

**Robert W. Esler**

University of California, San Diego  
9500 Gilman Drive 0326  
La Jolla, CA  
92037 USA

resler@ucsd.edu

**Abstract**

*This essay is an investigation of methods for performing algorithmic computer music. Over the course of several years the author has been developing an interpretive context in which to approach algorithmic computer music. The author describes the different means used to interpret David Birchfield's Community Art: Resonant Energy a piece for percussion and computer. The piece is created using a genetic algorithm that controls all aspects of the work while the performer reads a score generated in real-time based on this process. The author has found that when an algorithm defines a musical process, that process determines the interpretive context of the piece.*

**Introduction**

As a performer, I am very interested in the various practices and interpretations of electro-acoustic music. Many of the works that I perform span all areas of real-time electronic music which include the 'traditional' IRCAM style, designer real-time software patches, and tape music. These methods of interacting with technology are no longer the frontier of computer music, rather they have become the standard. However, algorithmic computer music for performers has had fewer instances for thorough development of an interpretive context. It is the intent of this essay to investigate the implications of performing and interpreting algorithmic computer music, specifically David Birchfield's *Community Art: Resonant Energy* a work for percussion and computer. (Birchfield, 2002)

To begin, it is important to understand what algorithmic computer music means in the context of this essay. Algorithmic computer music is electro-acoustic music for a performer or performers that uses software to implement a specific automatic or real-time process as its musical foundation. A few examples of 'automatic

## Performing Algorithmic Computer Music: Creating an Interpretive Context for David Birchfield's Community Art: Resonant Energy

or real-time processes could be game theory, models of improvisation, water conservation, and genetic algorithms. Many of these compositions do not use traditional notation or follow standard modes of musical narrative. Therefore, the practices familiar to most musicians (i.e. interpreting notation) are not sufficient in this circumstance. Developing a formidable interpretation of this type of music is difficult. My problem in these situations was what should I do with the material of the process. I discovered that illuminating the salient characteristics of the process was the best solution.

This concept is not necessarily new. Performers have been interpreting pseudo-algorithmic musical processes since the inception of Western theory and practice. A simple, pre-computer, example could be a fugue. The fugue is an algorithmic composition. The subject, counter subject, sequences, etc. have specific functions within the piece. Each time they occur the performer will interpret them in certain ways. The processes of musical form, rhythm and harmony have guided the majority of Western Music's interpretation through history. Performers understand the ebbing processes of tonal harmony, the large phrase structures of sonata form, and the ways to appropriately shape a cadence through rhythm and phrasing. Mozart presumably used these concepts in his *Musicalisches Würfelspiele*<sup>2</sup>, or Melody Dice Game, where musical cells are placed in an eight-bar phrase. (Mozart, *Musicalisches Würfelspiele*, 1983) Each cell is composed within a certain chord quality (tonic, dominant, pre-dominant, and sub-dominant) and then can be constructed using a grid and a pair of dice. This process can generate many different versions of a simple Minuet that obey the rules of tonal

---

<sup>2</sup> This is usually attributed to Mozart, but there is some argument over the authenticity of the work.

harmony within the purported language of Mozart.

Similar arguments could be made concerning serial music and its algorithmic micro-structuring. Once a performer understands the various permutations of the row, they can make decisions regarding phrasing, voicing, and musical shape. Interpreting musical processes of this kind is intrinsically the same as performing any computer programmed musical automata. The processes are revealed, have specific functions, and collaboratively shape the interpretation of the piece. Many of these same aspects can be related to performing algorithmic music.

However, as I examined how performers interpret tonal music and early twentieth century music, I found myself needing a more relevant model of musical automata. The repertoire for this diminutive genre is scarce. There are very few pieces that represent these aspects in their entirety. Moreover, the catalogue for such works is not well defined and documentation for the few existing pieces is insubstantial. However, one poignant example is George Lewis' *Voyager*, developed in the 80's and 90's. In *Voyager*, Lewis creates a real-time environment where a musician improvises freely with a computer. This environment seemingly displays complex interaction between the performer and the computer. (Lewis, 1999) To perform this piece one would need to know various details about the system such as: 1) how the computer reads incoming musical data 2) what type of relationships are derived from this data, and 3) how the computer represents the results of this data musically. Without these concepts the effect of human-computer interaction could perhaps be unnoticed by the listener and the performance could be seen as an improvisation with a pre-recorded tape or some other non-interactive media. Lewis notes that computer music programs are not "'objective or 'universal,' but instead represent the particular ideas of their creators." (Lewis, 2000) Lewis relates these notions to cultural practice and communal collaboration, which instils a very important concept for performers using algorithmic computer music systems and performance environments. The algorithm is not music but it is the process in which music is created. Performers must consider the trenchancy of this practice; if the algorithm informs the creation of the music then it must inform the creation of the interpretation. After listening to many versions

of *Voyager*, I discovered that every interpretation was focused around this point. Each performance explores the environment through pitch, rhythm, density, form, time and texture. The performers tunnelled through all aspects of the piece uncovering the music they wanted to make.

I have used a similar approach to *Community Art: Resonant Energy* by mining the processes of the algorithm as a guide for an interpretation. This has been a useful method for myself in creating a viable interpretive context for algorithmic computer music. All aspects of the piece were filtered through this approach including macro-structure, micro-structure, instrument choice, phrasing, physical gestures, and spatialisation.

## Description of Musical Processes

*Community Art: Resonant Energy* (2002) uses a genetic algorithm as its core musical process to produce a score and computer playback in real-time. There is no live processing of the performer, only stored samples of percussion instruments. The performer reads and interprets the score that is re-drawn on screen every thirty seconds. Each new screen represents one generation of the algorithm (see Figure 2 for example of one generation). The computer plays samples from a database reflecting the instrumentation of the performer. The samples are mostly recordings of different noises made by percussion instruments without any of the original attack or decay. In general this piece can be designed for any set of instruments, however this essay will discuss a version for unpitched metal instruments.

Within the processes of *Community Art* the note supplants the gene that is defined by a genetic code, which controls the musical decisions and live realisation of the score. Traits of each note such as: amplitude, frequency, envelope, pitch clarity, and shift rates are stored in a database and determine how and when each note will be played as well as its relation to its neighbours. A larger scale hierarchical process defines the phrase structure of each generation, which normally spans thirty seconds. Notes belong to gestures, that are placed in phrases 7.5 seconds each (four per generation). The higher-level traits of the hierarchical process are defined by density, central frequency, harmonic-

ity, and rhythmic activity. The genetic features of these higher level traits dictate the overall process of the music for each generation in accordance with their role in the system in much the same way that families, communities, and nations influence their members. When a generation is parsed, each note is placed in time on

a score from which the performer reads and interprets in real-time. Colours display note type and the relative thickness of each note gives a general indication of volume. A note is either a processed sample played by the computer, or a sound made by the performer.

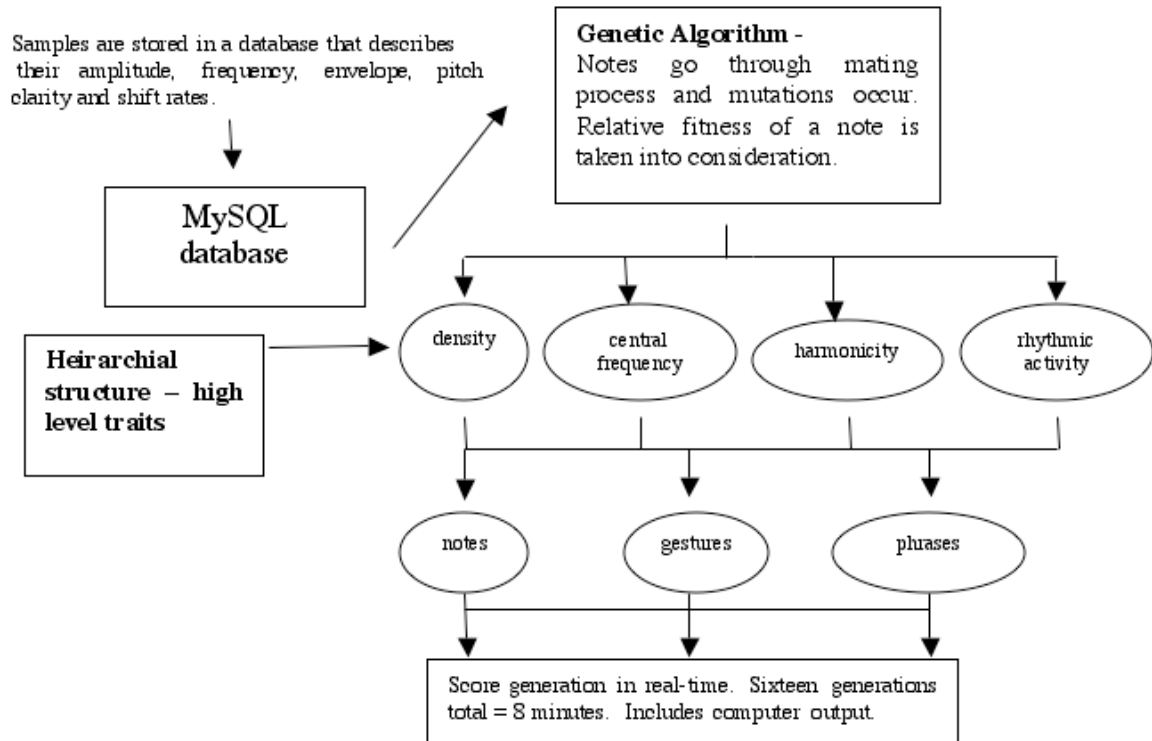


Figure 1. Flow chart defining the process of the algorithm

Figure 1 shows how samples are passed through each generation based on their genetically defined musical information. Notes, which are samples of percussion instruments, mate and create children who carry a recombinant genetic code. The computer analyses this code and searches the database for a suitable match and processes it according to the steps defined below. The performer's interface is a GUI that represents all the children that have been generated, some of which are heard, and others that are to be played by the performer. Figure 2 shows an example of this interface; there are several colours that represent how a child is processed. The GUI defines the musical geography of each generation and is used more like a topographical map rather than a literal score. The interface is most useful in understanding the changes that are taking place to the system through time. In general, the score is a

visual cue, or liaison, to the genetic processes of the piece

One particular element that the interface supplies is information about mutations. Mutations can change the genetic construction of the entire phrase or community. This may include loud violent gestures within a characteristically quiet phrase, or high frequency indications within a relatively low frequency phrase. In context these may appear as anomalies or even as dissonance to the evolving structure, however they mark an important phenomena of the genetic process. A mutation will either be adopted or selected from a community, which means it may eventually reshape the characteristics of a phrase or ultimately be eliminated. The interface displays mutations just as any other note, however, the performer can see trends within the system and gauge whether a particular series of notes is relevant to the community. This is the most useful attribute of

the interface and shows why interpreting each generation through a score is integral to the ex-

cution of the piece.

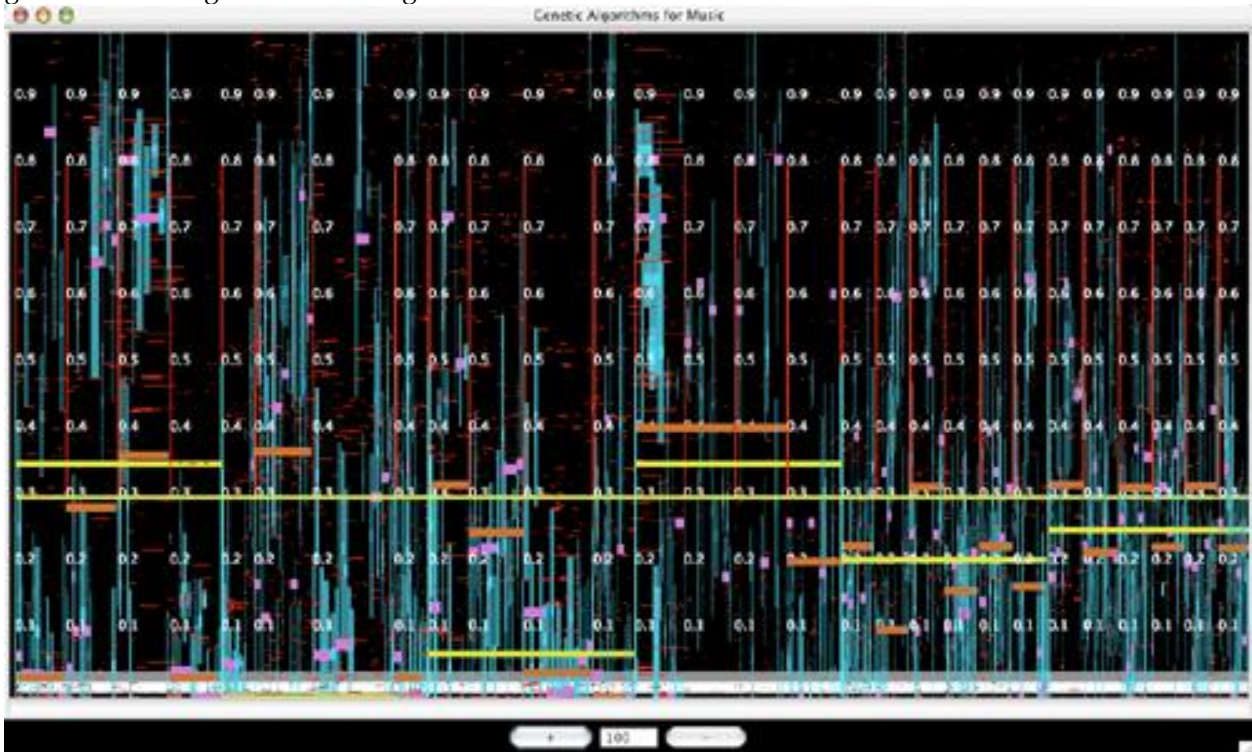


Figure 2. Score generation. X-axis = time, Y-axis = frequency

## Creating an Interpretive Context

The inherited structure of this musical organism allows for an open interpretation. Because the algorithmic process is complex and impossible to perceive on a micro-level, it is more important to realize the score's macro-structure and react to the way the computer executes notes for each generation. The macro-structure of the piece most often behaves in a certain fashion: 1) the initial generations are very dense with sound spanning all frequency bands and dynamic levels 2) the middle generations begin to reveal the musical community of each phrase, gestures start to homogenize within their neighbourhoods 3) Mutations occur periodically assigned either to the computer or the performer 4) the final generations remain stagnant and refined.

Birchfield comments on the structural element saying, "The form of this piece mimics the character of a percussion note such as a cymbal crash. A burst of energy slowly dissipates into a sustained resonance." (Birchfield, 2004) The 'sustained resonance' is the collective behaviour of each phrase. Together the piece lasts approximately eight minutes, sixteen generations total. The final four generations are most often perceived by the listener as identical, with very little change.

Once the macro-structure is clear the performer can begin a realisation. Instrument choice is the first process and is subjective; the sounds may or may not reflect the samples stored in the computer. Based on the structural qualities of the piece, the musical process is strengthened with the use instruments that are similar to the computer's output by providing: 1) shape to the gestures 2) a recall mechanism when phrases become refined 3) support to the characteristics of phrases as they evolve. Without these devices the piece loses its linear progression towards 'sustained resonance'. The instruments themselves are organisms that are part of this evolving musical community and must reflect this by remaining part of the sound summation; otherwise the listener's attention may be focused towards the dissonance of the two constituents and not the detail of the musical structure. As a whole the instruments act as sonic modifiers, they allow the shape of a generation to be heard and in turn provide textural variation and repetition, in tandem with timbral accentuation and saturation.

Together the instruments should render a broad range of frequencies, amplitude, dynamic envelope characteristics, timbral possibilities, and an array of physical gestures. A healthy selection considering each of these traits will facilitate the active role of a modifier. The instruments that I use include: large or small Chinese cymbals, small Chinese hand cymbals, metal bowls, metal shakers,

glass marbles, set of small bells, metal chains, etc. These instruments all contain variability in timbre, volume, frequency and gesture and collectively can be played either with hands, striking, scraping or even dropping on the floor. Mallets or sticks could be used as well, as long as they were able to provide enough sonic malleability and colouration in relation to the score.

Once a conception of sound is obtained then a method for executing the individual gestures in relation to phrases must be defined. As mentioned previously gestures are part of phrases, which are part of sections. Gestures must be identified and shaped to reflect the characterization of the phrase. Initially this is completely random as generations have not yet clearly revealed what these traits may be. However, as generations pass this becomes clarified more directly. For example, a gesture may include high frequencies, very quiet notes with short durations, and thin texture. This sort of characterization might be amplified if the performer remained very still, played very dry colours, with minimal intensity. These gestures are variegated based on the algorithm. They accent the phrase content if they are closely related to the instrumentation. However, the performer is given specific notes by the computer as detailed in the score generation. These notes are part of the community, and at times there may be thirty or forty notes in a given phrase. The performer's task is to select which notes to observe and which to ignore. The selection process is dictated by the originally defined method of executing gestures. In early generations, notes may not conform to the general architecture of the phrase, and that should be accounted for in the performance by remaining separated from the gesture. Later generations though will adapt to the evolving shape of the phrase, and the performer should observe this through instrument choice, duration, timbre and volume.

Mutations within a phrase should be accentuated when possible. These are most obvious within phrases that are locally homogeneous. If all notes assigned to the computer are generally soft and high frequencies a mutation could appear as an extremely loud note played by the computer. However, the performer can also be asked to perform a mutation within a gesture and if at all possible should be observed. Since the mutations are an important part to the virility of the phrase then the performer must accentuate their presence in any way possible. This helps sustain the evolution of the musical community over the course of the piece.

Physical gesture can be an adjunct to these processes. To perform the previous example of a phrase (high frequencies, very quiet, notes of short duration, and thin texture), there is an inherent

motion involved. The same is true for a contrasting phrase, perhaps one of low frequencies, very loud and dense texture. A much larger motion will be expected when the phrase is regenerated. To the audience, these physical gestures are a one-to-one relationship to the sonic characteristics of the phrase. This aspect can be cultivated not only because of its visual stimulation, but by its means of defining the resultant macro-structure of the genetic process. A computer can at no point give this reference. A sound from the computer is not associated with movement. However, a sound from the performer creates an association, and certain phrases render an affinity to specific physical gestures. There is an intrinsic metaphor present within this concept. Genetic processes must occur within a cell, whether it is musical or biological this cell eventually becomes a body. The inherited physical nature of this genetic musical process begins with the cell, and takes the form of human movement to make sound. In turn, the process must speak through a body otherwise it is incomplete and cannot survive.

The physicality of this piece was in part the inspiration to implement 5.1 channel spatialisation. This aspect was only recently added to the algorithm in August 2004. The original version was written for stereo with two microphones amplifying the left and right sides of the performer. The performer's left hand was mapped to the right speaker and the right hand to the left speaker, to give the audience a spatial image of the performer's physical space. This idea was expanded to four channels to give a three dimensional representation of this space to the audience. As the instrument choice saturates the sonic space, this treatment saturates the physical space. Both space and sound become ultimately saturated upon the initial generation. From this saturation a form and a structure can be derived, much like a sculpture is carved from stone. Without a saturated environment the musical process will not be heard.

## Conclusion

Realizing this piece is a developing study in process and how performers use process in interpretation and in real-time environments. The score is very specific while still allowing freedom to shape the fundamental processes. However, it is clear that without specific devices the process of this piece would be weakened.

This piece has been continually developing over the past three years, and I predict it will continue to develop. The piece is essentially a musical community in which a performer is a welcomed part. In the future it would be interesting to attempt a real-time analysis of the acoustic sounds and to integrate them within this system. That would perhaps be the ultimate realization of this

work, thus the performer would not only control the shape and structure of the music but also the behaviour and characterization of the computer processed sounds.

Personally I would like to continue pursuing works of this nature as they promote challenging performance issues. It is my goal to pass new modes of performance to people and new approaches to performers and composers. Special thanks must go to David Birchfield for his vision and creativity.

## References

- Birchfield, David. *Community Art: Resonant Energy* Unpublished, 2002
- Birchfield, David. Web Document <http://ame2.asu.edu/faculty/dab/index.php> [accessed 11.7.04]
- Lewis, George E. "Interacting with Latter-Day Musical Automata." *Contemporary Music Review* 18:3 (1999): 31-39.
- Lewis, George. "Too Many Notes: Computers, Complexity and Culture in Voyager" *Leonardo Music Journal* 10 (2000) 33-39
- Mozart, Wolfgang Amadeus. *Musicalisches Würfelspiele*. Brighton, Mass: Carousel, 1983.