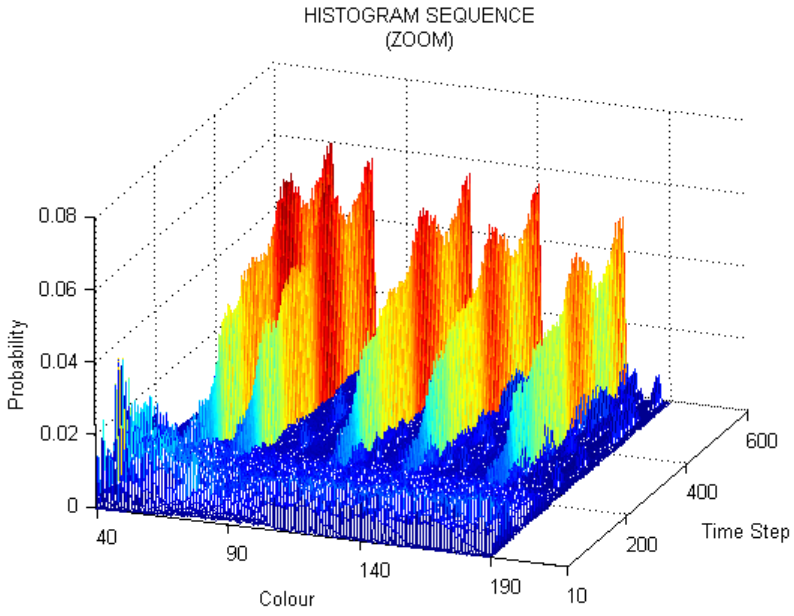

Figure 63: Zoom of the histogram sequence for the Spiral Behaviour of Figure 62.

As has already been mentioned, the Spiral behaviour produced by this thesis' modified HPM rules is better characterised by the histogram analysis than the same behaviour produced by the original HPM rules. Figure 64 illustrates this fact with a Spiral behaviour produced by the original HPM rules and with the parameter values suggested by the authors in the original paper. We can observe that the narrow bands of the histogram sequence, which characterise the spiral waves, are not as well developed as they are in Figure 63. The explanation of this phenomenon is not clear. As was mentioned, the only difference between the two sets of rules lies in the rounding processes. The number of spirals developed (three in Figure 64 vs. one in Figure 62) does not seem to be relevant here, since some experiments that produced multiple

spirals with the modified rules did produce prominent narrow bands in the histogram sequence. All in all, the cause-effect suggested here is not meant to be absolute. It is, in fact, impossible to survey all the parameter value combinations that produce Spiral behaviour. Although a considerable number of experiments have been carried out in this line, there may exist certain set of parameter values outside of the generalisation attempted here.

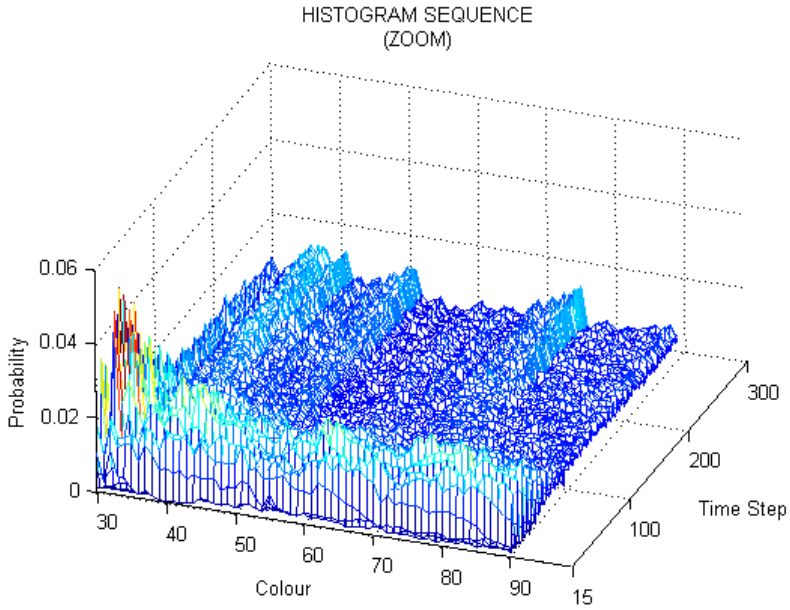

Figure 64: Spiral Behaviour with the original HPM rules.
(Top) 100x100 automaton with Moore neighbourhood. $V = 100$. $r_1 = 3$ $r_2 = 2$. $K = 20$.
(Bottom) a zoom of the histogram sequence.

The narrow bands obtained from this behaviour can be used for sound synthesis as is, without needing envelope estimation. Once the bands appear, their “amplitudes” can fluctuate, but they never return too close to zero as is found in, for example, the Quasi-Synchronic behaviour.

In conclusion, three behaviours of the HPM have been studied for sound synthesis with HMS: the Quasi-Synchronic, the Fast-Infection and the Spiral behaviour. Their histogram sequences have been studied with respect to the structural elements that resemble spectral components of sounds. It is important to note here that although a histogram analysis does not provide explicit information on the spatial dimension of the automaton, such information is, to a large extent, implicit in the histogram sequences. This is because the different types of emerging spatial patterns in the automaton produce different and characteristic histogram structures. Also, as has been mentioned, the spatial dimension is always available, and can be used, if desired, to complement the information gained from a histogram analysis.

5.4.3 Control

An important general aspect of controllability of the HPM is that the different behaviours that it can exhibit are determined by its parameter settings. Once a certain behaviour is achieved by a set of parameters, then, different runs with the same settings, but starting from different uniform random distributions of colours, will produce the same behaviour. Only a few exceptions to this have been encountered, all in relation to the Spiral behaviour. In those cases, a specific run expected to develop spirals did not develop such a pattern. HMS

benefits from this level of prediction, avoiding the unpredictability of cell values in favour of the predictability of behaviours. The particularities of different HPM runs produced with the same CA settings, as well as their applications for sound design, will be examined in the next subsection.

Regarding amplitude envelopes, the attack patterns that the automaton itself produces are remarkable. Sometimes the automaton also can produce an interesting decay between the attack and the sustain phases. Using the Fast-Infection behaviour, entire amplitude envelopes are generated. For other behaviours it is only needed to design releases patterns. This can be done by applying external envelopes and it is interesting the possibility of designing different envelopes for each structural element of the histogram sequence, in order to produce different release times for each component. Figure 65 shows two possibilities using the narrow bands of the Quasi-Synchronic behaviour, which will be considered as sound partials. The first makes the release times dependent on the amplitude of each “partial”, and the second makes them dependent on frequency (or histogram bin). See Sound Example 19.

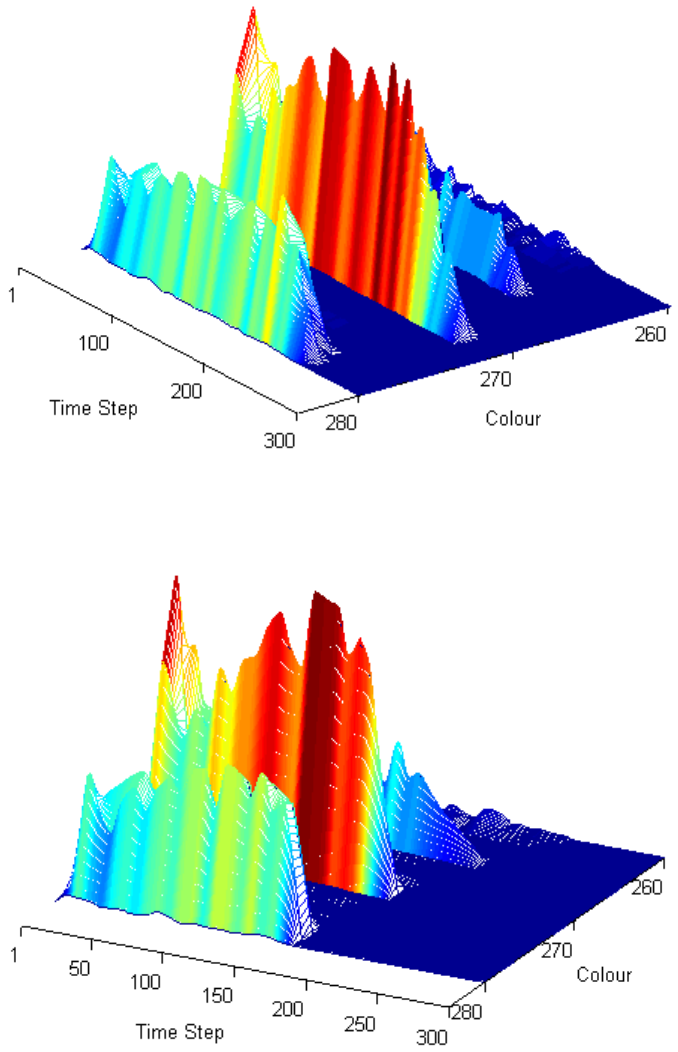

Figure 65: Different release times for each “partial”.
(Top) Release times dependent on amplitude.
(Bottom) Release times dependent on histogram bin.

It is worth mentioning another control possibility that evidences the potential of a DSP analysis for a better understanding of the self-organization process of CA. Once the Fast-Infection behaviour was analysed and an explanation of it was found with regards to the size of the neighbourhood and the fact that the divisor of Rule 2 is lower, it was possible to obtain the same type of behaviour with a Moore neighbourhood through modifying the rules. Simply by dividing the *A* and *B* parameters by two, it has been possible to obtain Fast-Infection behaviour using the Moore neighbourhood.

The next two subsections describe in more detail two more control mechanisms for sound design.

5.4.3.1 Invariance Property

An invariance property has been found by studying different runs of the HPM with the same settings but starting from different uniform random distributions of cells values. As we have seen before, it is assumed that the automaton will exhibit the same behaviour every run (i.e., the spatial patterns that the automaton develops are of the same type every run, but different from each other). However, the level of prediction goes beyond this. By looking at a zone of narrow bands in the histogram sequences, such as the NB2 zone of the Quasi-Synchronic behaviour, it is observed that the locations of the narrow bands remain identical in all runs. That is, in all runs the automaton self-organises through the same set of predominant colours. The relative “amplitudes” of the narrow bands remain similar. The time-varying “amplitude” evolution of each narrow band is slightly different every run. This invariance property remains largely present even for different sizes of the CA grid. Figure 66 illustrates all this graphically. We can also see that the smaller automaton (Run 3) produces more abrupt amplitude fluctuations. We could find an explanation of this on the weight with which every cell contributes to the histogram, which, as we already know, inversely depends on the automaton size.

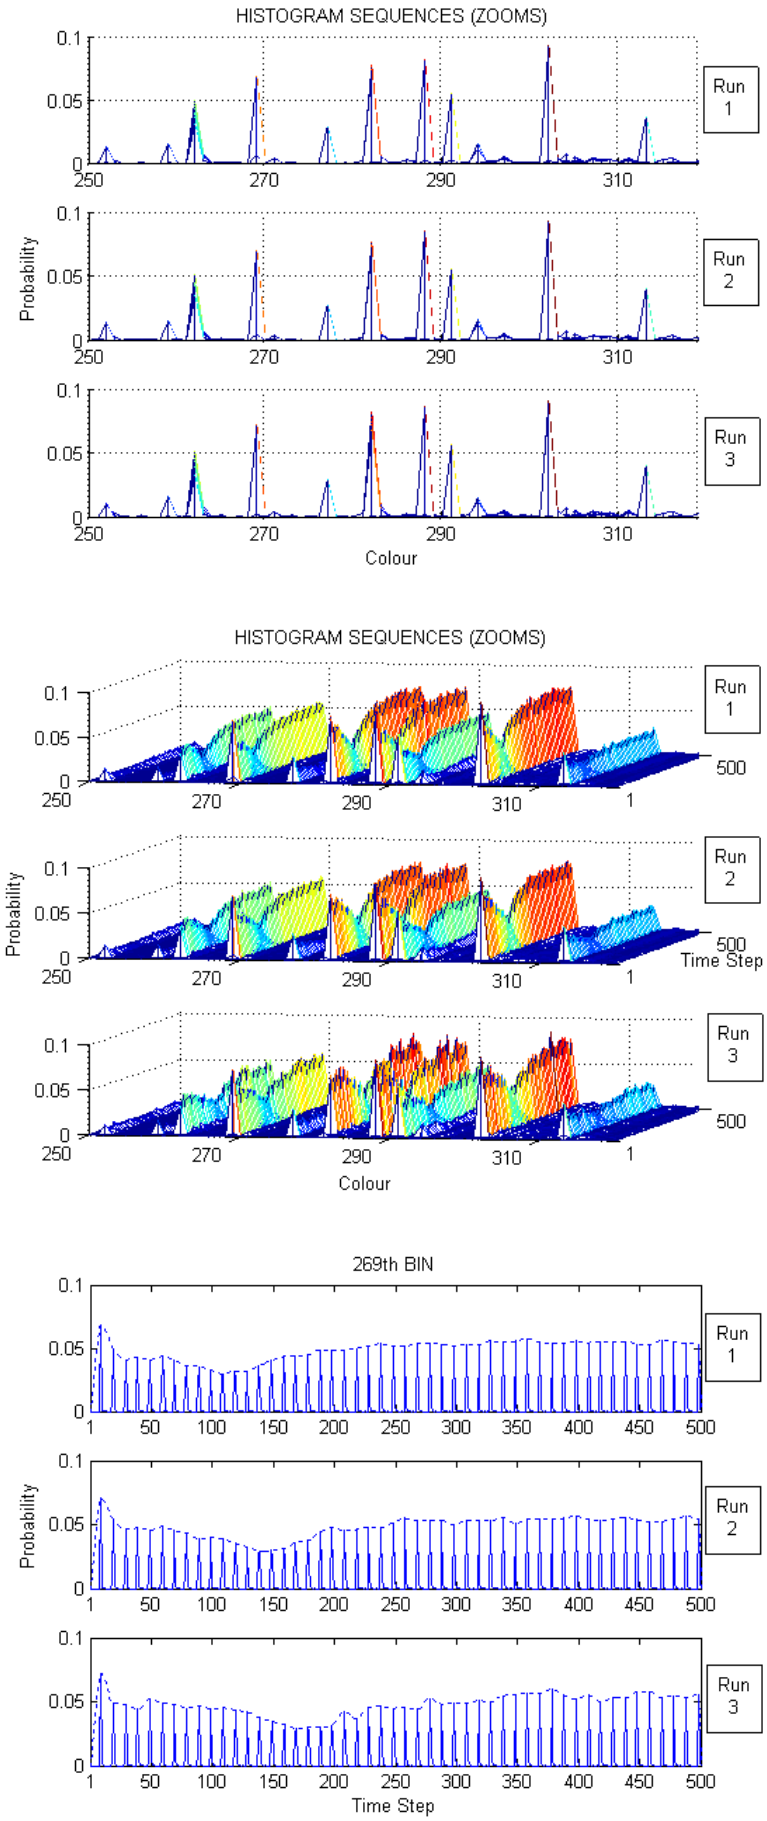

Figure 66: Invariance property.
Three histogram sequences from three different CA runs: Run2 has the same CA definition as Run1 but starts from a different initial configuration (uniform random distribution). Run3 has the same CA definition as Run1 and Run2, but with different CA size, where the CA size for Run1 and Run2 is 200x200, and the CA size for Run3 is 100x100 cells. (Top) Frontal plot. (Middle) Lateral plot. (Bottom) Time evolution of the 269th bin.

The level of prediction available through this invariance property is useful for instrument design in Automatic Mode. Once a timbre has been designed in Manual Mode, all the processes developed in this design stage can be replicated in order to automatically obtain multiple instances of the original type of sound (with different pitch, intensities and durations, as desired). All instances will share the same spectral structure, but with differences in the time-varying amplitudes of the spectral components. In conclusion, we can design instruments able to produce similar but not identical sounds.

5.4.3.2 Control by Coupled CA

This subsection describes a control mechanism for controlling over time the HPM by another cellular automaton, a growth model (GM). In a nutshell, this method firstly explores the possibility of evolving different simultaneous HPM behaviours within the same automaton by bringing into play different sets of parameter values. The number of parameter sets will be restricted to two. Following successful results, the next step will be to add dynamism to the process using a GM. The cells of the GM will have only two states and such an automaton will delimit two dynamic zones in the HPM, each governed by a different set of parameter values. Two HPM behaviours have been used in these experiments, the Quasi-Synchronic and the Spiral. Among all the possibilities that this control mechanism affords for the synthesis of dynamic sounds, highlights its use with the Quasi-Synchronic behaviour in order to make the attack portion of a sound dynamically more complex than the rest of the sound.

Firstly, a series of successful experiments were conducted to evolve two simultaneous behaviours in the HPM. This was done by delimiting two halves in the automaton each of which governed by a different set of parameter values. Only different values for r_1, r_2 and K were established; V was the same on both halves. From the experiments, it is observed that since the two behaviours are not isolated they are not independent, and it is visually noticeable that they influence each other. On many occasions it seems that one of the two behaviours has a stronger influence over the other. For instance, if we run a Quasi-Synchronic behaviour in one half and a Spiral behaviour in the other, the spiral waves emerging on one side of the automaton usually propagate to the other side but with a noticeable visual change so that we can differentiate two halves in the automaton. See Figure 67 and Video Example 9.

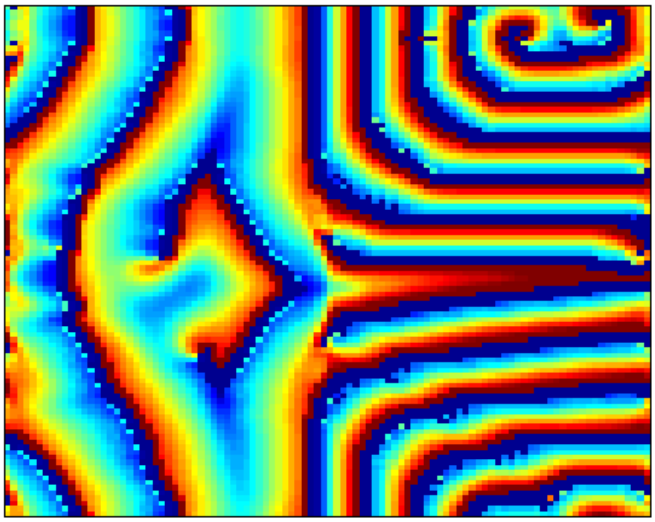

Figure 67: Simultaneous behaviours in an HPM.
The left half has the rule parameters adjusted to produce the Quasi-Synchronic behaviour whereas the right half has the rule parameters adjusted to produce the Spiral behaviour.

However, independently of the visual effect, what is interesting for us from the fact that each half is governed by a different set of parameter values is that

each half will evolve through a different set of predominant colours. This fact is better registered in the histograms than from visual perception. For instance, if we evolve the same type of behaviour in both halves, for example the Quasi-Synchronic, but with different sets of parameter values, it can happen that at the beginning we can visually differentiate the two halves but after a while, the influence of one side over the other makes that the automaton has only one visual appearance as a whole; we cannot distinguish two halves. See Figure 68 and Video Example 10.

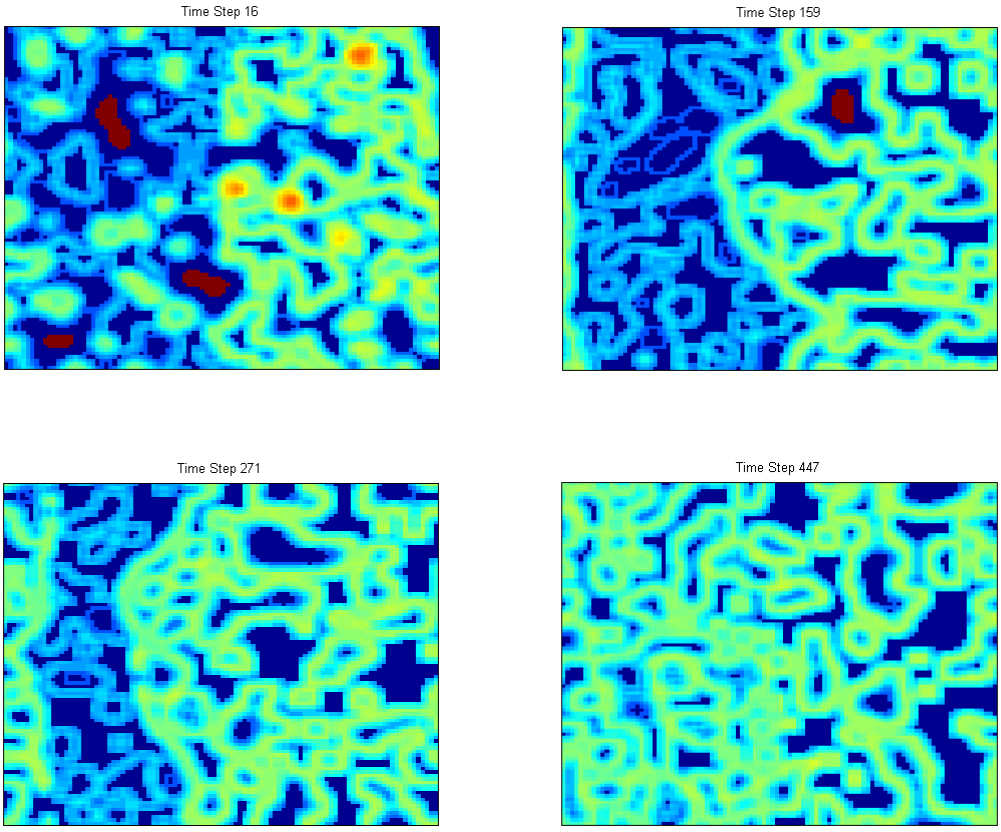

Figure 68: Two simultaneous Quasi-Synchronic behaviours.
In the beginning the left and the right halves are visually differentiated, but eventually, the two halves are not visually distinguishable.

However, from a histogram analysis we can always identify two superposed histogram sequences. This can be checked by examining the left-most narrow band of the NB2 zone, which as has been discussed, characterises Quasi-

Synchronic behaviour. In this experiment, two of these narrow bands are present all through the run, as we can see in Figure 69.

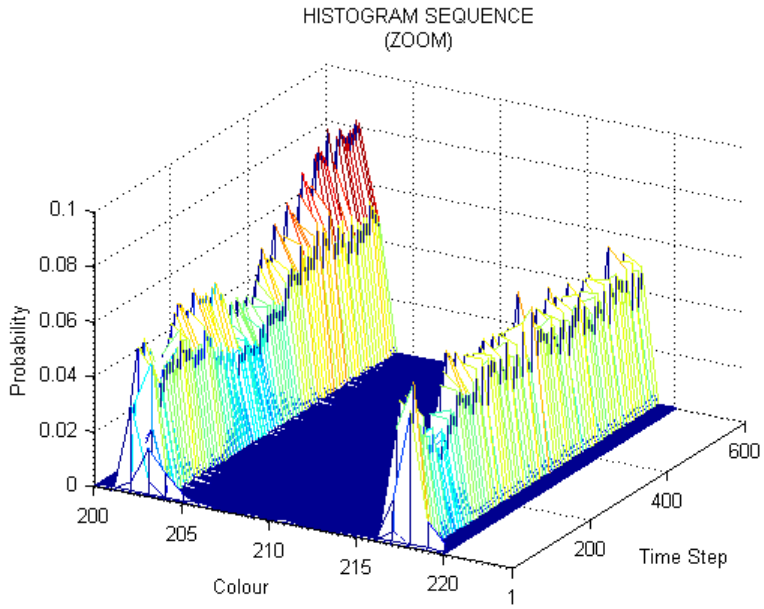

Figure 69: Histogram analysis revealing two simultaneous behaviours all through the run.

The next step to further explore the possibility of producing different simultaneous histogram sequences is to add dynamism to the process using a GM. The GM used here is the one developed in Subsection 5.3.3.1 for the "Attack Solution 1", as an extended version of the multitype voter model. Only two cell states are defined, empty and occupied. The GM is used to create and evolve different regions of parameter values for the HPM. Thus, the GM will control the HPM by determining the parameter values that govern each of its cells. This is accomplished as follows. Only two different sets of HPM parameter values are brought into play, one associated with empty GM cells (set of parameter values E) and the other associated with occupied cells (set of parameter values O). Both CA have the same size and therefore if a GM cell is empty that means that the cell located on the same coordinates in the HPM will

be governed by the set of parameter values E. On the other hand, if one cell of the GM is occupied, the corresponding cell in the HPM will be governed by the set of parameter values O. From now on, both two sets of parameter values are adjusted to produce the same type of behaviour so as to work with the same type of histograms.

The GM evolves through predictable transitions. An automaton like this with two states allows two types of transitions. The most immediate one is going from the set of parameter values E to O (Figure 70).

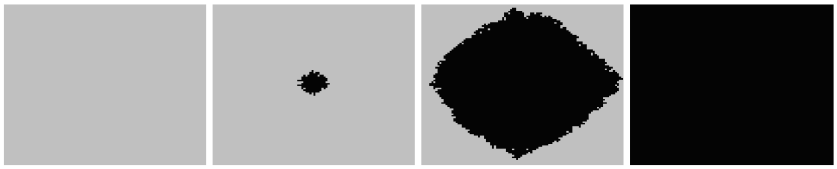

Figure 70: A GM evolution starting with all cells empty.
At some point an occupied cell is introduced in the middle of the CA lattice which will grow covering the whole CA grid.

That will produce a smooth transition between two HPM behaviours and will lead to the production of a cross-fade sound effect (Figure 71).

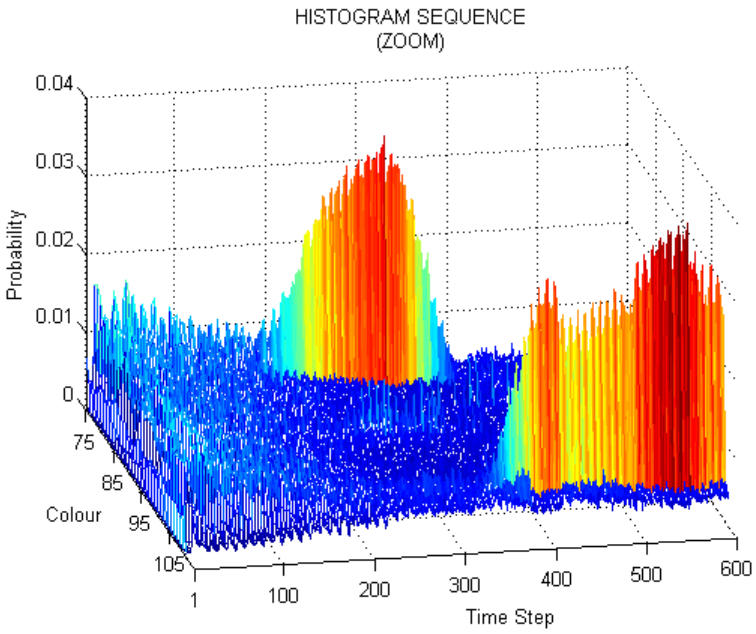

Figure 71: Cross-fade transition.
Zoom of a histogram sequence of an HPM controlled by a GM.
The HPM exhibits a transition between two consecutive Spiral behaviours.
The zoom shows a cross-fade transition between two narrow bands.

The second type of transition is of particular interest to sound synthesis. It consists in going from two simultaneous sets of parameter values, O & E, to O solely. An example of such a transition is illustrated in Figure 72.

Figure 72: A GM evolution starting with occupied cells in the left half and empty cells in the right half.

This implies starting with two superposed histogram sequences and ending up with only one. As a result of this, using the Quasi-Synchronic behaviour, in the NB2 zone there will be more narrow bands in the beginning of the run than after the GM has completely evolved. This process parallels the behaviour of sounds produced by acoustic instruments in regards to the fact that the spectral complexity in the attack is normally higher than in the rest of the sound. In

addition, this method simulates energy transfer between the partials. See Figure 73 and Sound Example 20.

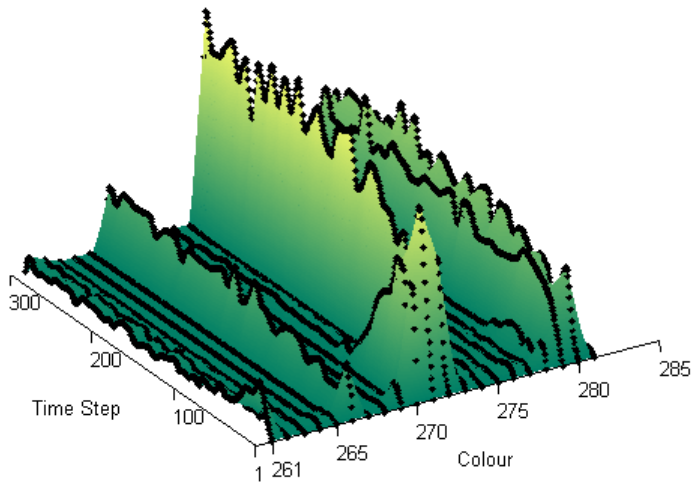

Figure 73: Zoom of an HPM histogram sequence after envelope estimation. The HPM exhibits two simultaneous Quasi-Synchronic behaviours and is controlled over time by a GM. The bins of the histogram sequence can represent the time varying amplitudes of sound partials, some of which are short-lived in the attack portion. We can also observe amplitude compensations between the disappearing “partials” and the permanent ones.

5.4.4 Cell Seeding Experiments

A number of preliminary experiments have been carried out regarding the initial cell seeding of the HPM. The motivation is to explore the potential of the HPM for producing different types of histogram sequences by initialising the automaton with configurations different to the usual uniform random distributions. One focus has been on initial configurations obtained by drawing pictures or directly from photos.

This possibility was not explored with the MVM automaton because in MVM only those colours that are present in the initial configuration will evolve during

the run. However, the HPM can be seeded with any distribution of colours and then, its cyclic nature will make the automaton evolve through certain predominant colours, some of which may not be present in the initial cell seeding.

The initial results have been promising and suggest that new types of HPM histogram sequences can be developed by means of different cell seeding. Figure 74 shows the result of evolving a 256 grey-level image consisting of a white background with a black square in the centre. The automaton's rule parameters were set up for the Quasi-Synchronic behaviour. However, the resulting histogram sequence is novel in interesting ways, compared to previous plots.

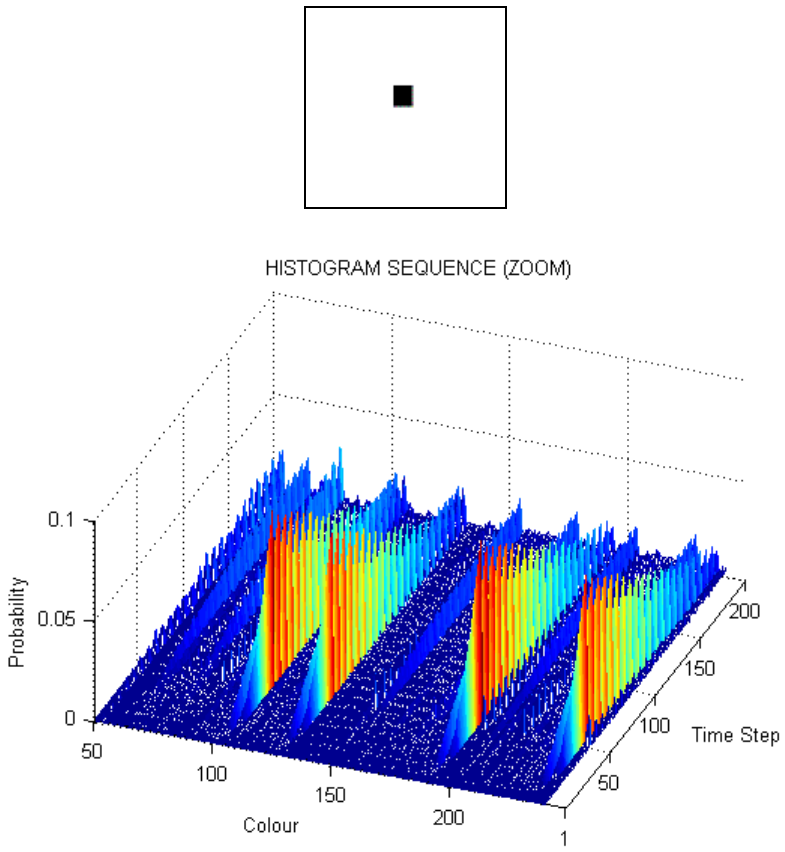

Figure 74: Different cell seeding produce different histogram sequences.

In the context of HMS, one important reason for seeding the automata with random distributions of colours is because it is immediate to produce different distributions which could produce similar but not identical sounds. In this sense, it may be interesting to devise an algorithm for modifying these new initial configurations (images, pictures, photos, etc.) so as to also produce similar but not identical sounds.

5.5 Frequency Trajectories

We have seen how to produce amplitude trajectories with HMS. However, sounds produced by acoustic means also involve frequency trajectories. Frequency trajectories in HMS can be generated by the automaton and/or by the user as will now be shown.

Regarding the latter possibility, leaving the generation of amplitude trajectories to the automaton, any kind of frequency trajectories can be specified or generated by the user. The possibilities could range from the design of quasi-random frequency fluctuations, which are often found in sounds produced by acoustic instruments, to the production of glissandi that the user can specify at will.

Frequency trajectories can also be generated by the automaton itself. Different automatons provide different possibilities and the discovery and practical use of these is a matter of experimentation. One approach that has been tried, with both, the MVM and the HPM, is modulating the angle of the sinusoidal components (using for example, frequency modulation) with their same

amplitude trajectories. The results range from interesting frequency fluctuations, which can be obtained from using amplitude trajectories similar to those found in Figure 73 (see Sound Example 20), to glissando gestures using, for example, the amplitude trajectories of Figure 43 (see Sound Example 21).

Finally, another possibility is to identify “frequency” trajectories in the histogram sequences. Interesting results have been obtained using modified versions of the HPM. See Figure 75.

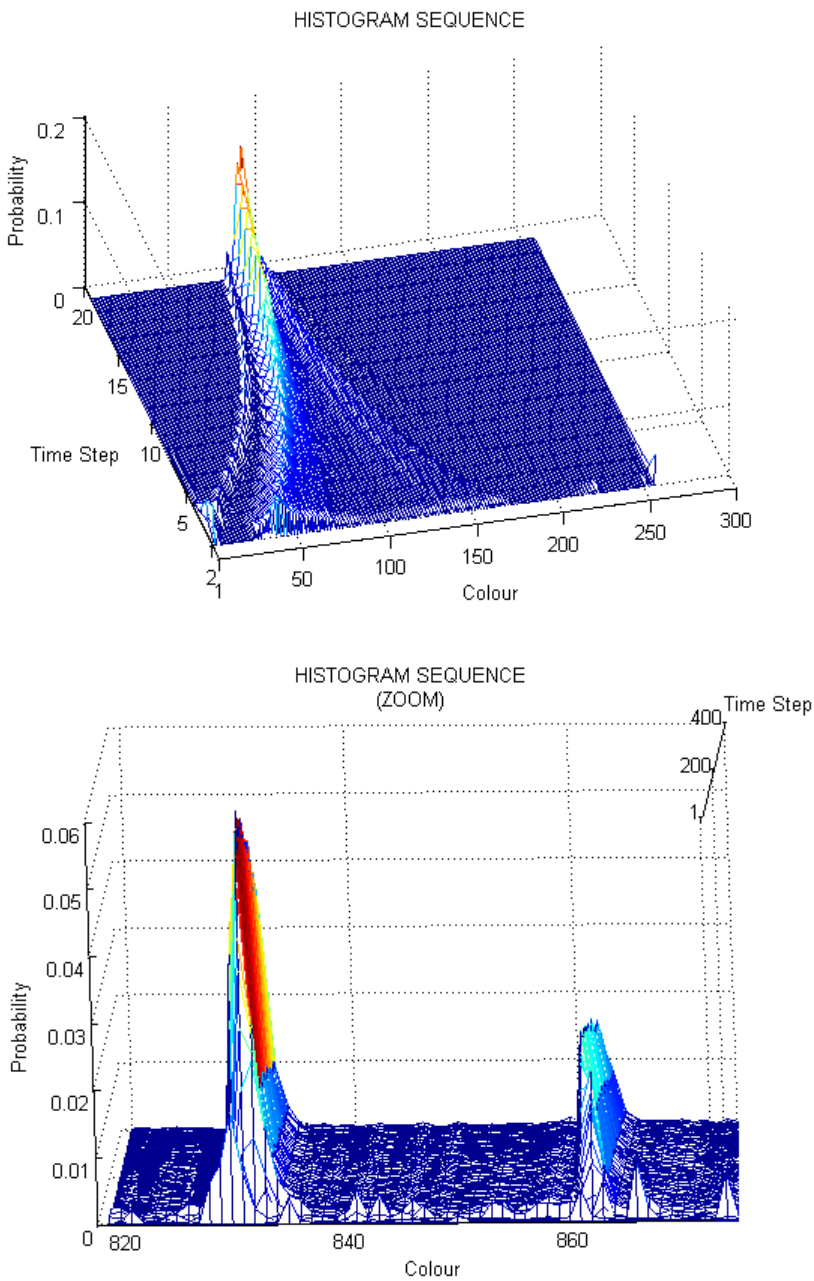

Figure 75: “Frequency” trajectories found in the histogram sequences.
(Top) Convergent glissandi.
(Bottom) Glissandi in the “attack” portion of the narrow bands of an NB2 zone.

5.6 HMS with Totally or Partially Unsuitable CA

One should not expect successful results from every automaton when using HMS. We will now examine this issue through the use of two case studies. In the first study, using the plurality vote rule, the first stage of developing a

mapping process will be possible because the automaton produces interesting histogram sequences. Therefore, the synthesis of interesting dynamic sounds will be possible. However, on the other hand, the design of control mechanisms for sound design will be hindered by the nature of this automaton. In the second case study, using the rug rule, not even a mapping process will be possible because no clear resemblances with spectral components of sounds will be found in its histogram sequences.

5.6.1 HMS with the Plurality Vote Rule

The next automaton to be considered for sound synthesis with HMS belongs to the same family as the MVM, that is, those CA that can be interpreted as voter models. It was named plurality vote (PVo) by David Griffeath in (Griffeath 1994), where the automaton is developed as a multi-state extension of the binary model called majority vote.

In the original majority vote model the opinion of each individual is replaced by the majority opinion in its neighbourhood. The PVo rule extends the number of possible opinions from two to as many. When estimating the majority opinion, in case of a tie, the individual in question will not change its opinion. In other respects, most of the experiments in this thesis have been carried out using a Moore neighbourhood of range 4, that is, a 9x9 box. The PVo has been implemented with periodic boundary conditions.

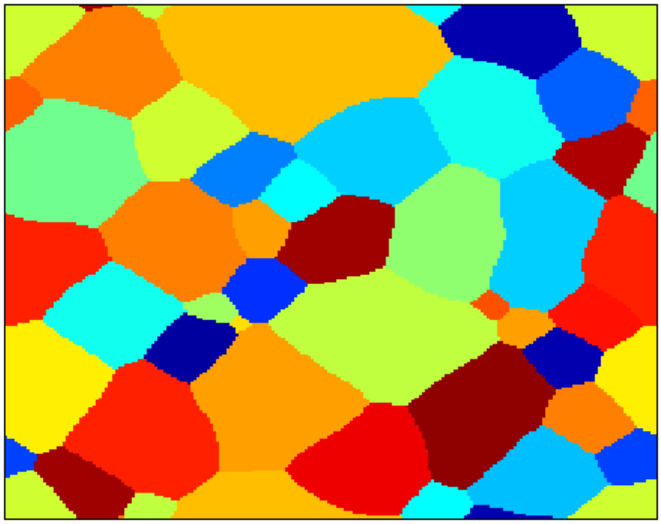

Figure 76: A configuration of the PVo showing homogeneous areas.

From a uniform random distribution of cell values, or colours, the PVo self-organizes into homogeneous areas of a single colour as shown in Figure 76. As the rule is iterated some areas will increase in size at the expense of others that can disappear. This dynamic resembles the one of the MVM but with the key difference that the PVo not always reaches consensus. On the contrary, usually the automaton fixates eventually, meaning that a certain achieved CA configuration remains constant thereafter. An explanation for this phenomenon seems to lie in the fact that the PVo approximates ‘a nonlinear partial differential equation called motion by mean curvature (mmc) that exhibits surface-tension clustering over time’ (Griffeath 1994). This “freezing” of the automaton is its main disadvantage in the design of control mechanisms in the context of HMS.

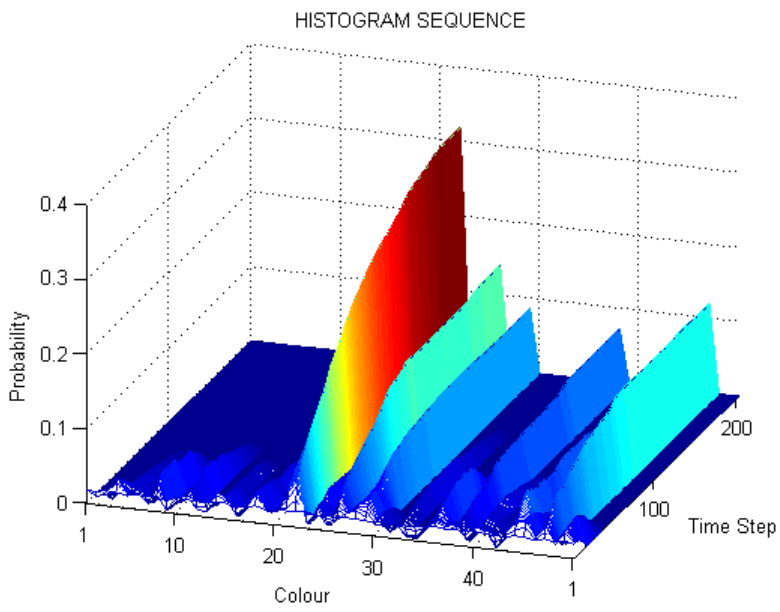

Figure 77: A histogram sequence of a PVo evolution.

Nevertheless, the histogram sequences obtained are suitable for sound synthesis. Figure 77 serves to illustrate the type of structures that can be obtained from a histogram analysis of a PVo evolution. We can see that, like the MVM histogram sequences, the disappearance of colours will lead to the production of sounds with a dynamic spectrum and with greater initial spectral complexity. Perhaps, one difference compared to the MVM is that amplitudes of "partials" are smoother; probably because the coloured areas of the PVo automaton are homogeneous and increase/decrease smoothly across their boundaries.

A common sound synthesis practice has been the application of external amplitude envelopes for attacks and/or releases. Using this, interesting plucked string sounds have been synthesised (see Sound Example 22). It is worth mentioning that on one occasion, the PVo did not fixate but did achieve consensus (see Figure 78). As a result, the synthesis of damped sounds was possible without the need to apply external envelopes through discarding the

“partial” corresponding to the colour that covered the whole automaton grid. See Sound Example 23.

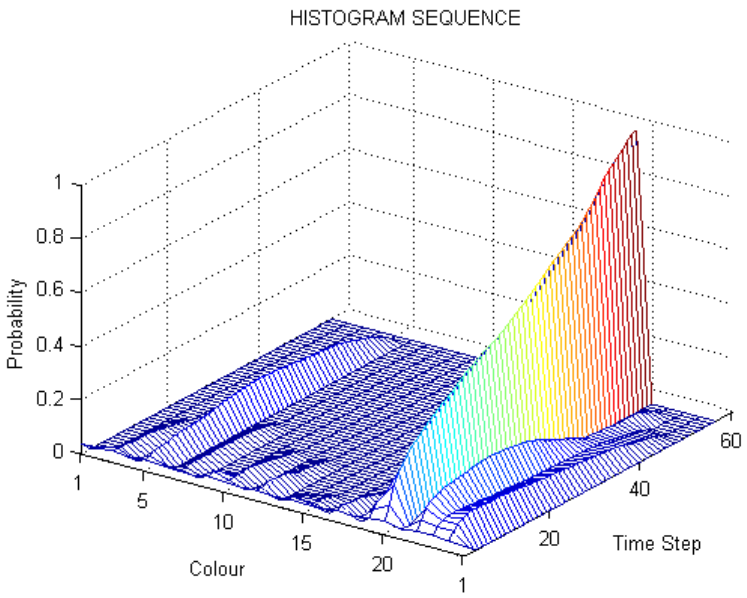

Figure 78: PVo achieving consensus.

In conclusion, the PVo is an interesting automaton for HMS sound synthesis. However, the fact that the automaton normally fixates at some point in its evolution makes it unfeasible to design any control mechanism. Thus, the automaton is not reliable for Real-Time modes of use. Nevertheless, the histogram sequences are very interesting and therefore, the automaton is suitable for experimental practice in Non-Real-Time modes. For future research it may be interesting to follow the direction suggested by David Griffeath (personal communication), regarding the prevention of fixation by means of randomness or a nonlinear ingredient in the PVo dynamics. Successful results along these lines make an investigation on control mechanisms for sound design more feasible.

5.6.2 HMS with the Rug Rule

The last automaton to be considered in this chapter exemplifies the fact that not every automaton is suitable in HMS sound synthesis. The rug rule was designed by Rudy Rucker and has been implemented according to the following definition:

> ‘Rug rules are averaging rules using the full range of 256 possible states. To update itself in a Rug rule, every cell takes four steps. 1) Every cell calculates the sum of its eight nearest neighbours' states. 2) Every cell calculates the average neighbour state by dividing the sum by eight and throwing out any remainder. 3) Every cell computes its new state by adding an Increment (usually the increment is 1) to the average neighbour state. 4) As a final step, new state is taken modulo 256’ (Wójtowicz 2001).

The rug rule is a cellular automaton inspired by physics as the author explains in the following:

> ‘Rug produces an effect which has been compared to boiling. If all of a cell's neighbors are at the maximum value of 255, then that cell's new value will be 256, which gets wrapped down to zero. At the next generation, the presence of this zero-valued cell will lower the values of that cell's nearest neighbors. The process is analogous to the way in which a hot enough region of water gives up some heat by forming a bubble of steam. The water right around the steam bubble cools off for a moment’ (Rucker 2010 p.50).

Figure 79 shows a configuration of the automaton which exhibits the abovementioned boiling behaviour.

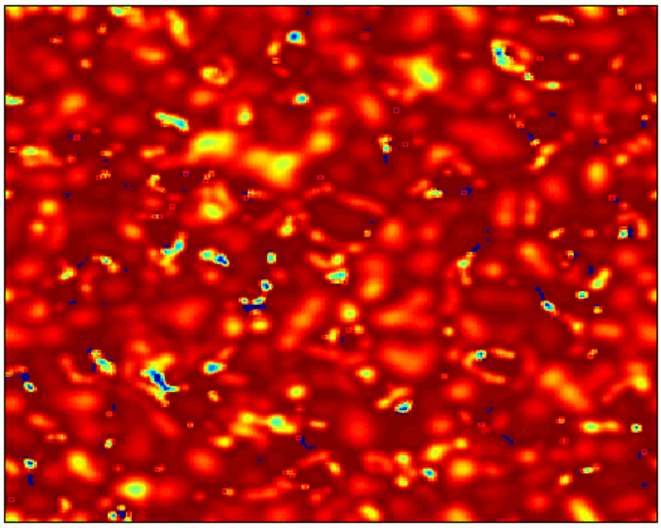

Figure 79: A configuration of the rug rule showing the boiling behaviour.

Figure 80 shows a histogram analysis of this behaviour.

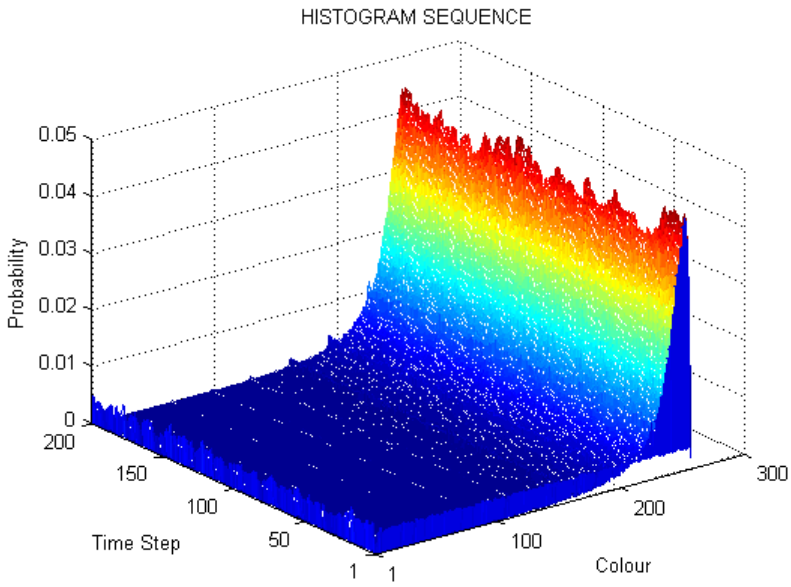

Figure 80: A histogram analysis of the boiling behaviour.

As we can see, no clear resemblances with spectral components of sounds can be found in this histogram sequence. Different variations of this automaton have been explored, such as different CA boundaries (periodic and fixed zero) and different increments (step 3 of the rule). However, results achieved have been no better. Nevertheless, the crest of the wide band on the right of the above

histogram sequence incorporates interesting oscillations over time that may be useful for other musical purposes.

Chapter 6 Conclusions

This chapter contains the following sections: conclusions, a summary of contributions to knowledge, and suggestions for further research. All three sections will be developed in numbered concise paragraphs for the sake of organization and clarity.

6.1 Conclusions

1. In Chapter 4 the main research problem that this thesis tackles was formulated in the following terms: the development of a sound synthesis technique based on CA capable of allowing a sound design process. A solution to this problem has been developed in Chapter 5, with the invention and development of a new spectral technique called Histogram Mapping Synthesis, which is capable of allowing a reasonable sound design process.

2. HMS is a result of a novel approach, introduced in Section 5.1, which proposes considering the output of CA as digital signals and using DSP analysis methods to better understand the self-organization process. Such an approach was proposed in order to identify mapping possibilities (between the CA domain and the synthesis domain) for making the synthesis of sounds possible, and control possibilities for allowing a sound design process to take place. The DSP method on which HMS is based is the statistical analysis of CA evolutions by histogram

measurements. This proved to provide very useful information on the evolution of multi-state CA. This CA analysis results in a histogram sequence that can be displayed as a 3D plot. On the basis of this DSP method, it has been possible to identify mapping possibilities for sound synthesis and control possibilities for sound design, as will be detailed below. A general description of the process is as follows. The CA provide structured numerical data. This data is analysed by histogram measurements which hopefully will produce characteristic structures in the histogram sequences. The CA behaviour is captured and characterised by these histogram structures. The spirit of HMS is to design mappings based on resemblances of the structural elements of the histogram sequences with spectral components of sounds. Based on these resemblances, the histogram sequence structural elements can be combined and manipulated in different ways in order to design different sounds. Different control mechanisms have been also designed for controlling the synthesis process and allowing a sound design process. It should be noted that for virtually all the mappings and virtually all the control mechanisms, the function of the CA is the same: to provide structured numerical data to be analysed. There can be some exceptions to this generalisation. For instance, the control mechanism by coupled CA, where the CA growth model has the function of controlling another automaton.

3. Controllability is the main criterion established in Chapter 4 to determine the sound design possibilities of a CA-based system. And it is the main characteristic of HMS, since a considerable number of HMS-based

control mechanisms have been identified and developed. For instance, in Subsection 5.2.1 a control possibility for sound design was presented which could be applied regardless of the automaton used. It was based on the idea of making decisions before rendering the sounds. Thus, an offline sound design process can be developed with the structural elements of the histogram sequences. As well as this, other control mechanisms have been identified and developed specifically for two of the automatons presented in this thesis: the multitype voter model and the hodge podge machine. These are introduced in Subsections 5.3.3 and 5.4.3 respectively. The achieved level of controllability demonstrates that HMS allows a reasonable sound design process.

4. Controllability is a concept that goes hand in hand with predictability. In terms of predictability, HMS largely avoids the unpredictability of CA cell values in favour of the predictability of CA behaviours with those CA for which the histogram sequences are a means of characterising their behaviours and, whose behaviours are predictable by their parameter values. This level of predictability allows a special type of sound design, namely, instrument design. This involves designing a timbre and being able to reproduce it with different pitches, intensities and durations. Instrument design can be boosted with specific CA in which it is possible to predict the location and characteristics of the structural elements within histogram sequences resulting from different CA runs. One example of this feature is the Invariance Property studied in Subsection 5.4.3.1.

5. HMS also fulfils all four criteria introduced in Chapter 4 for producing a variety of timbres, namely, allowing the use of a variety of: CA, mappings, synthesis engines and types of synthesis. It is highlighted the usage of a variety of synthesis engines because none of the previous systems studied in Chapter 3 provides this capability. In the next four paragraphs, HMS will be evaluated according to these four criteria.

6. Regarding the usage of a variety of CA, HMS can be used with different automatons. This thesis has explored the possibilities of three CA that have offered interesting results with HMS. These were the multitype voter model (Section 5.3), the hodge podge machine (Section 5.4) and the plurality vote rule (Subsection 5.6.1).

7. Regarding the usage of a variety of mappings, HMS is open to different mapping possibilities depending on the resemblance that can be found between structural elements of the histogram sequences and certain sound spectral components. For example, see Subsections 5.3.2 and 5.4.2. These possible mappings based on such resemblances are distinctive because in most other synthesis techniques based on CA there is not an intuitive correspondence between the components of the automaton and the components of a sound. Additionally, decisions taken in the offline sound design process are normally related to mapping actions such as, for instance, the assignment of frequencies to the structural elements of the histogram sequence. The explanations of Appendix A: Sound Examples, give insight into mapping-related control possibilities for sound design.

8. A variety of synthesis engines are useable with HMS, because the structural elements of the histogram sequences can be used as control data to drive different synthesis techniques. In this thesis sinusoidal additive synthesis has been the main approach, with and without angle modulations. There has also been the use of subtractive synthesis by means of FFT Convolution filtering. It may be possible to consider other synthesis engines but they were not investigated in this thesis.

9. The last criterion concerns the type of synthesis: pure synthesis or audio processing. In this thesis most of the experiments were around pure synthesis. However, audio processing is possible as well, for instance with the envelopes of the complex dynamic beats studied in Subsection 5.3.3.3.2.

10. HMS also stands out from other techniques because of its different modes of use: Manual or Automatic, Real-Time or Non-Real-Time and Interactive or Non-Interactive. Manual and Automatic Modes refer to the mapping process between the CA domain and the synthesis domain. In Manual Mode there is user involvement at the point between the CA run and the sound synthesis, whereas in Automatic Mode there is not user involvement at the point between those two stages. Real-Time and Non-Real-Time Modes refer to whether the sound is rendered while the automaton is evolving or after the run is completed. Interactive and Non-Interactive Modes refer to whether there is user control over parameters (of the automaton or the synthesis engine) on-the-fly, while the sound is

being rendered. These modes of use were studied in more detail in Subsection 5.2.2

11. Intuition and creativity have been determinant in the development of this thesis. For instance, the main DSP approach was based on initial intuitions for dealing with general problems, as explained in Section 5.1. Also, intuition and creativity have been relevant in defining the synthesis goal and evaluation in some of the sections of this thesis. For instance, the decision of exploring textures based on dynamic and complex beats (presented in Section 5.3.3.3.2) was taken after experimental sessions rendering sound textures with manipulated histogram sequences. Thus, the first seconds of Sound Example 13 were the inspiration for achieving this aesthetic goal. Also intuition and creativity have been relevant in the conception of some of the control mechanisms presented in this thesis such as the control by Coupled CA (presented in Section 5.4.3.2).

12. The sounds produced with HMS range from those that are novel to those that are imitations of sounds produced by acoustic means. All the sounds obtained present dynamic features and many of them, including some of those that are novel, retain important characteristics of sounds produced by acoustic means. As an example of this, some control mechanisms lead to the production of sounds with more complex spectra in the attack portion than in the rest of the sound. This parallels sounds produced by most acoustic instruments. This was possible by achieving a control over time of the spectral complexity of the sound. Mechanisms of this kind include those found in Subsection 5.3.3.2 with the multitype voter model or Subsection 5.4.3.2 with the hodge podge machine.

Further interesting results along these lines are the brass-like attacks of Subsection 5.3.3.1, and the energy transfer between partials that parallels the process observed in bell sounds (Subsection 5.3.3.3.1).

13. As a final point, most of the control mechanisms allow the production of sounds which are similar but not exactly identical, behaving in this sense like an acoustic instrument. As an example to clarify this point, let us consider a performer of an acoustic instrument who is asked to play the same melody several times. All the performances would sound similar. However, if a spectral analysis were carried out, it would reveal subtle differences between the performances because they are not identical. Chareyron expressed the desirability of having this capability in his system as follows:

    ‘In musical applications it may be a good idea to add small random modifications to the wavetables to allow the repetition of similar but not identical tones’ (Chareyron 1990 p.30).

6.2 Summary of Contributions to Knowledge

1. Proposed a novel DSP approach for the design of CA-based synthesis techniques (Section 5.1). Such an approach consists in considering the output of CA as digital signals and using DSP procedures to analyse them. It opens a variety of possibilities for attaining better understanding of the self-organization process of CA with the view of identifying, on the one hand, analogies with the sound synthesis domain which can be the basis of potential mappings for making possible the synthesis of sounds

and, on the other hand, control possibilities for allowing a sound design process, something most existing CA-based systems lack.

2. Developed Histogram Mapping Synthesis, a synthesis technique capable of allowing a reasonable sound design process (Chapter 5). HMS is a result of the abovementioned approach. It is based on a statistical analysis of CA evolutions. In HMS the functioning of a multi-state automaton is considered as a sequence of digital images, each of which is analysed by means of histogram measurements. Such a DSP analysis produces a histogram sequence that can be displayed as a 3D plot. This thesis documented that from an appropriate automaton, in the histogram sequences it is possible to identify time-varying structural elements resembling sound spectral components. These structural elements can be used as control data to drive different synthesis techniques. Depending on the resemblance of such structural elements, different mappings onto appropriate synthesis parameters can be established. As well as mapping possibilities based on these analogies, HMS offers control possibilities as explained below.

3. Identified and developed a considerable number of control mechanisms within the context of HMS (Subsection 5.2.1, Subsection 5.3.3, and Subsection 5.4.3). For instance, in Subsection 5.2.1 a control possibility for sound design is presented which can be applied regardless of the automaton used. It is based on the idea of making decisions before rendering the sounds. Thus, an offline sound design process can be developed with the structural elements of the histogram sequences.

Additionally, other control mechanisms have been identified and developed specifically for two of the automatons presented in this thesis: the multitype voter model and the hodge podge machine. These are introduced in Subsections 5.3.3 and 5.4.3 respectively. The achieved level of controllability demonstrates that HMS allows a reasonable sound design process.

4. Developed a textural concept based on dynamic and complex sound beats (Subsection 5.3.3.3.2). This textural concept is based on the control of the amplitude trajectories of a number of slightly detuned sinusoidal components. The amplitude trajectories are created using HMS and a modified version of the multitype voter model. The relative amplitudes of all the beating partials change constantly making the interference patterns change over time continuously, gradually and non-deterministically. In this way, the dynamic complex beats are produced. These beats are perceived by ear as a “bumpy” texture, in contrast to the invariable rhythmic sequence that common synthetic beats produce. An important feature is that the amplitude trajectories are compensated in such a way that the sum of all of them is constant and equal to one. This makes the resulting sound have constant overall amplitude that can be modified at will by the user. Endless duration and Real-Time Mode are potentially possible. The modified version of the MVM is the result of a control mechanism based on a DSP monitoring of the automaton which informs us on when and how to alter its autonomous evolution. Furthermore, it was shown that this textural concept can be also applied to existing sounds by means of audio processing.

5. Achieved the production of similar but not identical sounds (for example in Subsection 5.3.3.2). Most of the control mechanisms developed in the context of HMS allow this production of similar but not identical sounds, behaving in this sense like an acoustic instrument. Chareyron expressed the desire of having such capability in his system (as mentioned in the previous section, with his quotation).

6. Identified the key role of the weight with which every cell contributes to the normalised histograms (Section 5.2). This weight inversely depends on the CA size and implies that a change of colour of a cell in a small automaton will cause a larger variation or a more abrupt change between the corresponding pair of histograms than in the case of a larger automaton. In the sound synthesis domain this is normally translated in more abrupt amplitude fluctuations. In this thesis this fact has been demonstrated with different CA.

7. Presented a critical review of previous work based on criteria of control/prediction and variety of results (Chapter 3 and Chapter 4). The main conclusion of this review is that existing sound synthesis techniques based on CA have limited possibilities for allowing sound design. Regarding the control/prediction criteria, the review showed that in order to develop a sound synthesis technique capable of allowing a sound design process, more research is needed regarding the control of CA. This is because only a few pieces of research, those presented in Section 3.2 (Research Dealing with CA Control), showed concern on the

sound design problem taking into account the control of CA. Of these, only LASy achieved certain level of control/prediction. As for variety of results, four criteria were established in order to estimate the potential of a system for producing a variety of sounds namely, allowing the use of a variety of: CA, mappings, synthesis engines and types of synthesis. It was shown that none of the previous systems involved a variety in the synthesis engines used. This critical review is summarised in Table 1.

8. Devised a number of new CA as a result of these new control mechanisms (for example in Subsection 5.3.3.3.2). Some control mechanisms presented in this thesis are in fact modified versions of an original automaton. Thus, they can be considered as new CA in their own right. For example, the control mechanism developed in Subsection 5.3.3.3.2 modifies the UP parameter of the multitype voter model by means of time-varying colour-dependent UPs whose values are given by a self-regulation algorithm. Thus, the resulting model is a new automaton; the main difference being that the new automaton never achieves consensus, and can evolve endlessly.

6.3 Recommendations for Future Research

1. The exploration of new possibilities of HMS using other multi-state CA. In this sense, HMS has significant potential because CA research is very active and new CA development is ongoing. The search for CA suitable for Histogram Mapping Synthesis is largely a matter of experimentation. It is only after histogram analysis that one can determine the suitability of

the automaton in question. This is because on many occasions the results of a histogram analysis are by no means possible to expect in advance from the CA visual perception. A case in point is the hodge podge machine, where unexpected time-varying structures with “sound-like” features were discovered in the histogram sequences. Expectations for discoveries like this will always be present.

2. The exploration of CA models with much larger grid sizes and larger number of cell states. These could potentially lead to the emergence of more complex behaviours. Also, less speculatively, they may have the interesting application for the additive synthesis of thousands of time-varying sinusoids. For example with the multitype voter model. With this number of sinusoids evolving in time we may achieve highly rich and complex sounds along the lines of orchestral or choral timbres. It would be appropriate to make use of High Performance Computing (HPC) resources for this type of work.

3. Further exploring the textural effects of the dynamic complex beats. The first milestone would be to run the system in real time so that one could experiment with banks of oscillators producing dynamic complex beats. The real-time approach would allow a more dynamic and interactive exploration of the various parameters of the system (i.e. set of beating frequencies, CA parameter settings, etc) when searching for new textures. In addition, the compensated curves used to produce these beats may also be useful for audio processing applications; for example

to improve chorus algorithms. Such compensated and non-deterministic curves could be ideal for producing the required uncorrelated differences in pitch, time delay, amplitude and timbre, between the different sound instances used to compose a chorus mixture.

4. Performing Formant Synthesis with the wide bands of the Quasi-Synchronic behaviour. Some preliminary experiments have been carried out with these wide bands used as is, without any type of manipulation. To improve the results, more experiments would be needed to enable successful modelling of vocal tract or instrument resonances. It would be necessary to take into account the separation of each band, as well as their subsequent location at specific center frequencies with appropriate relative amplitudes and bandwidths (through interpolation).

5. Exploring the use of other CA for controlling the hodge podge machine using the coupling method studied in Subsection 5.4.3.2. In that subsection, the hodge podge was controlled by a growth model with only two states and whose behaviour was highly predictable. Further work exploring more complex and less predictable controllers may be of interest. For example, the multitype voter model can produce more than two dynamic areas, evolving with less predictability than the growth model.

6. Further exploring mutation processes in the multitype voter model. It will be recalled that the main disadvantage of sound textures synthesised with the MVM was their finiteness. This was because the automaton

stopped evolving when consensus was achieved. However, if new species entered the system through genetic mutation, the textures could be endless and more dynamic in terms of their frequency content.

7. Investigate the fixation in the plurality vote rule. This may be preventable through the inclusion of randomness or a nonlinear ingredient in the dynamics (as suggested by David Griffeath in a personal communication). Successful results along these lines would provide means to further investigate control mechanisms for sound design.

8. Investigating HMS in relation to non-CA computational models. Motivation for this includes studies in the area of Grain Growth using Potts models. These lead to computer simulations that may be suitable for HMS sound synthesis because they are similar to the evolution of the plurality vote rule.

Appendix A: Sound Examples

This Appendix presents sound examples rendered from histogram sequences of the CA studied in this thesis. Note that these examples represent only a few of the possibilities of HMS, since virtually infinite sounds can be rendered from a single histogram sequence by developing other offline sound design processes different than the ones explained below.

Also, different sound engines can endow the final sounds with a different character. In the below examples, the structural elements of the histogram sequences have been used to control two types of synthesis engine: sinusoidal additive synthesis, with or without angle modulations, and subtractive synthesis by means of FFT convolution filtering. The additive synthesizer has been implemented in Matlab. The FFT convolution filter has been implemented in Max/MSP.

The sound files can be found in the accompanying CD. Below, explanatory notes will be given for each sound example. “Harmonic” and “non-harmonic” are terms used here to distinguish two ways of producing sound with sinusoidal additive synthesis. The former refers to the assignment of frequencies to the partials following the harmonic series, and the latter refers to the assignment of random frequencies to the partials. These random values are bounded between the lowest and the highest frequency values selected by the user.

Sound Example 1: This example is an exception in the sense that is not generated using HMS. It consists of a piano tone, followed by the sound of its

residual after removing its sinusoidal components. The sounds correspond respectively to the spectrograms of Figure 3 and Figure 4.

Sound Example 2: A Harmonic sound rendered from the histogram sequence in Figure 33. This example also serves to illustrate the possibilities of an offline sound design process, that is, in Manual Mode: in order to create releases, the histogram sequence has been duplicated, flipped and joined to the end of the original one, creating a symmetrical structure.

Sound Example 3: This example is composed of six sounds, all rendered from the same histogram sequence in Figure 34. The first sound is harmonic, with only ten partials synthesised. The second sound has all twenty partials synthesised. The third and the fourth sounds are to be compared: the latter is synthesised after an offline sound design process attenuating the even partials, resulting in a more hollow timbre. The last two sounds are non-harmonic, with different frequencies and durations.

Sound Example 4: This example is rendered from the histogram sequence in Figure 35 and it should be compared to the first sound of Sound Example 3. The difference between them lies in the type of releases, as explained in the thesis.

Sound Example 5: This example illustrates the possibility of creating similar but not identical sounds. It is rendered from histogram sequences similar to those in Figure 38. To facilitate the perception of this nuance, the melody begins with a sequence of groups of three notes with the same pitch. The purpose of this

example is to illustrate this similarity feature, though it is admitted that, perhaps, the timbre of the designed instrument is not particularly interesting. It was initially designed in Manual Mode by establishing the relationship between the frequencies and the relative energies that the histogram bins will contribute to the sound, in a rather simple way: harmonic frequencies were assigned in ascending order to the histogram bins, which were ordered from largest contribution to smallest contribution. Following this sound design process in Manual Mode, all the sounds of the melody were then synthesised in Automatic Mode. The attacks and releases were created by applying external envelopes.

Sound Example 6: A non-harmonic sound rendered from a histogram sequence similar to those in Figure 38, i.e. with a greater timbral complexity in the attack than in the rest of the sound. The attack and the release were created by applying external envelopes.

Sound Example 7

and

Sound Example 8: These are two sets of sounds which illustrate, through comparison between the two sets, the possibility of producing similar but not identical damped sounds. The first set of sounds, Sound Example 7, consists of the exact same sound played six times. Then the second set of sounds, Sound Example 8, consists of 8 similar but not identical instances of the same sound. The sounds are harmonic and are an imitation of plucked string tones. Note that HMS can produce more notable differences between the similar but not identical sounds. However, Sound Example 8 is deliberately synthesised to create very similar sounds with very small differences. This is in an attempt to

simulate acoustic instruments which, on the one hand, never play the same sound twice, but on the other hand, have performers who are trained to play homogeneous timbres (i.e. minimizing timbral differences between the notes of a melody).

Sound Example 9: A set of three non-harmonic damped sounds designed to imitate bell tones. The sounds are rendered from the structure in Figure 43. The random frequencies of the partials are sorted in ascending order before being assigned consecutively to the partials of the histogram sequence. As a result, the higher partials are more short-lived than the lower partials.

Sound Example 10: This example presents four sounds obtained from the histogram sequence in Figure 45. The first two sounds are harmonic, with no external envelope being applied. The second sound is designed to be longer than the first, and therefore the attack (which will also be longer) is perhaps more interesting. External envelopes have been applied to the third sound, resulting in a violin-like tone. The fourth sound is non-harmonic.

Sound Example 11: This example consists of five sounds rendered from the histogram sequence in Figure 46. The first four sounds are non-harmonic. The last sound is harmonic. The frequencies are sorted in ascending order before being assigned consecutively to the partials of the histogram sequence. These partials, in turn, have been sorted according to the energy they will contribute to the sound. In this way, the higher partials are more short-lived than the lower partials. For all sounds, external amplitude envelopes have been applied for creating attacks and releases.

Sound Example 12: A non-harmonic texture rendered from the histogram sequence in Figure 47. Note that not all of the twenty partials were included in the synthesis. Two of them, the second and the last ones, were excluded in the sound design process because the random frequencies they were assigned were unpleasantly prominent (note that the last partial is the most prominent one).

Sound Example 13: A non-harmonic texture rendered from the structure shown in Figure 48 –a manipulated histogram sequence (as explained in the thesis). A long fade-in (up to the midpoint of the sound) and a long fade-out (from the midpoint to the end) were applied externally. Also, the random frequencies were sorted in ascending order before being consecutively assigned to the histogram bins.

Sound Example 14

and

Sound Example 15: Here we apply the textural concept of dynamic complex beats to an existing sound by means of audio processing. Sound Example 14 is the original sound of an electric guitar and Sound Example 15 is the processed version.

Sound Example 16: A harmonic sound rendered with a number of narrow bands from the NB2 zone of a Quasi-Synchronic behaviour.

Sound Example 17: This example consists of three sounds rendered from the non-sustained structures in Figure 61 (which are obtained from the Fast-Infection behaviour). The first and the second sounds are rendered from the top figure, consisting of narrow bands. The first sound is non-harmonic and the second is harmonic. The third sound is rendered from the bottom figure, which can be considered as a sequence of spectral envelopes. This sound has been synthesised using subtractive synthesis of white noise by means of FFT Convolution filtering.

Sound Example 18: A non-harmonic sound rendered from a Spiral behaviour histogram sequence. Only the most prominent narrow bands have been included in the synthesis, which have been considered as sinusoidal components.

Sound Example 19: This example illustrates two ways of designing release patterns for the narrow bands of the Quasi-Synchronic behaviour (as shown in Figure 65). The first sound has release times dependent on the amplitude of each partial. The second sound has release times dependent on the frequency (or histogram bin).

Sound Example 20: Two sounds rendered from the structure in Figure 73. The first is harmonic and the second is non-harmonic. These sounds present greater spectral complexity in the attack portion than in the rest of the sound. Before rendering the sounds, the following offline sound design process was carried out. The bins of the histogram were sorted according to the energy that each partial would contribute to the sound. Then, the third bin of the newly sorted

structure was doubled in amplitude. The frequencies were assigned in ascending order, so that the higher partials are the short-lived ones. This is an example where a small angle modulation has been applied to the sinusoidal components. The modulation is driven by the amplitude trajectories so that the amplitude fluctuations are translated into frequency fluctuations.

Sound Example 21: A non-harmonic sound rendered from the histogram sequence in Figure 43, and in which an angle modulation has been applied to the sinusoidal components. The modulation is driven by the amplitude trajectories of the histogram bins, so that these are translated into frequency trajectories in the form of glissandi.

Sound Example 22: Three different harmonic sounds are created as a result of three different offline sound design processes being applied to a plurality vote histogram sequence, similar to that seen in Figure 77. The sounds are an imitation of plucked strings. External envelopes were applied. The difference between the three sounds is a result of the number of partials that have been rendered for each sound: 17, 20 and 50 partials respectively for each sound.

Sound Example 23: Two sounds, one harmonic and one non-harmonic, rendered from the histogram sequence in Figure 78. Before rendering the sounds, the histogram sequence was sorted according to the energy that each partial would contribute to the sound. Then, the frequencies were sorted in ascending order before being assigned consecutively to the partials of the histogram sequence. As a result, the higher partials are more short-lived than

the lower partials. It should be noted that the partial corresponding to the colour that covered the whole automaton grid was discarded in the sound synthesis.

Appendix B: Video Examples

Video Example 1: This example shows the first generations of an MVM evolution. Starting with a uniform random distribution of cell values, or colours, as the initial configuration, this automaton self-organizes into clusters, or areas of a single colour.

Video Example 2: An MVM evolution achieving consensus.

Video Example 3: A CA evolution for creating entry delays for the partials, as seen in Figure 33.

Video Example 4: Self-explanatory video where similar but not identical sounds are generated for a video game. In addition to being on the accompanying CD, this video is available at:
http://www.youtube.com/watch?v=R6pKArUaTxM

Video Example 5: Self-explanatory video with examples of dynamic complex beats. Note that the final texture in the example is the same as Sound Example 13 (which constitutes an example of dynamic complex beats itself). In addition to being on the accompanying CD, this video is available at:
http://www.youtube.com/watch?v=XADlXSP5TkQ

Video Example 6: This video shows the Quasi-Synchronic behaviour of the HPM.

Video Example 7: This video shows the Fast-Infection behaviour of the HPM.

Video Example 8: This video shows the Spiral behaviour of the HPM.

Video Example 9: Evolving two simultaneous behaviours in an HPM. The left half of the automaton has the rule parameters adjusted to produce Quasi-Synchronic behaviour whereas the right half has the rule parameters adjusted to produce Spiral behaviour.

Video Example 10: Evolving two simultaneous Quasi-Synchronic behaviours in an HPM. Initially, the left and the right halves are visually differentiated, but over time, the two halves are visually indistinguishable.

Appendix C: Record of Activities

Publications

Serquera, J. and Miranda, E. R. (2008). "Cellular Automata Sound Synthesis: From Histograms to Spectrograms". In A. Adamatzky et al. (Eds.), Automata 2008: Theory and Applications of Cellular Automata. Luniver Press. Frome (UK).

Serquera, J. and Miranda, E. R. (2008). "Spectral Synthesis and Control with Cellular Automata". Proceedings of the International Computer Music Conference. Belfast (UK).

Serquera, J. and Miranda, E. R. (2010). "Algorithmic Sound Composition using Coupled Cellular Automata". Proceedings of the CMSIM International Conference. Chania (Greece).

Serquera, J. and Miranda, E. R. (2010). "Cellular Automata Sound Synthesis with an Extended Version of the Multitype Voter Model". Proceedings of the AES 128th Convention. London.

Serquera, J. and Miranda, E. R. (2010). "Evolutionary Sound Synthesis: Rendering Spectrograms from Cellular Automata Histograms". Proceedings of EvoWorkshops / EvoMUSART, Lecture Notes in Computer Science. Springer-Verlag. Berlin.

Serquera, J. and Miranda, E. R. (2011). "Cellular Automata Dynamic Control for Sound Design with Histogram Mapping Synthesis and the Multitype Voter Model". Proceedings of the International Computer Music Conference. Huddersfield (UK).

Presentations

"Non Linear Dynamical Systems from Micro to Macro Composition". Postgraduate research seminar, Plymouth University. (May 2006)

"Real Time considerations of LISP based systems for Algorithmic Composition". Postgraduate research seminar, Plymouth University. (November 2006)

"Cellular automata sound synthesis: From histograms to spectrograms". Oral paper presentation in Automata 2008 - EPSRC Workshop Cellular Automata Theory and Applications, Bristol, UK. (June 2008)

"Spectral synthesis and control with cellular automata". Poster paper presentation in the International Computer Music Conference - ICMC 2008, Belfast, UK. (September 2008)

“Cellula Automata Sound Synthesis with the Multitype Voter Model”. Postgraduate research seminar, Plymouth University. (March 2010)

"Evolutionary Sound Synthesis: Rendering Spectrograms from Cellular Automata Histograms" Oral paper presentation in EvoMUSART, Istambul (April 2010)

"Cellular Automata Sound Synthesis with an Extended Version of the Multitype Voter Model". Poster paper presentation in AES 128th Convention. London (May 2010)

"Algorithmic Sound Composition using Coupled Cellular Automata". Virtual oral paper presentation in the CMSIM International Conference. Chania, Greece (June 2010)

"Cellular Automata Dynamic Control for Sound Design with Histogram Mapping Synthesis and the Multitype Voter Model". Poster paper presentation in the International Computer Music Conference. Huddersfield, UK (August 2011)

“Cellular Automata Sound Synthesis”. Postgraduate research seminar, Plymouth University. (December 2011)

References

Adamatzky, A. (1994). *Identification of Cellular Automata*. London, Taylor & Francis.

Adamatzky, A. (2001). *Computing in Nonlinear Media and Automata Collectives*. Bristol and Philadelphia, Institute of Physics Publishing.

Aldoshina, I., Pychkov, S., Matcievski, I., Nicanorov, A., Tovstik, P. and Cherniav, S. (2003). "The analysis of peculiarities of Russian bells acoustic parameters". Proceedings of the 114th AES Convention. Amsterdam.

Ariza, C. (2007). "Automata Bending: Applications of Dynamic Mutation and Dynamic Rules in Modular One-Dimensional Cellular Automata." Computer Music Journal 31(1): 29-49.

Bessell, D. (2011). "Ophidian and the Uncanny Valley". Proceedings of International Computer Music Conference. Huddersfield (UK).

Beyls, P. (2004). "Cellular Automata Mapping Procedures". Proceedings of the International Computer Music Conference. Miami.

Boker, S. M. (2007). "Specifying Latent Differential Equations Models". In Steven M. Boker and Michael J. Wenger (Eds.), Data analytic techniques for dynamical systems. New Jersey, Lawrence Erlbaum Associates.

Bowcott, P. (1989). "Cellular Automation as a means of high level compositional control of granular synthesis". Proceedings of the International Computer Music Conference. San Francisco. 55-57.

Burraston, D. and Edmonds, E. (2005). "Cellular Automata in Generative Electronic Music and Sonic Art : Historical and Technical Review." Digital Creativity 16(3): 165-185.

Camazine, S., Deneubourg, J.-L., Franks, N. R., Sneyd, J., Theraulaz, G. and Bonabeau, E. (2003). *Self-organization in biological systems*. Woodstock, Oxfordshire, Princeton University Press.

Chareyron, J. (1990). "Digital Synthesis of Self-Modifying Waveforms by Means of Linear Automata." Computer Music Journal 14(4): 25-41.

Chinneck, J. W. (1988). "How to Organize your Thesis". Available online at: http://www.sce.carleton.ca/faculty/chinneck/thesis.html (Accessed: 25/10/2011).

Clifford, P. and Sudbury, A. (1973). "A Model for Spatial Conflict." Biometrika 60(3): 581-588.

Comajuncosas, J. M. (1998). "Implementing Cellular Automata in Csound, with Some Musical Applications". Available online at: www.csounds.com/jmc/Articles/cellular/synthesis.html (Accessed: 29/10/2010).

Cook, P. R. (2002). *Real sound synthesis for interactive applications*. USA, A. K. Peters, Ltd.

Cox, J. T. and Durrett, R. (2002). "The Stepping Stone Model: New Formulas Expose Old Myths." The Annals of Applied Probability 12(4): 1348-1377.

Fisch, R. and Griffeath, D. (1996). "WinCA: A Cellular Automaton Modeling Environment (version 1.0)". Available at http://psoup.math.wisc.edu/ftp/pub/winca_10.exe from www.archive.org.

Ganguly, N., Sikdar, B. K., Deutsch, A., Canright, G. and Chaudhuri, P. P. (2003). "A Survey on Cellular Automata". Technical Report, Center for High Performance Computing. Dresden University of Technology.

Gardner, M. (1970). "Mathematical Games: The fantastic combinations of John Conway's new solitaire game "Life"." Scientific American 223: 120–123.

Gaylord, R. J. and Nishidate, K. (1996). *Modeling Nature: Cellular Automata Simulations with Mathematica*. New York, Springer-Verlag.

Gerhardt, M. and Schuster, H. (1989). "A cellular automaton describing the formation of spatially ordered structures in chemical systems." Physica D 36(3): 209–221.

Gonzalez, R. C. and Woods, R. E. (1993). *Digital image processing*. Reading, Massachusetts, Addison-Wesley.

Griffeath, D. (1994). "The Cook Book: Soap Bubble clustering of a Plurality Vote Rule". Available online at: http://psoup.math.wisc.edu/archive/recipe8.html (Accessed: 12/10/2011).

Hibbert, B. (2000). "The Sound of Bells - Variation of intensity over time". Available online at: www.hibberts.co.uk/time.htm (Accessed: 4/3/2011).

Holley, R. A. and Liggett, T. M. (1975). "Ergodic Theorems for Weakly Interacting Infinite Systems and the Voter Model." The Annals of Probability 3(4): 643-663.

Hunt, A., Kirk, R. and Orton, R. (1991). "Musical Applications of a Cellular Automata Workstation". Proceedings of the International Computer Music Conference. Montreal. 165-168.

Katrami, A. I., Kirk, R. and Myatt, A. (1991). "A Phase Vocoder Graphical Interface for Timbral Manipulation of Cellular Automata and Fractal Landscape Mappings". Proceedings of the International Computer Music Conference. Montreal. 106-109.

Kimura, M. (1953). "Stepping-stone Model of Population." Annual Report of the National Institute of Genetics 3: 62-63.

Kirke, A. and Miranda, E. R. (2007). "Evaluating Mappings for Cellular Automata Music". Proceedings of ECAL 2007 Workshop on Music and Artificial Life (MusicAL 2007). Lisbon.

Kreger, T. (1997). "Cellular Automata in the Spectral Domain". Proceedings of Interface 1997. University of Auckland, Auckland.

Levitin, A. (2007). *Introduction to the Design and Analysis of Algorithms*. USA, Pearson / Addison Wesley.

Liggett, T. M. (1985). *Interacting Particle Systems*. New York, Springer-Verlag.

McAulay, R. J. and Quatieri, T. F. (1986). "Speech analysis/synthesis based on a sinusoidal representation." IEEE Transactions on Acoustics, Speech, and Signal Processing ASSP-34(4): 744-754.

McLaughlin, S. and Tremblay, P. A. (2010). "SpectralConway: Cellular Automata Off The Grid". Proceedings of the International Computer Music Conference. New York. 68-71.

Miranda, E. R. (1995). "Granular Synthesis of Sounds by Means of a Cellular Automaton." Leonardo 28(4): 297-300.

Miranda, E. R. (2001). "Evolving Cellular Automata Music: From Sound Synthesis to Composition". Proceedings of the Workshop on Artificial Life Models for Musical Applications - ECAL 2001. Prague.

Miranda, E. R. (2002). *Computer sound design: synthesis techniques and programming*. Oxford, Focal Press.

Miranda, E. R. (2005). "Artificial Phonology: On Synthesising Disembodied Humanoid Voice for Composing Music with Surreal Languages." Leonardo Music Journal 15: 8-16.

Miranda, E. R., Correa, J. and Wright, J. (2000). "Categorising complex dynamic sounds." Organised Sound 5(2): 95-102.

Morimoto, Y. (2010). "Hacking Cellular Automata: An Approach to Sound Synthesis". Proceedings of the SuperCollider Symposium. Berlin.

Neumann, J. v. and Burks, A. W. (1966). *Theory of Self-Reproducing Automata*. Champaign, IL, University of Illinois Press.

Ostrov, D. and Rucker, R. (1996). "Continuous-valued Cellular Automata for Nonlinear Wave Equations." Complex Systems 10: 91-119.

Pratt, W. K. (2007). *Digital image processing*. Hoboken, New Jersey, Wiley.

Resnick, M. (1994). *Turtles, termites and traffic jams*. Massachusetts, MIT Press.

Risset, J.-C. (1985). "Computer Music Experiments 1964 - ..." Computer Music Journal 9(1): 11-18.

Roads, C. (1996). *The Computer Music Tutorial*. Cambridge, MA, MIT Press.

Rocha-Iturbide, M. (1999). "Les techniques granulaires dans la synthèse sonore". Doctoral thesis. University of Paris VIII, Paris.

Rucker, R. (2010). "Artificial Life Lab (online version)". Available online at: http://www.rudyrucker.com/pdf/artificiallifelab.pdf (Accessed: 14/10/2011).

Sarkar, P. (2000). "A brief history of cellular automata." ACM Computing Surveys 32(1): 80-107.

Serquera, J. and Miranda, E. R. (2008). "Cellular Automata Sound Synthesis: From Histograms to Spectrograms". In A. Adamatzky et al. (Eds.), Automata 2008: Theory and Applications of Cellular Automata. Frome (UK), Luniver Press.

Serquera, J. and Miranda, E. R. (2008). "Spectral Synthesis and Control with Cellular Automata". Proceedings of the International Computer Music Conference. Belfast (UK).

Serquera, J. and Miranda, E. R. (2010). "Algorithmic Sound Composition using Coupled Cellular Automata". Proceedings of the CMSIM International Conference. Chania (Greece).

Serquera, J. and Miranda, E. R. (2010). "Cellular Automata Sound Synthesis with an Extended Version of the Multitype Voter Model". Proceedings of the AES 128th Convention. London.

Serquera, J. and Miranda, E. R. (2010). "Evolutionary Sound Synthesis: Rendering Spectrograms from Cellular Automata Histograms". Proceedings of EvoWorkshops / EvoMUSART, Lecture Notes in Computer Science. Springer-Verlag. Berlin.

Serquera, J. and Miranda, E. R. (2011). "Cellular Automata Dynamic Control for Sound Design with Histogram Mapping Synthesis and the Multitype Voter Model". Proceedings of the International Computer Music Conference. Huddersfield (UK).

Serra, X. (1997). "Musical Sound Modeling with Sinusoids plus Noise". In C. Roads, S. Pope, A. Picialli, G. De Poli, editors. Musical Signal Processing. Swets & Zeitlinger Publishers.

Serra, X. and Smith, J. (1990). "Spectral Modeling Synthesis: A Sound Analysis/Synthesis System based on a Deterministic plus Stochastic Decomposition." Computer Music Journal 14(4): 12-24.

Smith, J. O. (2011). *Spectral Audio Signal Processing, 2011 Draft*, http://ccrma.stanford.edu/~jos/sasp/, online book, Accessed: 8/12/2011.

Smith, J. O. and Serra, X. (1987). "PARSHL: An Analysis/Synthesis Program for Non-Harmonic Sounds based on a Sinusoidal Representation". Proceedings of the International Computer Music Conference. San Francisco.

Thornburg, H. (2005). "Detection and modeling of transient audio signals with prior information". Doctoral thesis. Stanford University, Stanford.

Toguchi, S., Akamine, Y. and Satoshi, E. (2008). "Research into the Generation of Sound Effects Using a Cellular Automaton". Proceedings of the 8th international conference on Cellular Automata for Reseach and Industry. Yokohama (Japan). Springer-Verlag.

Tovar, A., Patel, N., Kaushik, A. K., Letona, G. A. and Renaud, J. E. (2004). "Hybrid Cellular Automata: a biologically-inspired structural optimization technique". 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference. Albany, New York.

Truax, B. (1988). "Real-Time Granular Synthesis with a Digital Signal Processor." Computer Music Journal 12(2): 14-26.

Ulam, S. M. (1960). *A Collection of Mathematical Problems*. New York, Interscience.

Vaidhyanathan, S., Minai, A. and Helmuth, M. (1999). "ca : A System for Granular Processing of sound using Cellular Automata". Proceedings of the 2nd COST G-6 Workshop on Digital Audio (DAFx99). Trondheim.

Verma, T. and Meng, T. H. Y. (2000). "Extending spectral modeling synthesis with transient modeling synthesis." Computer Music Journal 24(2): 47-59.

Wójtowicz, M. (2001). "Cellular Automata rules lexicon". Available online at: http://psoup.math.wisc.edu/mcell/rullex_udll.html (Accessed: 14/10/2011).

Wolf, T. D. and Holvoet, T. (2005). "Emergence Versus Self-Organisation: Different Concepts but Promising When Combined." Lecture Notes in Computer Science 3464: 1-15.

Xenakis, I. (1992). *Formalized music: thought and mathematics in composition*. Stuyvesant, NY, Pendragon Press.

Zolzer, U. (2002). *DAFX - Digital Audio Effects*. West Sussex, England, John Wiley & Sons.