2007

Cellular automata as spectra: beyond sonification into composition

Warren A. Burt
University of Wollongong, wburt@uow.edu.au

Publication Details

Burt, W. A. (2007). Cellular automata as spectra: beyond sonification into composition. In A. Riddell & A. Thorogood (Eds.), Australasian Computer Music Conference (pp. 27-33). Canberra: The Australian National University.

Warren Burt
University of Wollongong; Illawarra Institute of TAFE
PO Box U-27
Wollongong, NSW 2500
Australia
waburt@melbourne.dialix.com.au

Cellular Automata as Spectra: Beyond Sonification into Composition

Abstract

Although Cellular Automata (CA) have been used in music composition for a number of years, most of the writings on them have either concentrated on the mechanics of the CA themselves, or else have dealt with implementations of CA in particular methods of sound synthesis. We instead concentrate here on the use of CA in "Scrabbles," a six-minute long musical composition, dealing mostly with the way the composition was structured using CA materials, and the algorithmic program that was written to choose parameters for the individual sounds of the piece.

Introduction

Although Cellular Automata (CA) have been used in music composition for a number of years, most of the writings on them have either concentrated on the mechanics of the CA themselves, or else have dealt with implementations of CA in particular methods of sound synthesis. Here we will instead concentrate on the use of CA in one particular musical composition, and talk mostly about the way in which the composition was structured using CA materials. Part of the maturation of areas of computer music is when the understanding of a technique is taken for granted, and is augmented by a deeper consideration of how and why particular techniques are used in particular compositions. Consider all the papers written on, for example, frequency modulation in the 70s and 80s. Today, fm is simply one more technique in the synthesis toolkit: the compositional use of the technique is now of primary concern.

History

There is an extensive literature on Cellular Automata and their use in music. The most extensive consideration of the inner structure of CA and musical applications is Dave Burraston's recently completed PhD thesis, *Generative Music and Cellular Automata.* (Burraston, 2006). Stephen Wolfram's *A New Kind of Science* is an exhaustive survey of CA (Wolfram 2002), and his research has recently been incorporated into an interactive website, *WolframTones,* in which one can hear musical realizations of various CA (Wolfram Research 2007). Unfortunately, I find most of the music generated with the algorithms on this site to be disappointing. An early software synthesizer which used CA in its granular synthesis engine was Eduardo Miranda's *Chaossynth*. Currently being rewritten, the earlier version is now available free from www.nyrsound.com (Miranda 1999). These are but 3 references to an enormous field. For more information, refer to the massive References section of Burraston's thesis.

As interesting as this research is, it did not engage me for this project. I was not so much interested in the mechanisms of CA, other than understanding them well enough to be able to know what was going on inside them, but was more concerned with using them as spectra, in order to get sounds with a certain degree of complexity. I have used CA before in a number of compositional contexts. They were used to choose pitches, durations and timbres, making clouds of slowly moving sound in *Voices, Tuning Forks, and Accordion* (1986), a composition for choir with microtonal tuning forks and microtonal accordion. The CA used for this was a mutant version of Conway's *Game of Life* written by my friend Dr. Henry Hunter. I also used CA to assemble extremely slowly changing inharmonic spectra in the computer composition *The Easy Beauty of Parallel Lives*, a movement of *Graphic Descriptions* (2003). The CA used for this piece was Peter Meyer's Q-Life. I also incorporated a one-dimensional CA into my early DOS interactive program *Randie,* described in "Interactive Improvisations with Electronic Music Systems," and with which I performed real-time interactive music on a laptop from 1991 on (Burt 1991). *Randie* produced MIDI output, which was performed with a variety of instrumental, electronic, and sampled timbres.

Sounds with Complexity – A Priority

For the current work, however, I was interested in using CA to assemble sounds with more complex, constantly changing timbres. This would entail finding CA which would have the right kind of changing graphics to produce sounds of the sort I was interested in. The CA produced would then be converted to sound using a graphic synthesis program. There are a number of these, such as the freeware *Coagula* for the PC (Ekman, 2003) and the commercial programs *Atmogen* (Sonorous Codes, 2004) for the PC, and *Metasynth* (Wenger, 2006) for the Mac. They all work on a similar principle, which is an extension of Western Music notation's up-and-down for frequency and left-to-right for time paradigm. That is, lines drawn closer to the top of the graphic produce higher frequencies, while lines drawn more to the right of the graphic produce later sounds. Additionally these programs add colour for placement in stereo space, and brightness for amplitude. All the programs generate sound with sine waves, but some, such as *Metasynth*, allow other samples to be used, while *Coagula* adds the feature that while pixels on the red-green continuum produce sine waves, the addition of blue to any colour adds a proportional amount of band-limited noise to the sound. For this composition, *Coagula* was used.

Notice that the conception of the kind of sound I wanted came first. I wanted to create sounds with rapidly changing, complex spectra. It occurred to me that CA might produce diagrams that, when synthesized with *Coagula*, could do this. *Sound* came before *technique*. Computer music, in general, has been, is, and, for the foreseeable future, as long as we are (rightfully) concerned with the developments of new tools and techniques, will continue to be, stuck in a scientistic rut. That is, as a field, computer music has the trappings of science (peer-reviewed papers, an emphasis on new uses of new technologies, etc), but lacks the ability to have rigorous proof that constitutes the essence of science. There is no way to prove music. (Success or failure in any conceivable market does NOT constitute proof of any kind whatever.) We should not pretend we are scientists. Nor do we need their approval (although their support would be nice). We are artists, and we operate with a different set of criteria. Only the quality of our compositional thought – that is, how we use our new tools – will make our efforts worth listening to.

Palette Choice as a Compositional Act

To see if any CA would produce the kind of diagrams that might be useful to me, I used *Mirek's Cellebration*, a freeware CA explorer for the PC (Wojtowicz, 2001). I wanted to synthesize spectra – that is, I wasn't interested here in synthesizing broadband noise, but sounds with discrete partials which changed frequency and amplitude rapidly. For just as most matter is empty space between fundamental particles, so most spectra (whether harmonic or inharmonic) consist of empty "frequency space" between partials. That is, a spectrum usually consists of a small number of partials at specific frequencies, not of partials at all available frequencies. This would mean that my CA diagrams would have to consist of mostly black, the colour used by graphic synthesis programs to indicate silence. Unfortunately, most CA programs completely fill in the image, as in Figure 1.

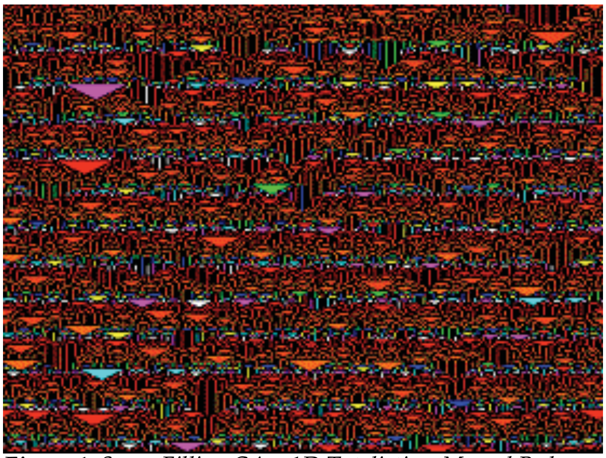

Figure 1: *Space Filling CA – 1D Totalistic – Marvel Rule*

In order to get diagrams that would be useful to me, I would have to either find a CA rule that generated mostly black, or else invent a colour palette that would "filter" the CA so that broad areas of black were generated. The "1D Totalistic" rules produced diagrams that consisted of mostly bands of activity occurring in horizontal lines, and I made a palette which consisted of alternating areas of black and colour, so that successive generations of the CA would be either sounding (coloured) or silent (black). By making the older generations all black, I guaranteed that most of the diagram would be black, with activity constrained to a certain number of horizontal regions. This would give me the kind of drawing I wanted. The "Coagula01" palette consisted of the following values and colours.

Generation	Colour
0	Black
1	Light Blue: R=128; G=255; B=255
2	Black
3	Red: R=252; G=0; B=0
4	Black
5	Green: R=0; G=255; B=0
6	Black
7	Light Yellow: R=255; G=255; B=128
8-24	Black

Table 1: *Colour listing of "Coagula01" palette*

Notice that the Light Blue and Light Yellow colours both have blue in them. This means that both those colours will produce some level of band limited noise. Applying this palette to the CA shown above in Figure 1 produced the result shown in Figure 2.

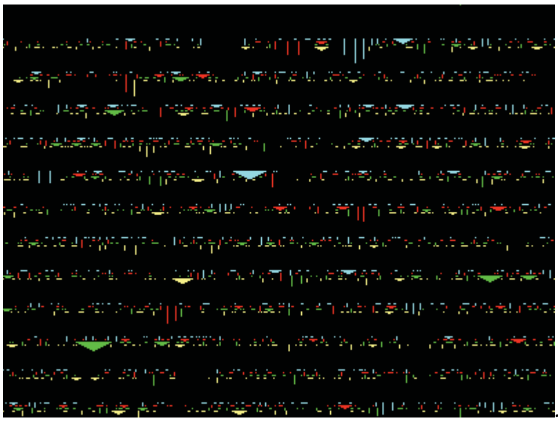

Figure 2: *1D Totalistic Marvel CA with Coagula01 Palette*

In this diagram, activity is restricted to horizontal bands, yet the kind of activity in each band is rapidly changing. This looked like it could produce the kind of sounds I was interested in. And in fact, when quickly processed by *Coagula*, it did indeed produce the kind of sonic complexity I had been searching for. My first impression on hearing the sound was that this was a sonic world I'd like to live in for hours. I quickly realized, however, that if I wanted the piece heard in the context of an academic conference, with its crowded schedules and severe time constraints, I would have to limit the duration of any piece made with these sounds. I therefore decided to make a short piece with many contrasting sounds, so that the range of sounds made with these graphics could be quickly heard. In my case, this is actually a welcome contrast, as my previous piece, *Proliferating Infinities* (2006), is just over 13 hours in length. This piece will probably remain unheard by all but a few friends until I can figure out a context in which it can be presented in such a way that its contemplative nature can be fully experienced. And as an aside, I wonder if it will ever be possible for ACMA, or similar organizations, to organize an event in which compositions of longer duration could be experienced in a non-pressured manner.

Exploration of the available 1D Totalistic CA rules gave me eight different rules which produced diagrams such as these. These rules, in Mirek's program, were called Class 4C, Forest, Gears 1, Marvel, No Name 1, Porridge, The City, and Tulips. All of these were different from each other, but all had activity occurring in parallel bands with changing activity within the bands. Of these, Class 4 C produced kinds of activity that changed over time, with alternating areas of straight vertical lines and smaller textures, as shown in Figure 3, in which a complex texture expands "out of" the area of vertical lines in the middle of the graphic. Each of the rules had their own intrinsic characteristics, and each produced differently detailed sonic textures with the larger family of sounds.

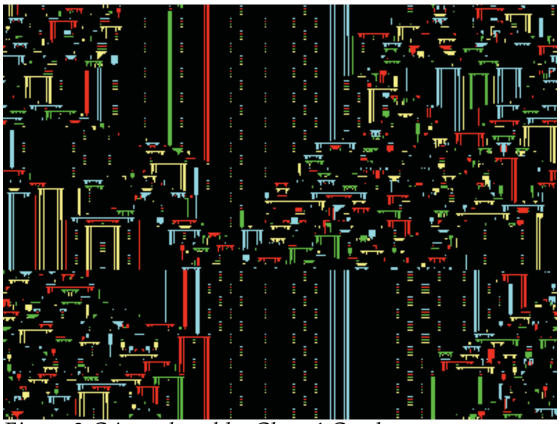

Figure 3 CA produced by Class 4 C rule

The City is similar to Class 4C. There are more areas of horizontal lines with small vertical "stalks", and less areas of large horizontal lines and empty space, but it has the same ability to have moving shaped frequency bands as Class 4C and Gears 1. The many long horizontal lines produce a series of rapidly sputtering very short noisebands.

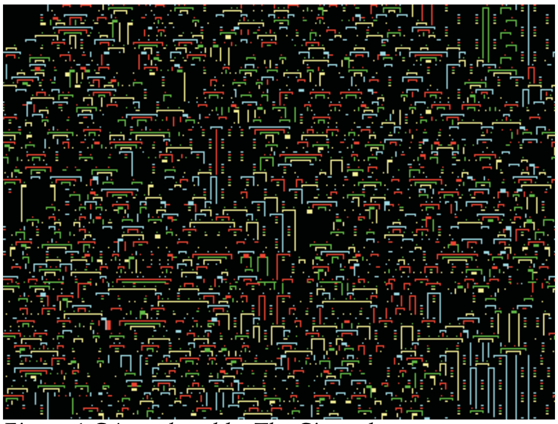

Figure 4 CA produced by The City rule

No Name 1 is similar to Marvel (see Figure 2), but has horizontal bands that are more narrowly spaced, and there are longer vertical lines which sometimes connect two adjacent horizontal bands. This make the spectrum thicker – more like a buzzing, dissonant tone cluster, than a well spread out chord with active spectral bands.

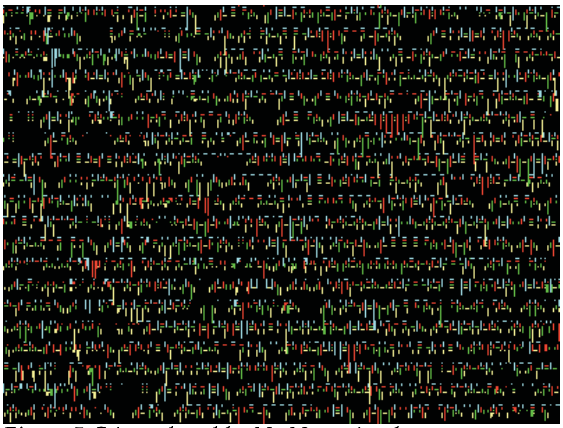

Figure 5 CA produced by No Name 1 rule

Forest is again, similar to Marvel, but extremely long vertical lines can appear. These make extremely short, sharp noise bursts, which make pleasing ornaments over the top of the overall active sonic texture. Forest can also have "holes" which by creating silence, give gestural shape to the overall texture.

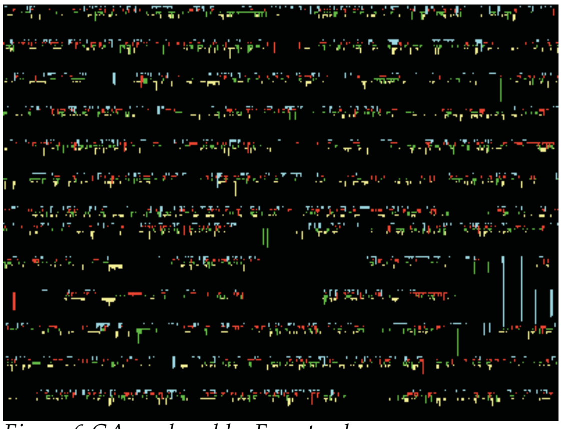

Figure 6 CA produced by Forest rule

Gears 1 produces bands of downward pointing triangles. This makes a quite surprising texture of rapidly attacking narrow noisebands that begin and change pitch unpredictably, while still retaining the sense of a sustained chord that occurred with Marvel. (It was this "sustained chord with great inner life" sound that first got me interested in exploring these particular CA further.) Very large "holes" can also appear in the Gears 1 diagrams, and these give a very pronounced sense of shaping of frequency and gesture to the overall spectra.

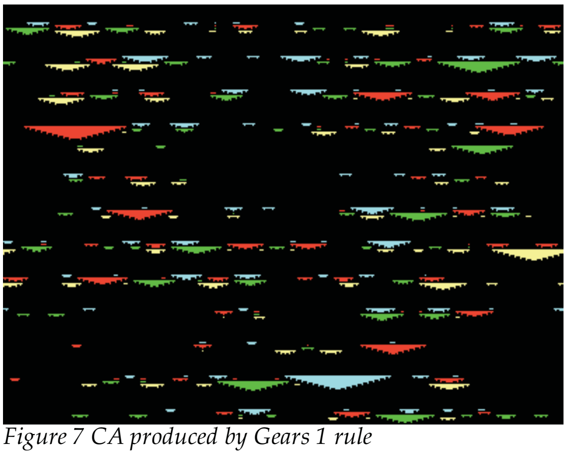

Figure 7 CA produced by Gears 1 rule

Porridge is a thicker, more regular version of Gears 1. Here the rapidly attacking narrow noisebands are very narrow, much more closely spaced, and the occur with much more regularity and faster than in Gears 1. The sound produced is somewhat of a cross between the sustained chord of Marvel, the sputtering noisebands of The City, and the tiny, expanding-then-contracting noisebands of Gears 1.

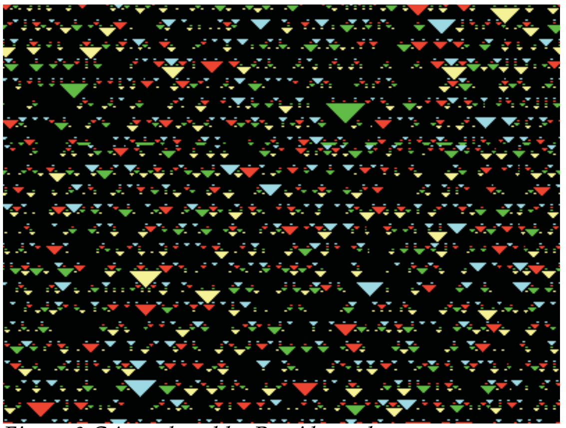

Figure 8 CA produced by Porridge rule

Tulips has very wide descending triangle shapes broken up by black vertical bands. These, surprisingly, only contribute a sense of wider noisebands to a texture that is basically very similar to Porridge or Gears 1. It should be mentioned again that all of these diagrams produce a similar family of sounds. The sonic details differ from rule to rule, but the overall sense of gesture – a sustained sound with a very rich and lively internal life – is similar over the whole family. This similarity is why these rules were chosen, and many other 1D Totalistic rules were rejected. I chose particular rules based on their ability to give me the kind of sound I wanted, trusting the mathematics of each CA to provide interesting internal details within my chosen overall sonic idea.

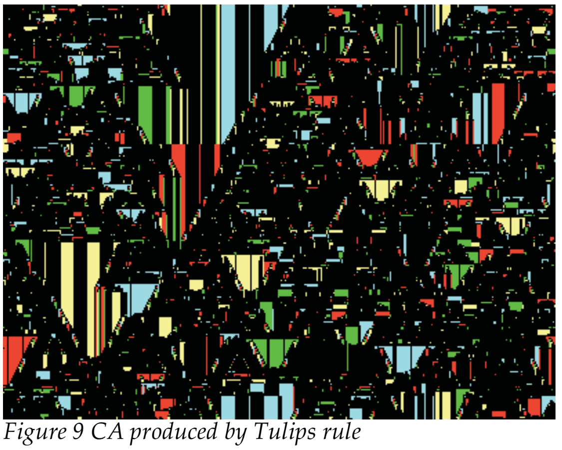

Figure 9 CA produced by Tulips rule

Choosing Parameters

Coagula turns graphics into sound, but has certain parameters that must be set. Among these are the duration of the sound, and the high and low frequencies that define the outer limits of the synthesis. Different settings of these parameters will radically change the nature of the sound. Additionally, I experimented with tilting the graphics slightly, by up to 5 degrees away from horizontal, to see if a sense of rising or falling pitch would be present in the spectra. It was, so I decided that some of the drawings would be tilted slightly to produce this as well.

After much experimentation and listening, I decided that there would be sequences of sounds made by these spectra, and that two sequences at a time would be mixed to make a more complex final result. Each sound would be made by a graphic produced by one of the eight possible rules. 70% of the time, the sound would be made by a straight graphic, while 30% of the time, the graphic would be tilted between 1 and 5 degrees in either direction. Each sound would last between 3.0 and 20.0 seconds, graduated in 10ths of a second. The selection of durations would be biased differently for each sequence of sounds. In some, short durations would be favored, in others, longer durations would be prominent. The use of graphics 400 pixels wide mean that at the longest duration, there would still be 20 pixels each second, thus avoiding the problem of the "rhythm of the pixels" becoming a predominant aspect of the sound. The lowest frequency at the start would be between 50 and 800 hz, while the highest frequency at the beginning would be between 800 and 12800 hz. The CA height of 300 pixels would insure that the spacing of possible frequencies would always be less than 35 cents per step. This would avoid the problem of inappropriate melodic material being formed through too wide spacing of frequencies. The ranges of the parameters are summarized below.

CA type	Class 4C, Forest, Gears1, Marvel, No Name 1, Porridge, The City, Tulips – chosen with equally-weighted random numbers
Straight or tilt?	70% - straight; 30% tilted – 3% each 1 to 5 degrees clockwise or 1 to 5 degrees counter-clockwise.
Duration	From 3.0 to 20.0 seconds. Biased random choice – linear 2x and linear 4x distributions (favoring low values), and linear inverse 2x and linear 4x distributions (favoring high values).
Low frequency at start	From 50 to 800 hz, chosen in octaves with equally-weighted random numbers
High frequency at start	From 800 to 12800 hz, chosen in octaves with equally-weighted random numbers

Table 2: *Parameter choices and ranges*

To make the decisions for each individual sound, an algorithmic gesture generating program was written in John Dunn's *ArtWonk* (Dunn 2007). I knew very well the range of sounds I wanted. Using algorithmic methods to choose the exact details of the sound allowed me to experience a sense of surprise as I observed the structure of the piece emerging, sound by sound. It also allowed me to get results I couldn't have gotten by picking parameters by instinct. And I have worked with algorithmic compositional methods so long that by now, my compositional instincts for choosing which method of choice to use for each parameter of a sound are highly developed. This led, for example, to using a different random distribution for choosing durations in each of the four lists. I knew the kind of spread of durations that would be generated by each distribution, and I chose each one accordingly.

Table 3 gives a listing of the durations produced by each of the four different Linear Nx distributions. Durations for each list were chosen from between 3.0 and 20.0 seconds, and were totaled up for each new sound. When the total duration exceeded 180 seconds, or 3 minutes, that list was ended. Reflecting the increasing average durations of the sounds in each list, the four lists had 30, 20, 14 and 11 sounds respectively, as the average duration increased from short to long.

Looking down each list, one can quickly see how, for example, Linear 4x chooses mostly durations at the very low end of the range, while Linear 2x chooses a low weighted selection but one more biased towards the middle, while the inverse distributions simply reverse this bias (I was delighted, by the way, that the first list began with the highly improbably long 18.3 second duration).

It might be mentioned that these distributions were part of a package of about 30 different random distributions, functions and attractors I wrote for John Dunn's *ArtWonk*, which have since been incorporated into the program. Since designing these functions, I have spent a considerable time playing with each, until I now almost intuitively know which distribution is going to give me what kind of control for a given parameter.

List #	1	2	3	4
Random Distribution	Linear 4x	Linear 2x	Linear Inv. 2x	Linear Inv. 4x
Durations	18.3	6.4	18.3	19.6
	3.4	6.2	17.6	11.2
	8.4	3.4	18.5	19.6
	8.1	11.1	16.1	15.7
	5.2	8.4	7.9	16.2
	4.3	14.7	16.8	15.1
	10.2	8.1	11.6	15.5
	4.7	14.4	9.8	16.3
	6	10.9	7.3	16.4
	6.9	6.8	10.6	19.8
	4	14.4	6.8	11.4
	3.8	5.2	14.4	
	3.5	6.3	10.2	
	3.7	14.3	16.6	
	3.7	4.3		
	3	10.2		
	3.8	10.8		

	10	10.8		
	3.2	4.7		
	5.6	13.8		
	5.3			
	6.2			
	4.3			
	5			
	8.3			
	5.7			
	14.2			
	3.6			
	4.6			
	4			

***Table 3:** Durations chosen by different random distributions*

In each subsequent list, as well, increasing frequency ranges were chosen. Frequency was chosen in a two stage manner. In the first stage, a random number, 1, 2, 4, or 8, is chosen with an equally weighted random distribution. In the second stage, a random number between either 50 and 100 (for low frequencies) or 800 and 1600 (for high frequencies) is chosen. These two numbers are multiplied together to get the required low or high frequency limit for the synthesis. This two stage method generates frequencies with an even spread of frequencies among all the octaves. Simply choosing a random number, even with equal-weighting, between the low and the high frequencies of a particular band would tend to not give frequencies evenly distributed among all the possible octaves (especially in the short term, when only 30 or less values are being chosen). To insure that all octaves were visited, the two stage method described here was made. In Lists 2, 3, and 4, the high frequency band was increased by one octave. In Lists 3 and 4, the low frequency band was dropped by one octave. In list 4 as well, choices were made from only the top two and the bottom two octaves. This meant that each subsequent list covered, on average, a wider pitch band than its predecessor. When the lists were mixed into the final sections, this meant that the piece would not only slow down, it would cover a wider frequency range at the end than it did at the beginning.

In the *ArtWonk* patch, pressing a button gave new values for each new sound. As shown, these were written into lists. The interface for the program is shown in Figure 10.

Figure 10: ArtWonk program for generating sound parameters

To sum up the overall process of making the piece: the specified CA was first generated in *Mirek's Cellabration*. The CA graphic was then processed and, if necessary, tilted in a graphics program. The processed graphic was then converted into sound with *Coagula*, using the parameters generated by *ArtWonk*. The generated sounds were spliced together end to end in a sound editing program. Lists 1 and 2 were mixed together to make a 3 minute section. Lists 3 and 4 were also mixed together to make another 3 minute section. These two 3 minute sections were then spliced together to make the final composition. Since the durations of Lists 1 and 2 are shorter than the durations for Lists 3 and 4, this means that the section 1 has a faster composite rhythm than section 2. This gives a slight sense of deceleration to the complete work. This sense of deceleration is one which I seem to be exploring in a number of current works. Maybe it reflects my fervent desire to slow down the pace of my life.

Conclusion

At each stage of composing the piece, I first made decisions as to what kind of sound I wanted, or the range of values wanted, or the desired distribution of values. Parameter ranges or distributions were then evaluated by working interactively. Knowledge of the kind of results produced by each selection method allowed me to pick methods appropriate for each parameter, and gave me the ability to proceed with composing rapidly and intuitively.

While composing, I was afraid that the "event-oriented" nature of my decision making would make a piece that would seem too "bitsy" or too rapidly changing. On listening to the completed piece, I realized that my fears were groundless. The different frequency ranges, rhythms, and durations of each sound simply contribute to a sense of rapid variation within an overall highly active sound, all the composite sounds and frequencies fusing into a single lively texture.

The final composition, "Scrabbles," achieves my desire – a science of complexity is used to make a complex sound object, one with great inner life and dynamism. Its immediate predecessors might be pieces such as Iannis Xenakis' *Bohor* (1962) and Larry Austin's *Caritas* (1969), although with its short duration and rapid textural changes, it perhaps can be seen as a rapid slideshow of snapshots in contrast with the vast frescos of those two magisterial compositions.

References

Austin, Larry (1969) Caritas, computer music on tape. Source Records, SR17, CPE, 1970.

Burraston, D. (2006) Generative Music and Cellular Automata. PhD Thesis, Univ. Technology Sydney, Australia. http://www.noyzelab.com/research/research/thesis.zip

Burt, W. (1991) Interactive Improvisations with Electronic Music Systems. *Sounds Australian*, Summer 1991-92.

Burt, W. (1986) Voices, Tuning Forks and Accordion, on *Music for Tuning Forks*, Tropicapricorn CD, www.tropicapricorn.com.

Burt, W. (2003) The Easy Beauty of Parallel Lives from Graphic Descriptions, on *Road Works – Some Numbers – Graphic Descriptions – Stretti,* Tropicapricorn CD, www.tropicapricorn.com

Burt, W. (2006) *Proliferating Infinities*. Unpublished 13 hour composition for harp samples in 264 different microtonal scales.

Dunn, J. (2007) ArtWonk (3.0 beta) http://algoart.com

Ekman, Rasmus (2003) Coagula – graphic sound synthesis freeware http://hem.passagen.se/rasmuse/Coagula.htm

Miranda, E. (1999) Chaossynth, software synthesizer. http://www.nyrsound.com/

Sonorous Codes (2004) Atmogen – graphic sound synthesis software. http://www.sonorouscodes.net

Wenger, E. (2006) MetaSynth4 – graphic sound synthesis software. http://uisoftware.com/MetaSynth/

Wojtowicz, M. (2001) Mirek's Cellabration – freeware cellular automata explorer. http://www.mirwoj.opus.chelm.pl/ca/index.html

Wolfram Research (2007) WolframTones: An Experiment in a New Kind of Music – interactive website. http://tones.wolfram.com/

Wolfram, S. (2002) A New Kind of Science. Wolfram Media, Champaign, Il.

Xenakis, Iannis (1962) Bohor, musique concrete on 8-channel tape on *Xenakis: Electronic Music*, EMF CD, EM102 www.cdemusic.org