PERFORMING ALGORITHMIC COMPUTER MUSIC: REAL-TIME SCORE INTERPRETATION OF DAVID BIRCHFIELD'S *COMMUNITY ART: RESONANT ENERGY*

Robert Esler
Music Department, University of California, San Diego
resler@ucsd.edu

ABSTRACT

This essay provides a description of the author's real-time score interpretation of David Birchfield's *Community Art: Resonant Energy*. The piece is composed for percussion and computer and uses a genetic algorithm to generate a score and computer playback of pre-recorded percussion samples in real-time. The author discusses the details of the genetic algorithm such as hierarchical form, mating, fitness evaluation, and contour definition. These processes are related to the understanding of the work in regards to the issues of performance execution, instrumentation, musical gesturing, and real-time score analysis. These issues are thoroughly illustrated with examples from the piece. The author offers a specific approach to performing within Birchfield's co-evolutionary musical system.

1. DESCRIPTION OF *COMMUNITY ART: RESONANT ENERGY*

Community Art: Resonant Energy is for solo percussion and computer processing. The computer uses a genetic algorithm to create in real-time a score, and process pre-recorded samples stored in a database. Both stereo, and 5.1 surround sound versions exist. [1] The percussionist is free to choose their instruments and the manner that they are played. There are two different instrumentations for the piece: a) metal instruments that are shaken, struck, scraped, etc. b) shakers and membranes. The description of each version reflects the types of samples that are stored in the database; the performer may choose similar or contrasting instruments.

Once the piece begins it runs continuously for a total of eight minutes. The trajectory of the piece is complex, non-linear, and unpredictable. All processes of the piece are generated through a genetic algorithm. From the algorithm samples are chosen from the database based on their suitability to the features of a note, then processed and played back. Alongside the samples, notes are allocated to the performer expressed in frequency range, amplitude, and duration. When the piece begins samples appear to be chosen at random and collectively create a thick texture of noise. The number of samples played back, or notes played by the performer, in a thirty-second period (length of one generation) can be a maximum of nine hundred and a minimum of forty-eight, all with varying durations, dynamics and processing. The middle stages of the piece display the progressive homogenization of the sample selection and note allocation. In the final stages, as the system evolves, a clear structure emerges that is stable and consistent.

The performer plays along with this system in efforts to shape the evolving structural process based on the information generated in the score (refer to section 3 for a detailed description of the score). Like the samples, the instruments chosen by the performer reflect a wide variety of timbres, frequencies, durations, and dynamics. Overall, interpreting the piece depends largely on how the algorithm changes over time. Each realization of the piece is different. The algorithm creates a new structure and sound environment at each performance.

2. GENETIC ALGORITHM

As a composer Birchfield is interested in naturally occurring time-cycles and the interaction of multiple, independent cycles in different time scales. [3] From observations of these cycles Birchfield has created a generative system that defines musical surfaces from which music can be constructed. Using this system Birchfield intends to achieve two goals: 1) slow evolving, large-scale structures that are balanced by variegated local activity and energy. 2) musical processes that interact through time coexisting independently engendering specific musical features. [3] With this motivation Birchfield has created an automated compositional model of a co-evolutionary genetic algorithm.

2.1. Formal Structure

The algorithm uses a hierarchical structure to define large-scale characteristics of form and small-scale characteristics of density. There are five tiers to the hierarchy all which nest the following tiers: 1) meta-sections – approximately 1-4 minutes in length 2) sections – 30 seconds in length 3) phrases – 7.5 seconds, 4 per section 4) gestures – between 3-9 per phrase 5) notes – 4-25 per gesture. From this description the note is the basis of the structure, gestures supplement the aggregation of notes and provide a textural map of the evolving musical material. Each of these tiers is a population. Each population mates within themselves (i.e. sections mate with sections) and share a feature space. A feature space means musical attributes such as amplitude, density, bandwidth, etc. are common among each population. In addition, higher-level tiers define the number of lower-level components (components refers to tiers). For example, phrases determine the number of gestures, which in turn define the number of notes. In *Community Art: Resonant Energy* sections and phrases are fixed (16 meta-sections, 4 sections and four phrases each). Through this architecture, each population evolves concurrently but still retains an implicit relationship amongst the hierarchy.

2.2. Mating Process

As members of each population die, children are born and take their place. Children are hybrids of their parents. Parents are chosen randomly. One of the parents will carry a higher-level fitness function for certain traits (the fitness function is defined in the next section. Traits may include amplitude, duration, density, frequency, etc. The most successful traits will be passed to the child along with the relative fitness level of these traits. This element provides a continually improving population with diverse musical attributes.

Mutations occur when genetic traits are errantly copied. They help insure the diversity of the system. If a child is born with a mutated gene then it could persist until the next generation, ultimately changing the proceeding population. Each mutation is subject to a probability, which is dictated by a random number generator. If the random number is below the probability then a new value for the mutated feature is chosen at random. For example, if a note is set at 100Hz but subject to mutation, then it will be assigned a new value within a specified frequency range.

2.3. Fitness Evaluation

The fitness function determines the success of a musical organism and determines whether particular traits survive through the next generation. Fitness is determined by environmental constraints, which are defined respectively by the previous tier of the hierarchy. Gestures define the environmental constraints for notes, and phrases respond the same to gestures.

Each child is given a primary value, which they received from their parents Moreover, Parents provide a bandwidth within which the primary value can be located before the fitness level will decrease. The function is defined as:

*Absolute value((Parent bandwidth/2.0) – (Component value – Parent value)) * scalar value + small value of absolute distance from primary value, irregardless of bandwidth* [2]

As the primary value of a feature moves away from the parent's bandwidth, the fitness level of the feature will decrease. Provided with environmental constraints, there is an acceptable range that a feature may exist, however, the further the value of the feature is from its ideal bandwidth, the less likely it will survive. In addition, fitness is judged through peer evaluation. If a feature is placed within the lower range of its peer group it is has a low chance of survival. This means that if four out of five gestures exceed their parents' bandwidth they will not all die, leaving only one gesture to sustain the population. Rather the relative fitness of their peers determines their individual outcome.

2.4. Contouring

Aside from a primary value, features are assigned contour definitions. Each feature of a component carries an independent contour. Contouring provides a direct way of shaping musical features, and insures that particular features will not remain static through time. Contours are scaled to fit within a feature's bandwidth provided by the parent. For example, a phrase may include nine gestures that occur close together in time for 1/3 of its duration and decrease in time the remaining 2/3. However, phrases can only define contours of lower-level components and cannot directly affect the microscopic details of those components. Phrases cannot determine the contour of notes. In *Community Art: Resonant Energy* contours are limited to frequency range, amplitude, density, duration, and pitch clarity (filtering). This is due to the concrete nature of the samples and their immutability.

3. SCORE INTERPRETATION

While learning this piece there was always a degree of uncertainty that what I was doing did not relate to the computer. After many hours of work and reflection, I decide to mine the processes of the algorithm in efforts to find an interpretive context. This context became the aural implications of the musical system. The piece reflects a certain affinity. Although it is always different, it still retains similar traits from realization to realization. It became clear that the performer's role was to reveal this evolution to the audience for them to hear the progression of Birchfield's musical system.

3.1. Description of Score

The window for the performer into this system is a graphical score that is generated in real-time on a computer screen. The score displays thirty-second sections before a new screen is drawn and a new generation is simulated. Figure 1 shows an example of one thirty-second meta-section. The score is a grid that represents time on the x-axis and frequency on the y-axis. Phrases are marked by solid white lines that extend vertically down the entire length of the screen. Gestures appear as groups of notes and notes appear as colored lines that run horizontal and are of varying length. The thickness (vertically) of a line determines the amplitude of a note, the length shows its duration and the color displays its function within the system. A note can be an unprocessed sample, a processed sample (filtered for greater pitch clarity), or a note executed by the performer. A cursor moves horizontally across the screen indicating time.

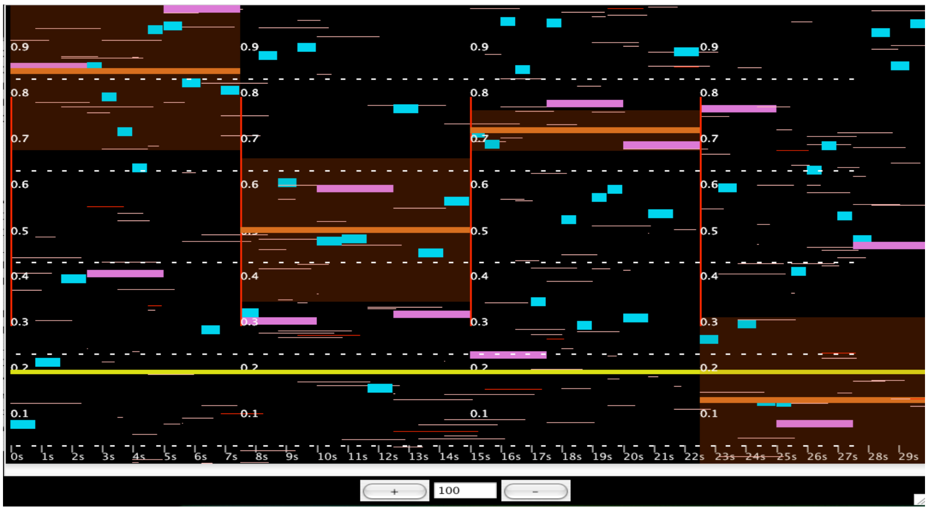

Figure 1. Score generation. This indicates one meta-section of the piece.

3.2 Performance execution

Even with a clear interpretive context, the real-time interpretation of the system through the score is extremely complex. Several issues arose regarding the execution of the piece: 1) What kind of instruments should be used to execute notes and how should they be played? 2) How should the notes in the score be realized in regards to the evolving structure? 3) How can the performer accurately assess the future progression of the system over the course of the work? 4) How can one perform with this system without appearing removed from the evolution of the piece? These issues are explained in the following segments.

3.3 Instrumentation

After many weeks of observing simulations of the piece I decided to choose instruments that could function several ways: 1) sound similar to the recorded samples 2) capable of being played in multiple ways creating many possibilities of timbre, dynamics, and durations 3) relatively small to enable rapid exchange between each instrument or allow many to be played at one time. With these functions in mind I chose a large Chinese cymbal as the fundamental instrument (this is in regard to the metal instrument version). This instrument afforded enough sonic malleability to provide low to middle frequencies, long to short durations and loud to soft dynamics. All the other instruments were centered around this instrument. From there I chose two small hand cymbals, which could be played separately or on the Chinese cymbal itself. This allowed easy exchanges between low to high frequencies and loud to soft dynamics. Once the idea of playing a variety of instruments on the Chinese cymbal was worth pursing I chose more items that could be played both on the Chinese cymbal or independently. These included: varying sizes of chains, strings of bells, small cow bells, large bells, and scrap metal. In addition, I chose several shakers that sound metallic and could reflect the contours and timbral similarity of the recorded samples. Examples include: tin can with metal tacks, metal bowl with small marbles or buckshot, glass container with small nails, and ankle bells. All, or combinations, of these instruments gave me enough freedom to interpret the score and provide an interesting and interactive sound environment.

3.4. Seeing the Evolution

It became obvious that trying to perform the score exactly as written was impossible. Instead different skills were needed. One skill that I developed in working on the piece was analyzing the evolution of each section visually during performance. As mentioned before, notes have different colors and congregate in gestures. As the system evolves gestures become more stable, and tend to change less over subsequent generations. It was important to plan ahead for upcoming phrases. If there were five gestures in the fourth phrase, each of which spanned a small frequency range, long durations, and soft dynamic, then I would prepare instruments that would function within this phrase.

Seeing the structure also helped in connecting the contours of the system with the contours of my instruments. Although the score does reveal which samples it will play, the system tends to develop communities of sounds within phrases that include several combinations of timbres, amplitudes, durations, etc. In many ways this is a didactic process. Once this type of behavior is learned, the performer can interact simultaneously with the computer. These types of musical gestures by the performer provide a deeper correlation to the system.

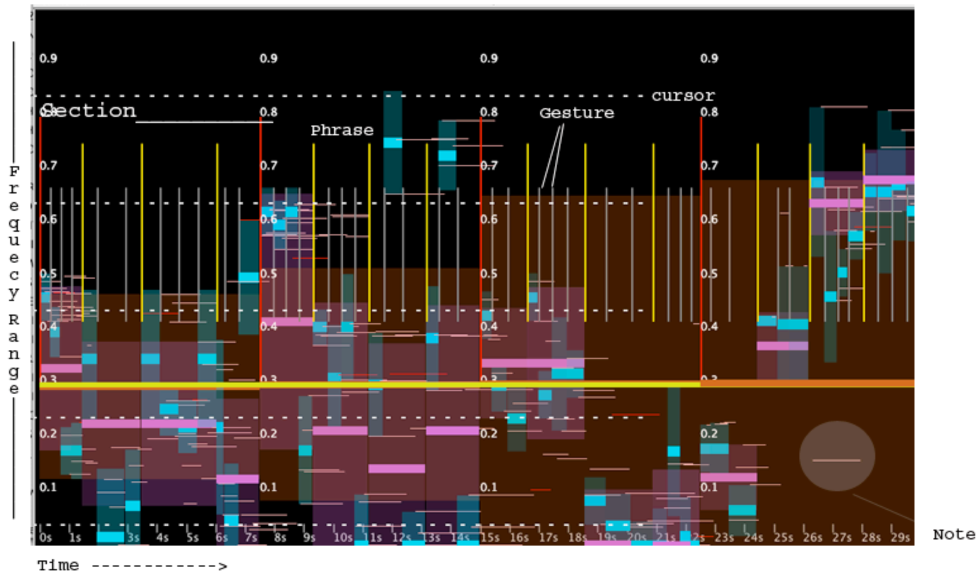

Figure 2. Mock-up of a score generation. One section in length. This example clarifies how certain aspects of the score appear. Large blocks indicate bandwidth of a gesture, phrase, or section..

3.5. Hearing the evolution

In tandem with seeing the structure, hearing the progression of the system was the most valuable asset. As mating changes the system it becomes more apparent which musical features are surviving and which have been selected to die. Hearing the evolution informed such decisions as instrument choice, instrumental color, instrumental contour, and the physical gestures needed to reflect the musical activity.

3.6. Imitating, accentuating, coloring, and shaping the algorithmic processes

The evolution of the piece is the most salient feature of the musical system and the most impressionable to the listener. However, Birchfield is also interested in the local activity and the intricate detail that is created through the algorithmic processes of the piece. There are several types of musical gestures that reveal the large-scale evolution and also reinforce the small-scale details of mutations, mating, and contouring. These include: imitation, accentuation, coloration and shaping.

Imitation creates a one-to-one correspondence between activity of the computer and activity of the performer. Whereas accentuation supports a gesture or phrase, but remains independent. Coloration provides a timbral scaffolding to individual notes and gestures and shaping remits the changes in duration and amplitude. These gestures intuitively sustain the processes of mating, mutation, and contouring. Imitation informs mutation by connection and repetition. Accentuation provides individuality amongst populations while coloration binds populations. Shaping defines the contouring of musical features and associates dynamics with duration. Most importantly, these relationships between the computer algorithm and the performance execution are not meant to truly define the micro-details of the genetic process. They are rather subsidiary elements to the evolution of the system.

4. CONCLUSION

The most difficult aspect of the piece is that every detail must be synthesized in real-time. There is no way of knowing how the system will behave and no time to plan how you will execute the piece precisely. Instead, the performer is left with developing a concrete interpretation and understanding of the processes involved with this work. Birchfield has created a performance system that is essentially boundless with musical possibilities.

This piece has been an exceptional challenge and is continuing to evolve as a musical project. Many thanks to David Birchfield for his vision and innovation.

5. REFERENCES

[1] Birchfield, David. *Community Art: Resonant Energy.* Published by composer. 2002.

[2] Birchfield, David. *Evolving Intelligent Musical Materials*. Doctoral Dissertation, Columbia University, 2003.

[3] Birchfield, David. *Genetic Algorithm for the Evolution of Feature Trajectories in Time-Dependent Arts*. Proceedings 6th International Conference on Generative Art, Milan, Italy, 2003.