

Form and Transience - Generative Music composition in practice

John Eacott B.A. (hons) Music

University of Westminster, London

eacottj@wmin.ac.uk / john@informal.org

“Music teaches us all that is beautiful in this world. Let us not hamper it with a machine that tells the story day by day, without variation, without soul, barren of the joy, the passion, the ardor, that is the inheritance of man alone. (John Phillip

Sousa 1906)¹

My route to becoming a composer of generative music could be described as conservative. I began my musical life as a trumpeter, playing mostly jazz and a fairly early age I learnt how to harmonise and arrange music using manuscript paper. These two skills have until recently formed the backbone of my career. Two years ago I switched, practically overnight, to creating generative music.

At first I eulogised about the extraordinariness and newness of the medium (comparing its significance to milestones like the development of notation around the 15th century, equal temperament in the 18th and the development of recorded media and electronic instruments in the 20th). Since that naive rush, my position has changed. There is nothing particularly extraordinary about using generative approaches to creating music. In fact it is a natural progression and one that simply uses the potential of new musical instruments, particularly ones that utilise digital technology.

Instruments

Musical instruments define music as much as the players and composers do, and historically, the development and evolution of musical instruments has always impacted on the creative opportunities of musicians.

Computer based musical instruments have already dominated the production of pop and some other areas of music for the last twenty years. Although this has been significant, computers have been used predominantly for relatively mundane tasks as compact versions of sequencers, recorders and synthesizers. During this period major amounts of research and investment not only by computer and software developers but also communications companies is focussed on Artificial Intelligence. Although still a long way from creating programmes that really compare with human thought, AI can offer much to the development of musical instruments.

In the field of generative music the distinction between instrument and the composition is not always clear and I would like to illustrate this with reference to two well known examples.

In the case of a system like KOAN, although the documentation that accompanies it encourages you to interact with the parameters to create your own unique settings and hence, create a new piece of music, what is mostly achieved (in my experience) could be described as a remix of pieces preprogrammed by the development team - not that that is a bad thing.

Jem Finer's 'Longplayer'² on the other hand, uses no interactivity, no random number generation and is completely predetermined (and Jem wouldn't call it a generative piece anyway!) but it shares some of the characteristics of a generative system by nature of its construction as computer algorithms which realise the 1000 years of music for him from source material of only 20 minutes.

Whether it be a wind chime, computer generated, algorithmic or other generative system, what we have essentially is an instrument which, like all musical instruments, has a greater or less restricted set of performance possibilities - a sort of 'embedded' musical content. The composition of a generative work therefore is a process of defining musical content through the instrument upon which it is realised.

Composing

The challenge facing a composer of generative work is a need to quantify issues about the nature of beauty itself in order to construct a system that realises this for you. In order to do this I looked for a model of beauty that I could study. The one I came up with is a large tree that stands in a garden near to my flat in Stockwell, South London!

Each morning when I wake I can see this tree from my bedroom. Every day its form is the same although there are changes to the detail. The backdrop of sky is always different, varying from clear blue to various types of clouds moving and reforming in various ways. The sun is normally behind the tree in the morning depending on when I wake, the time of year and the level of cloud. The tree itself changes appearance throughout the year by growing leaves which change colour in the autumn and fall away gradually leaving only the bare branches. The branches themselves can be still or swaying depending of the wind and the action of birds and other creatures. Recently a violent storm in the night left us wondering if the tree would still be standing in the morning. This potential for catastrophic change works alongside the regularity and consistency of the form itself to make us value it further.

Every time I look at the tree I find the consistency of the form reassuring and gratifying while the detail of the image is unique. Its continual growth and evolution mean that any beauty the view possesses is transient. I wish to draw parallels between this and music .

One of the features of playing the trumpet is that you never know exactly what sound will come out when you blow it! As hard as you may try to develop a perfect 'embouchure' and achieve consistent production, the reality is that each note you play is to a greater or lesser degree a unique event.

Similarly, notated music is not an exact art. The way that a note, phrase or entire work is interpreted and played varies according to many factors including the amount and level of interpretive parameters granted by the composer. In this way, score notation could be regarded as an *instrument* as well as a composition. This is by virtue of the fact that notation is a means by which a musician can realise an idea, in just the same way that a guitar or piano is. While the composition itself is essentially an idea (or complex group of ideas), the manuscript it is written on is a means of transmitting and communicating the idea to others.

The manuscript becomes part of the instrument of performance along with the conductor, musicians and their associated instruments. At all levels, from the production of a single note on a single instrument to the whole event, it is a work's form along with its potential for uniqueness, transience and catastrophe that makes live music a powerful experience.

In contrast to the fluidity and variation of live performances, the last hundred years have been dominated culturally by recorded media such as the gramophone record, film, videotape and CD. The degree to which these forms have dominated our cultural awareness is such that for some people linear recordings are regarded as the primary form of musical expression. Although the recordings, films and videos of the 20th century constitute a vast and important body of work and ideas we are possibly at a point when their dominance is evolving in favour of non-linear arts.

In my own composition, whether generative or conventional it is important to mix form and transience. Composing on manuscript paper is clearly about form and uses the notion of transience through the performance of musicians. In a generative work both categories are written in to the algorithms. I use the terms 'fixed' and 'fluid' content which differentiate between elements of form, sound, pitch, rhythm or whatever.

'Fixed' elements will certainly feature in the piece while 'fluid' elements may be subject to change in some way or possibly not occur at all. In order to illustrate how I have applied this analysis in the creation my own music I shall refer to two recent and contrasting works.

The Street

The street is an interactive / generative music installation situated at the University of Westminster's Harrow Campus. The piece takes its name from the central thoroughfare which links the various buildings of the campus and through which most users of the university must pass.

The space is a striking design in brick, glass, steel and teflon with a modern, functional and uncluttered aesthetic appeal. The idea for the piece designed by myself and computer scientist Ross Clement was for a generative music piece that would respond to the movements of

people using the space by which is detected by means of ultrasound movement sensors. The piece should sonically enhance the experience of the users and appeal (rather than annoy!) to most users at a basic level. The installation should run continuously for 2 weeks.

Creative vision

A central issues in creating a permanent music installation in a public space is finding something that will enhance the space and not simply pollute. Given the wide range of musical tastes and cultural diversity of people using the space this was a significant challenge.

The way in which the music should respond to the users was also critical. There is a tension between the control of content and form on one hand and interactivity which affects that content on the other. Somehow the music must work appropriately with any degree and type of interaction. This *can* be achieved by minimising the affects of interaction and hence retaining near total control over content. This is incompatible however with the desire to present a level of interactivity clearly perceptible to the user.

In very general terms it was decided that the piece should have a contemporary sound which reflected the styling of the building. In terms of interactivity we decided that the clearest and most obvious relationship between users and music would be best. In other words, the more movement in the space, the greater the level and density of the music. If there was no movement the music would stop.

As a large, simple clock hangs in a central, dominant position in the space we decided to include a functional time announcement which would also give the piece a clear metric structure.

The piece grew fairly logically from these decisions.

- The musical content would be synthesised from basic sound sources like sine, sawtooth and square waves and some white noise which reflected the elemental / modern materials used in the building.
- The music would be organised in 8 layers of sound which would be ‘uncovered’ according to an algorithm which read the density of movement in the space.

- A verbal time announcement every 15 minutes would begin and end each section of music helping to promote clarity of structure as opposed to endlessness. At this point new melodic, harmonic, rhythmic and textural content would be generated which would apply to the (slightly less than) 15 minute musical section to come.
- The piece would be constructed in SuperCollider³ sound synthesis environment as this is the most suitable environment for defining all our own controlling algorithms and sound synthesis.

(A full list of the programme algorithms is available in the appendix)

jnr tv - Dance Track

Since May 1999 I have been performing live under the name 'jnr tv'. The idea is to play club / dance style music like house and garage in club venues. So far, because of my changing ideas and developing technological skills each performance has had very different content.

The central idea of the project is to do live performances of 'dance music' that involve interactive and generative content. I use my trumpet and voice during the performance. Rather than cluttering the with other devices I have chosen to use only my trumpet as a controlling device.

On 29th Nov I was able to perform a version closer to my original vision. The performance consisted of a roughly 15minute set divided into 5 sections. Each section included elements that were interactive / generative or both. Only the final section 'Dance Track' was clearly in dance music style.

One of the features of dance music that I most definitely wanted to preserve are the elements of form that permit DJ's to lead dance records into one another. A normal dance track is succinct, lasting up to 5 minutes with a sequence of builds and breakdowns which provide the dynamic contrasts people on the dancefloor rely on.

Because I believe there is nothing 'special' about generative processes, it is important to me that my music works in the same way that any piece of music does. In other words I am not

prepared to make excuses that the piece 'may not always be good' because it is generative. That is not the way forward. To make successful generative music is to write algorithms that satisfy your compositional needs. In the case of Dance Track it was important that the material had a reguarity, repetition and overall identity from one playing to the next. I see no reason why this could not lead to mainstream use - a generative 'hit' single or TV theme! . In the case of my 'jnrtv' performance, (except for the clue in the title of my act!) I make no announcement of the fact that the music is generative. For me the music should work and be appreciated as music and that is all that matters.

In 'Dance Track' fixed elements of content include the song structure and most elements of the rhythm. Melodic / harmonic elements and some of the sound / textural and spacial elements in this piece may be fluid. I believe this permits an appropriate balance of form and transience for this type of work.

(See appendix for list of programme algorithms)

Conclusion

The compositions described above are working pieces. They demonstrate the incorporation of generative approaches into musical composition alongside conventional methods. I reject the notion that a generative music system 'composes the music for you' and simply embrace generative approaches as a natural addition to the whole armoury of composing methods. Although it is my aim to make generative music which works on any terms, I believe our aesthetic of music is also evolving from one that purports to value specific content to one that embraces a notion of composing as the process of formulating musical ideas with varying degrees levels of form and transience.

Appendix

Below is a brief summary of the algorithms used in each of these pieces. Although some detail is left out the basis of the algorithm including any decision making is complete.

the street (- installation)

Algorithms;

Interactive control

Each time an impulse from the ultrasound sensors (indicating movement in the space) is received by the computer a density variable 'd' is incremented. After a constant period (which was 'tunable' to achieve an appropriate response) d decrements. d is limited in range between 1 and 8.

The value of d at any moment affects the content of the music heard by 'un-gating' some of the 8 sound functions according to the table below.

<i>Value of d</i>	<i>sound functions heard</i>
1	1
2	1, 2
3	1, 2, 3
4	1, 2, 3, 4
5	2, 3, 4, 5
6	3, 4, 5, 6
7	4, 5, 6, 7
8	5, 6, 7, 8

You will see that as d increases up to a value of 4 new elements are introduced in addition to earlier ones. From 5 to 8 new elements are introduced and earlier elements are removed. This is to allow the musical texture to grow without becoming overly cluttered.

Metric structure

Every 15 minutes a time announcement is generated using an array of pre-recorded voice samples. At this time the following composition parameters are generated which will apply to the next (14 and a half minute) section of music.

- Tempo - range from 180 to 240 bpm (this is the tempo of the smallest subdivision and so it isn't that fast)
- Time signature - range from 3 to 8 beats in a bar.
- Beat subdivision⁴ - range from 2 to 6 subdivisions per beat.
- Harmonic structure - chooses 4 harmonic bass notes, each taken from the F pentatonic scale. Chord 1, 2, 3, 4.
- Melody - an array of up to 48 notes chosen from F pentatonic scale over 2 octaves. 50% of the notes (on average) are substituted with rests to provide rhythmic variation.
- Melody timbres - the synth voice used to play the melody is generated from within limits dictated by taste.

Sound functions

1. Windchimes - Using randomised pitched bell tones. Responds immediately upon trigger. Envelope: 5 seconds.

2. Texture / wind ambience - A mixture of atmospheric sounds. Responds immediately but longer envelope: 12 seconds.

3. Single synth pad - A single synth note responding to the harmonic structure as follows: A fifth above Chord root 1 held for 4 bars. A fifth above Chord root 3 held for 4 bars. Soft 'analogue' synth pad sound, slowly evolving envelope with 'phased' effect.

4. Synth Pad in Fifths - A pair of notes built on the root and fifth of each chord of the Harmonic structure. Duration 2 bars. Soft pad sound.

5. Percussion - A 'wood block' like tone to mark the first beat of each bar and a 'hi hat' sound to mark each subdivision

6. Off beat chords - Marimba-like tonal percussive sound playing on even numbered beats i.e. 2, 4, 6 and 8 (depending on Time Signature). Chords consist of root and fifth of Harmonic Structure chords.

7. Melody - The 48 notes of the melody array (including rests) played in sequence depending on the number of subdivisions required by the Time Signature and Beat Subdivision. i.e. if Time Signature is 5 and Subdivision is 3 only 15 notes of the melody will be played. Every even numbered bar the melody sequence is displaced 1 subdivision late. This is a simple (but effective) way of creating melodic / rhythmic interest while preserving economy of material.

8. Big Beat - Finally to signal 'maximum energy' of the street (and its occupants) a big 'bass drum' playing on each beat and a 'snare drum' playing on even numbered beats.

jnr tv - Dance Track

Here is a breakdown of the main features of the construction in 'phrases' of 4 bars each. (i.e 8 phrases = 32 bars)

Song structure: Intro 8, Chorus 1 8, Breakdown + build 8, Outro chorus 8 - fixed.

Drums: Bass Drum (3 patterns), Snare Drum (2 patterns), HiHats (2 patterns). Patterns arranged according to song structure - fixed.

Synth1: (2 patterns) arranged according to song structure. Pattern content and panning - fixed, Effects / filtering - fluid.

BassSynth: (2 patterns) arranged according to song structure. Content - fixed. Effects - fluid.

Crash: (An effect often at the start of a section - an electronic equivalent to a crash cymbal) arranged according to song structure - fixed. Filtering and delays - fluid.

TrumpetSynth: An algorithm that reads the pitch and velocity of notes played on the trumpet and plays as a synthesised sound. Repeats on a 4 bar (1 phrase) cycle. Subject to an overall envelope linked to the song structure. - fluid

SynthStream: A pattern of synthesized arpeggios. The actual pattern of notes combines fixed and fluid content. Effects and filters are fluid. Subject to an overall envelope linked to song structure.

References

¹ John Phillip Sousa in C. Roads; “The Computer Music Tutorial” MIT Press 1996 . p677

² Jem Finer / Art Angel “Longplayer” www.longplayer.org

³ SuperCollider; Sound synthesis environment. www.audiosynth.com

⁴ Using the ‘hierarchical structure’ of meter described by Longuet-Higgins and Lee 1984 in Rowe p152