Cellular automata sound synthesis: From histograms to spectrograms

Jaime Serquera and Eduardo R. Miranda

Interdisciplinary Centre for Computer Music Research (ICCMR),
University of Plymouth, Drake Circus, Plymouth, Devon PL4 8AA, UK
{jaime.serquera,eduardo.miranda}@plymouth.ac.uk

Abstract. We are investigating ways in which advanced unconventional computation methods may provide novel approaches to music technology. In this paper we report on a new technique to synthesise sounds from Cellular Automata (CA) models of reaction-diffusion chemical computers. We render the behaviour of CA into sounds using statistical analysis of CA cells as their values change in time. We estimate the probability distributions of the cells values for each cycle of the automaton by histogram measurements of the images produced by plotting the values of the CA cells as a matrix of coloured pixels. We have studied the behaviour of a CA model proposed by Gerhard and Schuster under a variety of different settings in order to gain a better understanding of the histograms, with a view on predicting the types of sounds that different CA behaviours would produce.

1 Introduction

The field of Computer Music is as old as Computer Science. Computers have been programmed to play music as early as the early 1950's when Geoff Hill programmed the CSIR Mk1 computer in Australia to play popular musical melodies [1]. Nowadays, the computer is becoming increasingly ubiquitous in all aspects of music. Uses of computer technology in music range from systems for musical composition to systems for distribution of music on the Internet. The implementation of such applications often demands the skillful combination of software engineering and artistic creativity. Whereas most current research into computers and music focuses on the development of media technology for delivering music to consumers (e.g., MP3 format, Internet search engines, and so on) our research focuses on the development of technology for musical creativity; that is, technology to aid musicians to create content for the media. We are particularly interested in investigating ways in which unconventional computation methods may provide novel approaches to music technology, particularly for the design of new musical instruments. In this paper we report on a new technique to synthesise sounds from Cellular Automata (CA) models of reaction-diffusion chemical computers [2].

In this paper, the CA model used to illustrate our new synthesis technique is based on an automaton proposed by Gerhard and Schuster [3], where states of a cell are 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 transition rules are expressed as follows:

- RULE 1: IF $m_{x,y}[t] = 0$ THEN $m_{x,y}[t+1] = int(\frac{A}{r1}) + int(\frac{B}{r2})$
- RULE 2: IF $0 < m_{x,y}[t] < V-1$ THEN $m_{x,y}[t+1] = int(\frac{S}{A}) + K$
- RULE 3: IF $m_{x,y}[t] = V-1$ THEN $m_{x,y}[t+1] = 0$

where the value 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 CA grid. A and B represent the number of infected and ill cells in the neighbourhood, respectively; r_1 and r_2 are constants (which can be set to different values); S stands for the sum of the values of all cells in the neighbourhood; and V is the number of possible values that a cell can adopt. A desirable property of this CA model is its cyclic nature, which allows us to work with different ranges of cell values (also referred to in this paper as *colours*).

2 Rendering spectrograms from cellular automata histograms

In a nutshell, our method for rendering sounds from CA involves a mapping from CA histograms onto sound spectrograms. We devised a method to render the behaviour of CA into sounds using statistical analysis of the CA cells as their values change in time. We estimate the probability distributions of the cell values for each cycle of the automaton by histogram measurements [4] of the images generated by plotting the values of the CA cells as a matrix of coloured pixels. The histogram of a digital image with levels of gray colour in the range $[0, L-1]$ is a discrete function, where r_k is the k^{th} gray level, n_k is the number of pixels in the image with that gray level, n is the total number of pixels in the image, and $k = 0, 1, 2, \ldots, L-1$. Loosely speaking, $p(r_k)$ gives an estimate of the probability of occurrence of gray-level r_k [4], where $\sum p(r_k) = 1$.

We have found zones in the histograms consisting of narrow bands (sometimes with a width of just one colour) clearly separated from each other. This is of interest to us because the automaton self-organizes through very specific sets of predominant cell values, or colours. These sets of colours vary depending on the settings of the of the transition rules parameters. By examining the evolution of these narrow bands, we surprisingly found that their envelopes (i.e., time trajectories) resemble the amplitude envelopes of the partials in the spectrum of sounds. With this in mind, we devised a method to generate sound spectrograms from CA histograms. Considering that the time domain is common for both the histograms sequences and the spectrogram, we map the histogram's sample space domain onto the spectrogram's frequency domain and, the histogram's probability domain onto the spectral magnitude domain.

The spectrograms are synthesised using a combination of additive synthesis and Frequency Modulation (FM) techniques [5]. Firstly, we select and extract the structures of the histograms sequence corresponding to the most predominant bins. Then, these structures are converted into partials of a spectrogram. We extract the amplitude envelopes and, without any kind of amplitude transformations, apply interpolation in order to define the duration of the sound (60 time steps per second). The frequency assignment for the partials is arbitrary; in our experiments we have assigned random frequencies to each partial in order to obtain complex ratios between them. By designing a bounded random frequency generator it is possible to obtain a wide range of different sounds. In order to render the partials we consider narrow bands of just one *colour width*. Then, we add frequency fluctuations by using FM. The FM is driven by the amplitude envelopes of each partial as follows: $f_p(t) = f_p + width * AmpEnv_p(t)$ where p denotes each of the partials.

3 Making sense of the CA behaviour

We have studied the behaviour of the automaton under a variety of different settings in order to gain a better understanding of the histograms, with a view on predicting the types of sounds that different CA behaviours would produce.

3.1 Quasi-synchronic behaviour

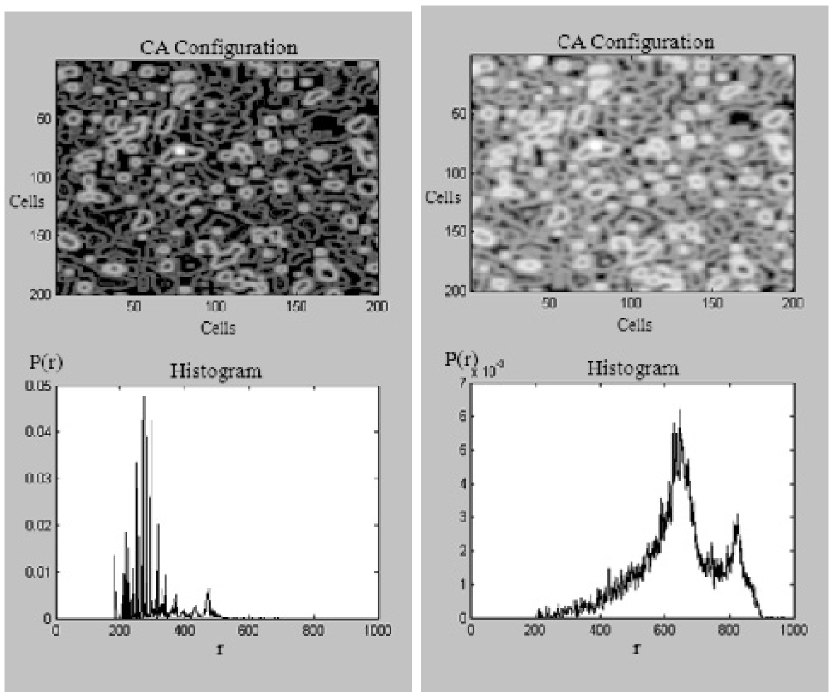

(a) (b)

Fig. 1. Different CA configurations in the 'quasi-synchronic' behavior and the histograms they produce: narrow bands (a) and wide bands (b).

We have discovered that quasi-synchronic CA behaviour (i.e., where all cells of the CA grid reach their maximum allowed value almost simultaneously) generates histograms that are very suitable for sound synthesis. Just after the cells reach their maximum value, patterns of *distorted circumferences* emerge. The contours of these patterns create narrow bands, or peaks, in the histogram. From here on, the cells values increase towards the maximum value and the boundaries of the distorted circumferences become less defined, creating wide bands in the histogram (Fig. 1). At each cycle of the automaton, this process is repeated but with slightly different distorted circumferences shapes, creating structures in the histogram with time varying amplitudes.

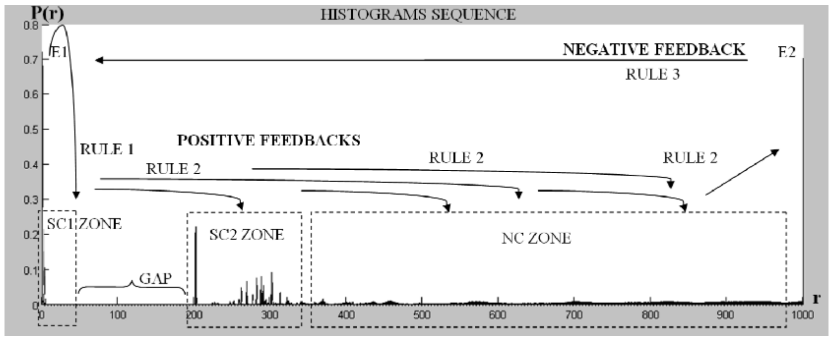

Fig. 2. Frontal plot of the histograms sequence of a CA evolution and the role of the CA rules.

Figure 2 shows the frontal plot of a typical histograms sequence produced by a quasi-synchronic CA run. It also shows how the CA rules and parameter values are reflected in the histogram. We can see the two feedbacks, one positive (due to Rules 1 and 2), and one negative (due to Rule 3). Peaks E1 and E2 are the two CA end points corresponding to the minimum and maximum cell values. The probabilities of these values (or colours) are always high. This is so for E1, due to Rule 3, and for E2 it is due to the need to impose a limit on (or *cap*) the maximum cell value (e.g., when the result of Rule 2 is a value that is over the maximum allowed value). SC1 is a zone of narrow bands that appear due to Rule 1. Next to SC1 there is a GAP. This is a region that always has zero probability values, meaning that values corresponding to this area are ever reached by the automaton. The beginning of the GAP is due to Rule 1, which has a maximum possible value depending of r_1 and r_2; this limits the size of the SC1 Zone. The GAP and its size are due to the constant K in Rule 2, which is like an offset. SC2 is a zone of narrow bands due to Rule 2 applied to cells with values (or colours) in the range of SC1 and SC2 itself. Finally, in NC Zone there are wide bands that appear due to Rule 2.

The behaviour described above is granted by a variety of different combinations of CA parameter settings. We found interesting time evolving structures by working with large ranges of values, from hundreds and usually thousands. The images produced by the automaton with thousands of different values (or colours) proved to be more suitable for our purposes because they seem to be much more lively and natural than with just a few values. Typical parameters settings (for someone wishing to replicate our experiments) are: CA grid of 200 by 200 cells, K equal to a value between 20% and 30% of V, and both r_1 and r_2 set equal to 2.

3.2 Long-term behaviour

The effect of the automaton's long-term behaviour can be of two types: long-term behaviours that produce structures for sustained sounds and long-term behaviours that produce structures for non-sustained sounds. The previously mentioned quasi-synchronic behaviour is an example of the former type. CA definitions with fewer neighbours are likely to generate structures suitable for non-sustained sounds. This is the case when considering a Neumann neighbourhood [3] or even fewer neighbours; e.g., the case where the central cell is not a neighbour. For instance, if we consider fewer neighbours, the divisor of Rule 2 is lower than if we consider Moore neighbourhoods, and therefore infected cells will have greater chances of getting immediately ill. Ill cells in the neighbourhood cause a higher numerator while the denominator does not grow due to ill or healthy cells. This causes the automaton to reach a long-term behaviour with cells whose value oscillates between 0 (corresponding to SC1 Zone) and $V-1$. While reaching this long-term behaviour, the rest of the histogram's bins fade out to zero, creating the effect of sound release. These structures are interesting because different release times emerge for different histogram bins (Figs. 3 and 4). This kind of effect occurs in the sounds produced by most acoustic instruments.

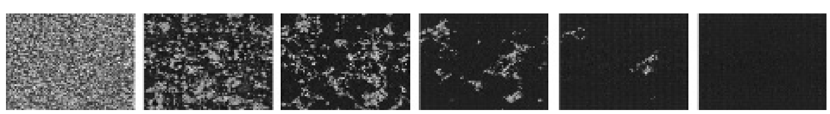

Fig. 3. CA behaviour that produces non-sustained sounds.

3.3 Spiral waves

A type of behaviour referred to as spiral waves create types of histograms that are of great interest for sound synthesis. Starting from an initial random configuration of cell values, the automaton often produces behaviour that resembles the quasi-synchronic behaviour mentioned earlier.

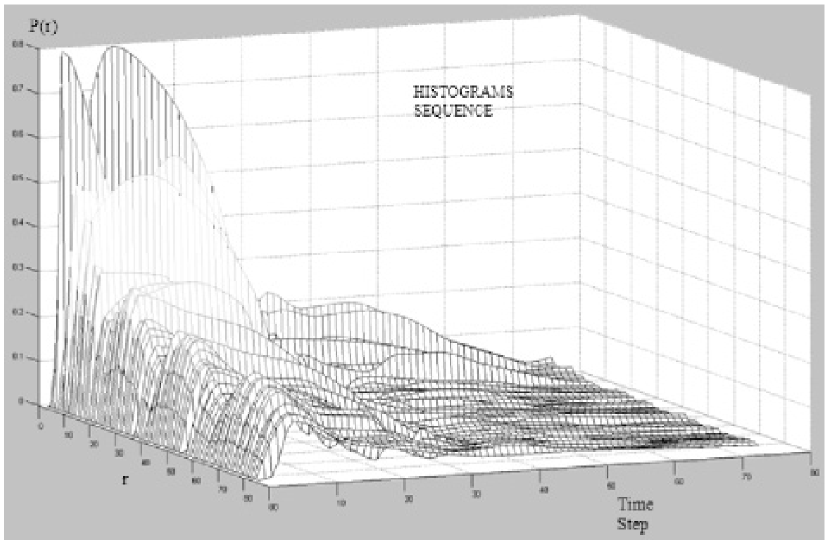

Fig. 4. Histograms sequence from a non-sustained structure obtained with a method for extracting smooth amplitude envelopes.

But then, spiral waves start to develop, creating structures that resemble sound partials with increasingly high amplitudes. Once the spirals have completely evolved, they often expand themselves, covering the whole CA grid. This creates sequences of images, which are cyclic and stable. When rendered to sound, the amplitudes of the partials often stop increasing and settle to relatively stable values. What is interesting here is that the relative amplitudes of the partials are very similar to each other but with small differences in their time evolution (or envelopes) (Fig. 5). This is a general property found in most sounds produced by acoustic instruments. For instance, a sound played on a violin starts noisy due to the attack of the bow and then it settles into a periodic vibration. (When the violin is played by experienced violinists, we often cannot hear this noisy attack because the strings settle into vibration very quickly.)

3.4 Time dynamics

Different CA behaviours can produce different types of spectral structures with different time evolutions. Time varying amplitudes can be considered in various ways. The original amplitude evolution of the histogram's bins usually display oscillatory behaviour, and it could be de-

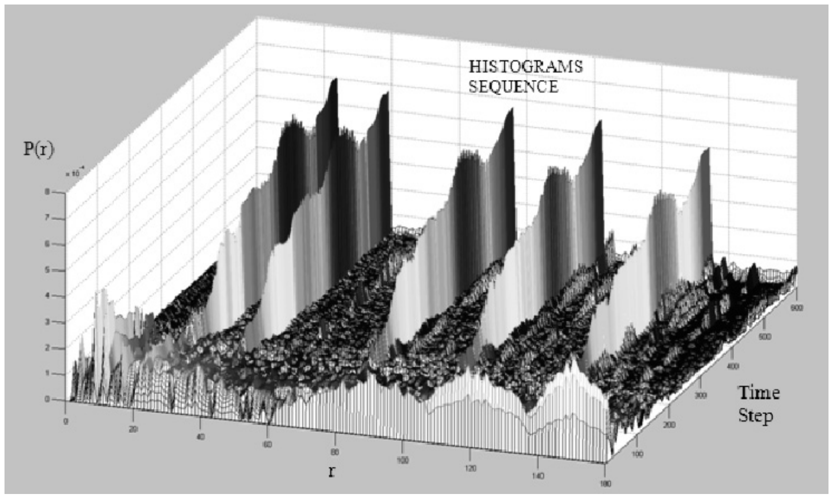

Fig. 5. Histograms sequence from spiral waves.

sired to perform an amplitude envelope extraction algorithm to get a smoother amplitude envelope (as the one shown in Fig. 4). Figure 6 shows the time evolution of one bin of a histogram corresponding to a quasi-synchronic behaviour (solid line). Note the prominent peaks due to the quasi-synchronic behaviour. The distance of the peaks is the period of the CA cycle. The dotted line indicates the amplitude envelope that could be extracted from this histogram.

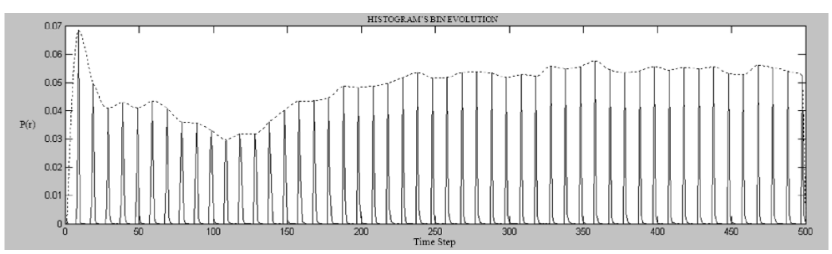

Fig. 6. Time evolution of one histogram's bin in the quasi-synchronic behavior (solid line), and its respective amplitude envelope (dotted line).

The self-organizing process of the automaton also produces interesting noisy structures. This is important for our research because in addition to the sinusoidal partials, noise is an important component of sounds, particularly at their onset (or *attack*) [5]. In Figs. 7 and 8 we can see a noisy structure at the beginning of the histograms sequence, which then disappears while more stable structures emerge.

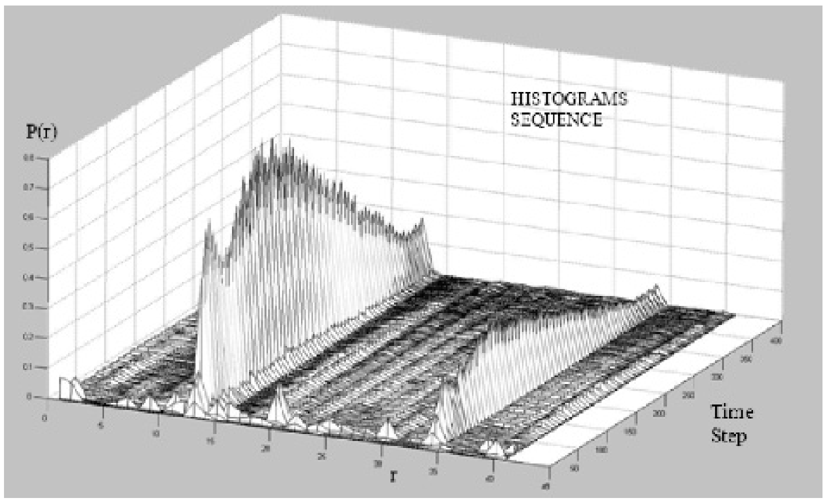

Fig. 7. Structures from SC2 Zone showing self-organization and correlated amplitudes.

When histogram structures have narrow bands, the respective spectra will contain deviations of partials frequency trajectories, which is another property commonly found in sounds produced by acoustic instruments (Fig. 8).

3.5 Invariance property

We have found an invariance property in the histograms sequences by studying different runs of the automaton (with enough time to reach the long-term behaviour) with the same settings, but starting from different initial uniform random configurations of cells values. By looking at a zone of narrow bands (peaks of prominent colours like SC1or SC2 in the quasi-synchronic behaviour) we have observed that the structures of the histograms in terms of peak locations remained identical for all cells; that is all cells ended up organized with the same set of prominent colours. The relative amplitudes of the peaks remained similar, but the

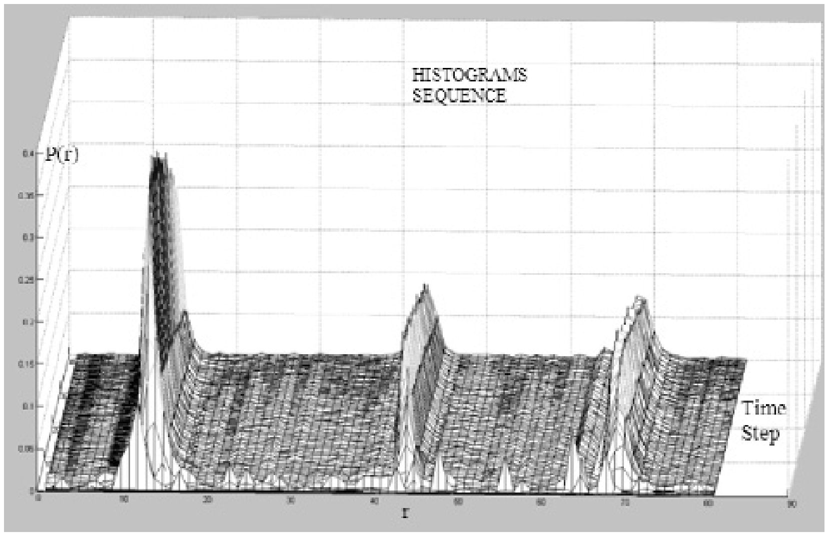

Fig. 8. Structures from SC2 Zone showing self-organization and deviation of structure trajectories.

time variations of the amplitudes (or envelopes) were slightly different for every run. This invariance property remains largely present even for different sizes of the CA grid (Fig. 9). It is therefore possible to automatically obtain multiple instances of a certain type sound. All instances will share the same structure, but with differences in the time-varying amplitudes. Thus, we can design an instrument that would not output the same exact sound twice and therefore, capable of generating more natural sequences of notes.

Although we can expect, and we assume, that every run of the same automaton would display identical behaviour, it is not straightforward to ascertain in advance the value that a specific cell would hold after a number of generations. But it is possible to ascertain the position of the predominant peaks in the histogram. Thus, this is a considerable degree of predictability with applications for example in the context of our mapping. This does not mean that we will have this predictability when starting from any initial pattern of cells values (or colours). When starting from an image with a certain level of organization of colours, the resulting histogram would probably be different from when starting from a random distribution of colours; we have not tested such cases systematically yet.

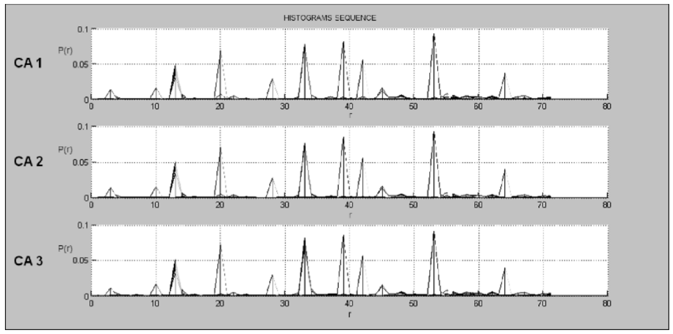

(a)

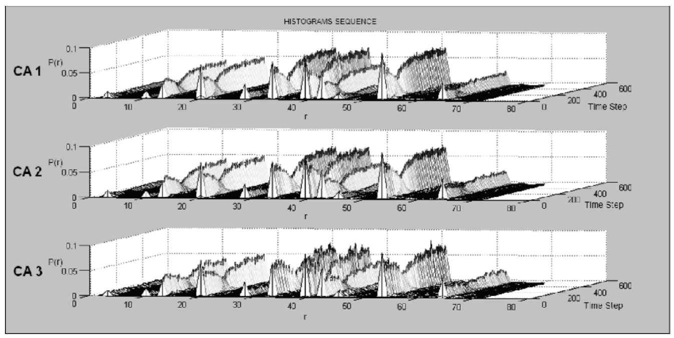

(b)

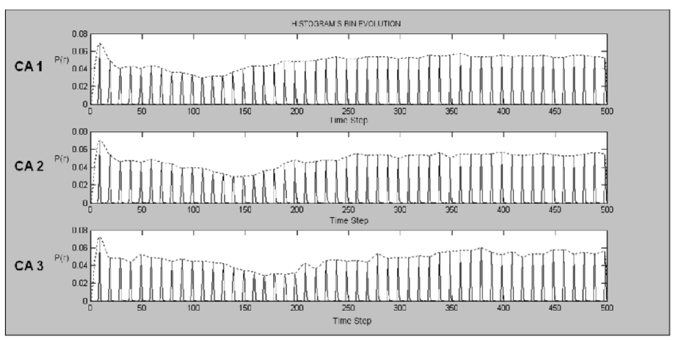

(c)

Fig. 9. Invariance property in histograms sequence for three different CA runs: CA2 has the same definition than CA1 but starting from a different initial configuration (uniform random distribution). CA3 is the same definition than CA1 and CA2, but with different size, where the size of CA1 and CA2 are 200 by 200, and the size of CA3 is 100 by 100 cells). (a) Frontal plot. (b) Lateral plot. (c) Time evolution of one bin of the histograms sequence.

4 Conclusion and ongoing work

The predictability of the outcome of CA evolution is an open problem [6]. Although a level of unpredictability is accepted, and often desired, in systems for generating music and sound, being under unpredictability conditions implies limited controllability. A lack of a reasonable level of control restricts the music or sound design process [7]. Our synthesis technique alleviates this limitation in many respects. As we have seen, the CA rules and parameters are very well reflected in the histogram. Thus, it is possible to find direct relations between the CA parameters values and their effects in the histogram. Most of them refer to the spectral dimension and some to the time dimension. For instance, the lower the value of K, the narrower is the GAP (Fig. 2). As a consequence, the quasi-synchronic behaviour produces more noise bands in the NC Zone. It is intuitive to see that a wide GAP implies less noise bands. The parameter K also contributes to the nature of the beginning of a sound; the lower the value of K, the longer it takes for the partials to settle into a well-defined histogram structure. The parameters r_1 and r_2 control the tendency of a healthy cell to become infected by infected and by ill neighbours, respectively. With the quasi-synchronic behaviour, the lower the vakues of r_1 and r_2, the wider the SC1 Zone. If these values are extremely low, the histograms sequence will evolve presenting only two peaks in the first and the last bins. Conversely, if the values are higher than the number of neighbours, the histograms sequence will evolve, blocking the first bin. With respect to the time domain, we have seen that by considering fewer neighbours it is possible to control the time-evolution of the amplitude envelopes. Once analyzed how non-sustained structures appear, it is possible to induce the same effect working with Moore neighbourhoods, by modifying the rules. One way to work with fewer neighbours is by dividing the rule parameters A and B by two. Then one can obtain the same kind of non-sustained structures that is obtained when working with Neumann neighbourhoods.

We are currently considering the possibility of using this technique to render sounds from real chemical reaction-diffusion processes.

References

[1] Doornbusch, P. (2005). *The Music of the CSIRAC: Australia's First Computer Music*. Victoria, Australia: Common Ground.

[2] Adamatzky, A., De Lacy Costello, B. and Asai, T. (2005). *Reaction-Diffusion Computers*. Elsevier.

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

[4] Pratt, W. K. (1978). *Digital Image Processing*. New York, NY: Wiley.
[5] Miranda, E. R. (2002). *Computer Sound Design: Synthesis Techniques and Programming*. Elsevier/Focal Press.
[6] Wolfram, S. (1984). Computational Theory of Cellular Automata, *Communications in Mathematical Physics*, 96:15–57.
[7] Miranda, E R. and Wanderley, M. M. (2006). *New Digital Musical Instruments: Control and Interaction beyond de Keyboard*. Middleton, WI: A-R Editions.