A more perfect union: Composition with audience-controlled smartphone speaker array and evolutionary computer music

Tate Carson
Experimental Music and Digital Media
Louisiana State University
tcarso2@lsu.edu

ABSTRACT

A more perfect union incorporates an audience-controlled smartphone speaker array with evolutionary computer music. A genetic algorithm drives the work and the performance practice that the audience follows.

1. INTRODUCTION

A more perfect union incorporates an audience-controlled smartphone speaker array with evolutionary computer music. It fits into a long list of evolutionary computer music works[9], but differs from these works due to the addition of real-time audience interaction facilitated by a distributed smartphone speaker array. Using smartphones, the audience controls the outcome of the piece. Instead of passively watching a performance on a stage, the audience becomes the performance by taking an active role in the work's creation. The intention of the work is to produce a platform for audience interaction that incorporates audience preference, not technique, to arrive at an unexpected musical result.

2. MOTIVATION

The audience takes part not solely in the diffusion of the work, but also in its creation. Participation in the work is solely dependent upon audience members' musical appreciation. A genetic algorithm takes participants' decisions and guides the melodic and rhythmic results. The hope is that given an ideal amount of time and willing participants, the composition will develop into something that all participants appreciate. The piece was designed with no preset aesthetic goals for the resulting composition, only that the participants agree on it. This end result may not be possible, however the work is intended to be predominantly a conceptual experiment where the choice of the audience is more important than any preconceived musical aesthetic.

The audience's role is easy to understand and participants need not comprehend the work's complex underlying algorithm. Because participants have a role in creating the work by expressing their preference, they have an opportunity to feel a greater connection to the work because they are driving its composition[3]. If a participant can identify an instance of one of his choices recreated in another part of the performance space, he will feel he is an active part of the experience. The challenge is to enable a participant to feel as though he is having the same level of engagement that one feels as a trained musician, entirely in command of his own choices.

A more perfect union is an experiment in distributing control of a composition to an unfamiliar group of people. The work lies within the experimental music aesthetic of algorithmic network music passed down from the early network music group the Hub[2]. The Hub was interested in the emergent qualities of algorithmic processes that were unpredictable. *A more perfect union* relates to the Hub's aesthetic, but instead of only the band members being networked together, the entire audience is networked together and invited to share control of the composition.

3. RELATED WORK

Daniel Shiffman, through his video tutorial series *Coding Train*, is the main inspiration for *A more perfect union*, which uses as its starting point Shiffman's example on interactive selection[1]. In his video, Shiffman cites Karl Sims' 1997 installation *Galapagos*[9] as his inspiration. Sims' work consisted of several video screens, each displaying a different virtual organism. The installation allowed spectators to take part in evolving virtual organisms by choosing the amount of time they spent in front of one video screen versus another. The longer a viewer stood in front of a screen, the more he increased the fitness of that virtual organism, and made it more likely for the organism to pass its traits onto the next generation.

Nexus[1] was one of the first distributed performance systems that used a browser-based interface; it solved the problem of cross-platform distribution. Unlike similar systems, it incorporated a user interface library, NexusUI, to aid in the rapid building and prototyping of user interfaces.

Another significant work in the history of distributed music systems is *Fields*[8]. The work is both a performance system and composition, developed with the needs of the composition in mind. An important aspect of Fields is its insistence on "using what's available." It does not require the user to download software, but instead uses a web browser for sound production, creating a more inclusive environment for audience interaction. *A more perfect union* uses Rhi-

[1] https://www.youtube.com/watch?v=Zy_obitkyOE

Web Audio Conference WAC-2018, September 19–21, 2018, Berlin, Germany.

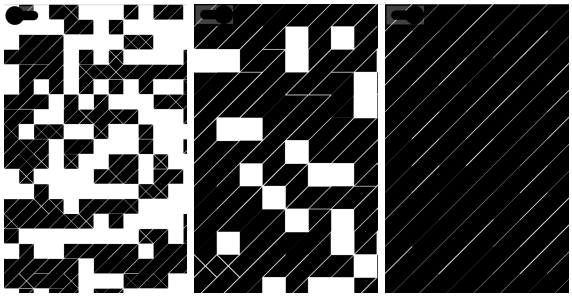

Figure 1: Three versions of the audience webpage

zome[7], the web server developed for *Fields*.

Other works related to *A more perfect union*, but less influential, include *Crowd in C[loud]*[4], by Antonio Deusany de Carvalho Junior and Sang Won Lee. This work allowed the audience to participate in the performance through smartphones using cloud-based technology. Also, *echobo*[5] by Sang Won Lee and Jason Freeman used audience participation practices, but required the audience to download an application.

4. PERFORMANCE

For *A more perfect union*, a projection screen guides audience members to a webpage on their smartphones where they see a toggle switch used to turn the sound on and off and a visualization of the genetic algorithm. The webpage displays program notes and instructions for participating in the work. The audience is encouraged to move around the performance space to experience the work from various vantage points.

Each audience member enacts the fitness function of the genetic algorithm by listening to a sound he likes for a desired amount of time. The longer he listens to a sound, the higher the sound's fitness value and the more likely it is to pass on its traits. The fittest sounds pass on their genes, combine with other fit sounds, and create new sounds with the traits of the parent sounds[6]. Over the duration of the piece, as audience members listen to sounds they like and skip ones they don't, they should ultimately hear an overall composition develop through their shared efforts.

5. TECHNICAL DESIGN

The technical design of *A more perfect union* consists of three parts: the controller webpage, the audience webpage and genetic algorithm, and the server that connects the audience members' devices.

5.1 Webpages

The controller webpage has buttons for starting and stopping the work. There is a box to type in a duration for the performance, which allows the controller to vary the running time for different performances.

When an audience member reaches the audience webpage, Tone.js[2], a JavaScript framework for working with the Web Audio API, downloads and prepares samples. The audience members control the starting and stopping of sounds with a toggle switch from the interface library NexusUI[3]. Each sound is accompanied by a visualization (see Figure 1) of the melody currently being heard; p5.js[4], a drawing library, displays the visualization.

5.2 Server

When an audience member stops a melody, the fitness function of that melody is transmitted to all other audience members using Rhizome. Amazon EC2[5], a cloud computing service, hosts the webpages and server. Cloud computing is a cheaper alternative to traditional hosting solutions because the hosting cost is based on time used during a performance.

5.3 Genetic algorithm

The genetic algorithm used in *A more perfect union* is based on evolution as Charles Darwin understood it, wherein a phenotype is the perceivable trait of a gene and a genotype is the actual representation of that genetic material. The major components of the genetic algorithm are[6]:

Genetic representation (genotype): An array of numbers that represents the genes.

Evaluation (fitness) function: The means by which fitness is evaluated; operates on the phenotype. Interactive selection is a type of fitness function that uses human judgment.

Population structure: A data structure containing the genotypes and their corresponding fitness values.

Selection method: The fittest genes of a population have a higher probability of passing on their traits.

Genetic operators: After the fittest parents are selected, two common operators are applied– crossover and mutation.

Replacement scheme: How new genetic material is introduced into the population.

In *A more perfect union*, the genetic representation is a random array of numbers between 0 and 1, which gives the audience a neutral starting point for evaluating melodies. The fitness function is interactive fitness. The audience members decide the fitness of each sound by toggling a switch to turn sounds on and off. The time each audience member spends listening to a given melody determines its fitness score; the longer they listen to that melody, the greater the fitness. Time listened is used as a driver of the fitness function because it is assumed that the longer someone listens to a given melody before moving onto the next one, the more they liked that melody and the relationship it had to the other ones happening at the same time. This method intends to determine the group's idea of what a the composition should ultimately sound like. Next, the population structure—a fixed array—is filled with the genes.

When an audience member turns off a sound, the selection method determines which traits are passed on (See Figure 2). Sounds that have high fitness values from other audience members create a mating pool of sound choices. Sounds with a higher fitness values are pushed to the mating pool more times, making them more likely to be selected and prioritizing positive choices over negative ones. Two parents are then picked from the mating pool at random.

[2]https://tonejs.github.io/

[3]https://nexus-js.github.io/ui/api/#intro

[4]https://p5js.org/

[5]https://aws.amazon.com/ec2/

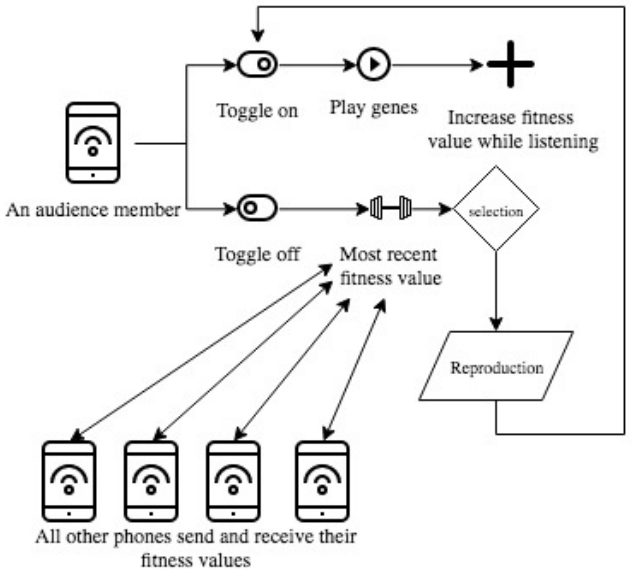

Figure 2: Fitness value data flow

The algorithm uses crossover and mutation as the genetic operators. The crossover function creates a child from the two parents; this child is a pseudo-random combination of both parents. Next, the algorithm applies mutation to the genes. It picks a random gene and inserts a 0 or 1 into a random index of the array; this makes certain that the gene pool will not converge too quickly to one extreme or another.

The phenotype is the genotype mapped onto musical parameters so that the audience can perceive the algorithm. Musical parameters used include tempo; sample choice of metal, bell, marimba, pots and pans; sample tuning, restricted to a predetermined scale; and sample rhythm.

6. EVALUATION

Thus far, the work has been performed four times—twice in traditional performance halls and twice in art gallery settings. The performance results were not quite what I had envisioned when I started the project. One reason is that the work was never performed for more than about ten to fifteen minutes each time, which wasn't enough time for a coherent composition to develop. I believe that it could take up to an hour for a cohesive composition to emerge, which is the overall goal of the piece. An hour is a lot to ask of an audiences' attention, and a near impossibility for programming on most concerts. Within ten to fifteen minutes there is little possibility that an audience member will hear his choices reflected in a meaningful way. Also, I do not think that the interactions the audiences have had with the smartphone interface was as conducive as it should be to encourage their participation. Enlightening the participants about how their selection preferences are contributing to the composition, but not overwhelming them with technical details, is a major challenge.

Although the performances did not turn out exactly how I had envisioned, I was pleased with how they sounded. I have received construction feedback from participants that indicates they enjoyed participating in the work.

7. FUTURE WORK

Optimizing the genetic algorithm in a way that is more specific to a musical composition would make the audience members' actions more noticeable. The selection, crossover, and mutation functions are currently somewhat generic, but could be made less random and more specific to a musical domain.

In addition, instead of presenting *A more perfect union* as a traditional musical performance on a stage in a concert hall, with a delineated beginning and ending, it would be interesting to present the piece as an installation. An instillation setup would allow the piece to last for a longer amount of time, letting the algorithm develop. When people walk in and out of the installation space, the sound would change over a longer period and be more unpredictable.

Finally, developing a system to visualize the inner workings of the algorithm would allow the audience to understand how the composition is being created. Thus far, the composition has been presented with little emphasis on the scientific aspect of the genetic algorithm, and rather more as a metaphor for voting on the outcome of a collective musical composition. A more detailed visualization would have the potential to reveal additional information about the algorithm without foregrounding the technical nature of the work.

8. CONCLUSION

A more perfect union explores a new method of audience interaction. The technical interface enables a group of strangers to collaboratively create a musical composition by individually expressing their musical appreciation. The current structure allows for further experimentation and improvement with the goal of delivering to audiences an engaging participatory musical experience in which they feel ownership.

9. REFERENCES

[1] J. Allison, Y. Oh, and B. Taylor. NEXUS: Collaborative Performance for the Masses, Handling Instrument Interface Distribution Through the Web. In *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME)*, pages 1–6, 2013.

[2] S. Gresham-Lancaster. The Aesthetics and History of the Hub. *Leonard Music Journal*, 8:39–44, 1998.

[3] D. Kim-Boyle. Network Musics-Play, Engagement and the Democratization of Performance. In *Proceedings of the 2008 Conference on New Interfaces for Musical Expression*, pages 3–8, Genova, Italy, 2008.

[4] S. W. Lee, A. Deusany de Carvalho Junior, and G. Essl. Crowd in C[loud] : Audience Participation Music with Online Dating Metaphor using Cloud Service. Atlanta, USA, 2016.

[5] S. W. Lee and J. Freeman. echobo: A mobile music instrument designed for audience to play. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 450–455, Daejeon, Korea, 2013.

[6] E. R. Miranda and J. A. Biles, editors. *Evolutionary Computer Music*. Springer-Verlag London Limited, 2007.

[7] S. Piquemal. rhizome, Feb. 2017.

[8] T. Shaw, S. Piquemal, and J. Bowers. Fields: An Exploration into the use of Mobile Devices as a Medium for Sound Diffusion. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 281–284, Baton Rouge, Louisiana, 2015.

[9] K. Sims. Interactive Evolution of Dynamical Systems. In *Proceedings of the 1st European Conference on Artificial Life*, pages 171–178, 1992.