

Algorithmic Sound Composition using Coupled Cellular Automata

Jaime Serquera and Eduardo R. Miranda

ICCMR, Interdisciplinary Centre for Computer Music Research, University of Plymouth, UK

Email: {jaime.serquera, eduardo.miranda}@plymouth.ac.uk

Abstract: In this paper we introduce a new approach to algorithmic sound composition using a bespoke technique combining coupled Cellular Automata (CA) and Histogram Mapping Synthesis. Two CA are used: a hodge podge machine and a growth model. The latter serves as control of the former. The hodge podge machine can exhibit different kinds of behaviour depending on the values of a set of rule parameters. Our method explores the fact that different simultaneous behaviours can be evolved within the same automaton if we bring into play different sets of parameter values. However, we restrict the number of parameter sets to two. Therefore, the CA growth model will have only two states and will delimit two dynamic zones in the hodge podge machine, each of which governed by a different set of parameter values. The predictable evolution of the two zones will produce a controlled dynamic sound spectrum. Among all the possibilities that this process affords for the composition of a variety of sounds algorithmically, we highlight its application to the attack portion of a sound, making it dynamically more complex than the rest of the sound.

Keywords: Sound Synthesis, Cellular Automata, Histogram Mapping Synthesis, Additive Synthesis.

1. Introduction

Computers have been used as a tool for composition since the dawn of Computer Science in the 1950's. The *Illiac Suite for String Quartet*, composed in the USA in late 1950's by Lejaren Hiller (composer) and Leonard Isaacson (mathematician), is often cited as the first piece of music involving materials generated by a computer; e.g., the fourth movement was generated using a Markov chain [5]. Whilst the computer can afford composers the design of very complex abstract structures by using a plethora of modelling methods (including chaotic systems), most of composers have been focusing on using these systems to process information at the so called "note level". In a nutshell, the computer is used here to process musical notes represented symbolically. In tandem to these developments, computing technology also fostered the development of digital signal processing tools, which is increasingly giving composers access to very fine control of sounds. In this case, composers are no longer limited to the capabilities of acoustic musical instruments to play their music. That is, now composers have the possibility to compose "their instruments". Nowadays, the practice of

Serquera and Miranda

musical composition may also include the synthesis of sounds at the micro-level; that is at the level of their spectral components and split second variations. In this paper we focus on using cellular automata to compose sounds at the crossroads of the art of algorithmic composition and the art of sound synthesis.

Cellular automata are mathematical/computational models normally implemented as a regular grid of cells in one or more dimensions. Each cell may assume any state from a finite set of V 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.

Cellular automata have been of interest to computer musicians because of their emergent structures –patterns not created by a single rule but through the interaction of multiple units with relatively simple rules. This dynamic process leading to some order allows the musician to explore new forms of organization. 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. In the latter respect, even though novelty is a goal, it is a common practice to model and simulate certain properties of acoustic instrument sounds. In sound synthesis CA are normally used for controlling over time the parameters of a synthesis instrument. Many of the synthesis techniques demand enormous amounts of control data for obtaining interesting results, making it difficult to be controlled manually. CA represent a solution to this problem because with few parameter specifications it is obtained massive amounts of structured data. The goal is to transfer the structured evolution of CA onto the sound synthesis domain. This is always done through a mapping, a set of correspondences between different domains.

There have been different mapping attempts [2] ranging from direct assignments of CA values, like in Lasy [3], to higher-level approaches intending to map the overall CA behaviour, like in Chaosynth [6]. We are interested in the second type of approach. Our research strategy is based on the analysis of CA evolutions by means of digital signal processing techniques in order to discover structural information of their organization. Then we proceed with the mapping of the analysis results onto appropriate synthesis parameters. Histogram Mapping Synthesis (HMS) is a fruit of this approach.

The control is the main problem 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. But all this does not mean that they are uncontrollable. In this paper we provide a mechanism for the control over time of a hodge podge machine by another cellular automaton, a growth model.

2. Histogram Mapping Synthesis (HMS)

Histogram Mapping Synthesis is a recently new sound synthesis technique. It was originally conceived working with CA but it is not restricted to them. Concisely speaking, the method involves a mapping from histograms onto sound spectrograms.

The functioning of a two-dimensional automaton is considered as a sequence of digital images and it is analysed by histogram measurements of every CA image. Such a CA analysis gives a histogram sequence.

The histogram of a grey level digital image is a graphical representation of the number of occurrences of each grey level¹ in the image. By dividing the number of occurrences by the total number of pixels of the image, the histogram is expressed in probabilistic terms giving an estimate of the probability of occurrence of each grey level in the image.

In general terms our mapping method works as follows: the bins of the histogram sequence are considered to be bins of a spectrogram. With an appropriate automaton, in the histogram sequence it is possible to find structural elements resembling spectral components of a sound. For example, from a histogram analysis of the hodge podge machine we discovered structural elements similar to sinusoidal components, others similar to noise components and others similar to transients [7]. This makes such a mapping process distinctive; in most other cases there is not an intuitive correspondence between the components of the automaton and the components of a sound.

With these structural elements we can design the time varying frequency content of a sound; we can build a spectrogram. This spectrogram can be rendered into sound using different synthesis techniques –the structural elements of the histogram sequences become control data for the synthesis program.

¹ Apart from this definition, in this paper we refer to colours instead of grey levels because we usually display the CA in the computer screen using a palette of different colours.

3. Two Cellular Automata

In this research two CA are coupled, a hodge podge machine and a CA growth model.

3.1 Hodge Podge Machine

The hodge podge machine is a mathematical model of the oxidation of carbon monoxide. During the numerical investigation of such a cellular automaton, 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-Zhabotinsky reaction [4].

In the hodge podge machine, 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 original hodge podge rules can be found in [4]. Our investigations include the experimentation with modifications in the rules in order to explore new musical possibilities [7]. For this research we have implemented the following transition rules:

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

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 $\in \mathbf{N}$; round stands for the nearest integer; and V is the number of possible states that a cell can adopt. For this research we have considered the Moore neighbourhood and periodic boundary conditions, i.e., a torus.

The hodge podge machine can exhibit different kinds of behaviour depending on the values of its rule parameters K , r_1 , r_2 and V . For sound synthesis purposes we identified three suitable behaviours [8], two of which are considered in this research: the quasi-synchronic and the spiral waves. Each of these behaviours can be obtained from a variety of different combinations of parameter values.

The quasi-synchronic behaviour (in which all the cells reach the maximum state almost simultaneously) generates a type of histogram sequences from which we can obtain structures similar to sinusoidal and noise components. After the maximum state is reached, patterns of distorted circumferences emerge. These shapes will create narrow bands or peaks in the histogram which are similar to sinusoidal components. From here, the cell values grow towards the maximum state and the boundaries of the distorted circumferences become less defined, creating wide bands in the histogram which are similar to noise bands. At each cycle of the automaton, this process is repeated but with slightly different distorted circumferences shapes, creating structures in the histogram sequences with time varying amplitudes. In the histogram sequences we differentiate a number of bin-zones according to the structures that we can find in them. In this research we focus on the structures similar to sinusoidal components (or partials) that are located in a zone close to the K^{th} bin.

The spiral wave behaviour creates a type of histogram sequence that is of great interest for sound synthesis. When the spiral waves start to develop, they create structures in the histograms that resemble sound partials with increasingly high amplitudes. Once the spiral waves have completely evolved, they often expand themselves covering the whole CA grid. This creates sequences of images which are cyclic and stable. At this point, the amplitudes of the partials often stop increasing and settle to relatively stable values.

Broadly speaking, the most relevant fact derived from the histogram analysis of our hodge podge adaptation exhibiting these two behaviours is that starting from a uniform random distribution of colours the automaton self-organizes through a very specific set of predominant colours.

3.2 CA Growth Model

The growth model used in this research is the one we developed in [9] as an extended version of the multitype voter model.

We implemented the multitype voter model with this transition rule: a number between 0 and 1 is chosen as to be the update probability for all cells. Then, for each cell in the grid, a random number between 0 and 1 is generated at every time step. If the random number generated for the given cell is higher than the update probability, then the state of the cell changes to that of one of its neighbours selected uniformly at random. Neighbour is defined as the four orthogonally adjacent cells: north, east, south and west. The boundary conditions are periodic.

The CA growth model springs out of the following modifications to the previous rule: we define only two cell states, empty and occupied. The

functioning is the same (the same rule), but in order to ensure a growth we impose that occupied cells can not become empty. With this, the occupied cells will grow covering the whole automaton.

4. Simultaneous Behaviours

We conducted an experiment delimiting two halves in a hodge podge machine each of which governed by a different set of parameter values in order to evolve two simultaneous behaviours. We only established different values for r_1 , r_2 and K ; V was the same on both halves.

From that experiment we 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 example, if we run a quasi-synchronic behaviour in one half and a spiral wave behaviour in the other, it usually happens that the spiral waves emerging on one side of the automaton propagate to the other side but with a noticeable visual change so that we can differentiate two halves in the automaton (Figure 1).




Figure 1. Simultaneous behaviours in a hodge podge machine. The left half has the rule parameters adjusted to produce a quasi-synchronic behaviour whereas the right half has the rule parameters adjusted to produce a spiral wave 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. And that fact is well reflected in the histograms. 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 differentiate visually 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 can not distinguish two halves. However, from a histogram analysis we can always identify two superposed histogram sequences.

5. Adding Dynamism

The next step to further explore the possibility of producing different simultaneous histogram sequences is to add dynamism to the process.

The CA growth model (GM) will control the hodge podge machine (HPM) by determining the parameter values that govern each of its cells. We use the GM to create and evolve different regions of parameter values for the HPM. We bring into play only two different sets of parameter values, one associated to empty cells of the GM (set of parameter values E) and the other associated to occupied cells (set of parameter values O). Both two CA have the same size and therefore if a GM cell is empty that means that the cell located in the same place in the HPM will be governed by the set of parameter values E. On the contrary, if one cell of the GM is occupied, the corresponding cell in the HPM will be governed by the set of parameter values O. Both two sets of parameter values will be adjusted to produce the same type of behaviour in order to work with the same type of histograms.

The GM will evolve in the form of predictable transitions. Such an automaton with two states allows two types of transitions. The most immediate one is to go from the set of parameter values E to O (Figure 2). That will produce a smooth transition between two histogram sequences and will lead to the production of a cross-fade sound effect.




Figure 2. A CA growth model evolution starting with all the cells empty. At some point it is introduced an occupied cell in the middle of the automaton which will grow covering the whole automaton.

The second type of transition is more interesting for sound synthesis. It consists in going from two simultaneous sets of parameter values, O & E, to O. An example of such a transition is illustrated in Figure 3.




Figure 3. A CA growth model evolution starting with occupied cells in the left half and empty cells in the right half.

That means starting with two superposed histogram sequences and ending up with just one. With this, working with the quasi-synchronic behaviour we can simulate an important characteristic of acoustic instruments: they usually produce more partials in the attack portion than in the rest of the sound. In addition, our system models another process observed in acoustic instruments in which energy passes between the various modes of vibration, some increasing in amplitude, some decreasing [1] (Figure 4).




Figure 4. Zoom of a histogram sequence of a hodge podge machine exhibiting two simultaneous quasi-synchronic behaviours and controlled over time by a CA growth model. 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.

6. Conclusions

In this paper we have provided a coupling method for controlling over time an adaptation of the hodge podge machine by another automaton, a CA growth model.

The choice of the latter automaton has been determined by the musical applications that it can afford with the HMS technique. We have achieved the

synthesis of sounds with more complex spectra in the attack portion than in the rest of the sound, an important feature found in sounds produced by acoustic instruments.

Other CA can be suitable for controlling the hodge podge machine with our method. For instance the multitype voter model can produce dynamic areas evolving with more unpredictability than the growth model.

References

- [1] Aldoshina, I. et al.: The analysis of peculiarities of Russian bells acoustic parameters. 114th AES Convention. Amsterdam (2003)
- [2] Burraston, D. and Edmonds, E.: Cellular Automata in Generative Electronic Music and Sonic Art: A Historical and Technical Review. *Digital Creativity* 16, 165–185 (2005)
- [3] Chareyron, J.: Digital Synthesis of Self-Modifying Waveforms by Means of Linear Automata. *Computer Music Journal* 14, 25–41 (1990)
- [4] Gerhardt, M. and Schuster, H.: A cellular automaton describing the formation of spatially ordered structures in chemical systems. *Physica D* 36, 209–221 (1989)
- [5] Hiller, L. A. and Isaacson, L. M.: *Experimental Music: Composition with an electronic computer*. McGraw Hill. New York (1959).
- [6] Miranda, E. R.: At the Crossroads of Evolutionary Computation and Music: Self-Programming Synthesizers, Swarm Orchestras and the Origins of Melody. *Evolutionary Computation Journal* 12, 137–158 (2004)
- [7] Serquera, J. and Miranda, E. R.: Spectral Synthesis and Control with Cellular Automata. *Proceedings of the International Computer Music Conference*. Belfast (2008)
- [8] Serquera, J. and Miranda, E. R.: Cellular Automata Sound Synthesis: From Histograms to Spectrograms. *Automata 2008: Theory and Applications of Cellular Automata*. Luniver Press, Frome (2008)
- [9] Serquera, J. and Miranda, E. R.: Evolutionary Sound Synthesis: Rendering Spectrograms from Cellular Automata Histograms. *Proceedings of EvoApplications / EvoMUSART, Lecture Notes in Computer Science*. Springer-Verlag, Berlin (2010)