Thoughts on Composing with Algorithms

Lori Speigler

Background

Throughout history there has been an evolution of means to define and record what we might call “music source” (instructions which when followed produce audible music). Common music notation has been for a few hundred years a successful and useful multidimensional representation of actions that players can do to realize sonic compositions. For quite a long time as well, although considerably lagging music notation’s representation of specific sonic events, there has also been an evolution in the conceptualization and description of musical process, of procedures of generation of musical data, instructions that can be followed by human beings or machines.

It is not the function of this brief essay to go into any specifics of the various attempts to compose music by pre-defined logical processes, either historical or personal. The evolution of today’s artificial machine-executable languages has given the description of music-as-process the jump start it has for so very long needed unleashing a vast variety of approaches. Algorithms are essentially shorthand notations for large numbers of specifics. A few operands and operators can instantiate a potentially infinite number of musical sounds. We can now relatively easily opt for a small number of relatively powerful variables instead of having to individually specify a very large number of weak variables one at a time by hand.

The distinction between ‘generative music’ (logic-based, algorithmically specified) and music composed by other means is vague. Composition that was to some degree rule-based was commonplace for centuries prior to computers. We composers have all studied species counterpoint, and many older ‘forms’ are actually process descriptions (e.g. canon and fugue) rather than abstract structures to fill in with material (e.g. sonata form, rondo or strophic song form).

Why - Personal

Although I have always been fascinated by abstract structures and at times attempted the design of algorithmic languages for music (Spiegel 1984, 1982-4), my own use of logically defined musical processes has often begun as self-simulation (perhaps a new form of self-expression), or alternatively as a sonification of an extramusical phenomenon, or at times an exploration of curiosity or hypothesis (“what would it sound like if…”). These attempts have run the full gamut from small logic modules that decide one specific aspect of something I am otherwise more intuitively creating (e.g. stereo placement) to stand-alone generative processes that, once set in motion, will go on potentially forever composing ever-changing musical material with no further human intervention.

By "self-simulation" I mean that I discovered early in my use of computers that some of my own sonic decision-making was predictable enough that I could describe it by rule. I wanted to be able to automate whatever aspects of my own musical decision-making

process I could delegate to logic in order to free myself to focus on and be engulfed in those aspects of the process I could not rationally explain. At times, I had the absurd fantasy of ultimately being able to automate enough of my mind's compositional processes that I could leave behind me at the end of my life an artificial simulator of my musical self that would be able to go on creating new Laurie Spiegel compositions long after I am no longer here to hear them or see the response.

Realistically, although I have been able to make logic-based musical entities that are able to play streams of ever-changing new material that embody some of my musical biases, these exercises fall far short of a true automation of my creative musical self. More often, algorithmic logic has functioned as only one part of my compositional process, combined with as-yet incomprehensible subjective components. What we might call “musical AI” was never my intent, only a byproduct of the desire to increase my ability to manipulate and generate musical material and to interactively have more real time musical power and control as a composer and as an improviser. No artificially constructed non-conscious logic-based entity will have the drive, passion, motivation or inner need to express itself musically.

Why - General

To make sense of using algorithms, i.e. descriptions of musical procedures encoded in logic, of process descriptions written in artificial languages as instructions that machines can follow, it would be wise to stay connected to our natural musicality and to go back to the question of why we create music at all.

Music may be background for other activities such as dance, theater, film or our ordinary lives. For such background we may want only a texture that has a certain rhythm, mood or quality of feeling. Music as foreground, as main focus of attention might be our personal moment-to-moment expression of emotion or of our individual sensuality, or it may attempt to capture, communicate or express subjective experiences that manifest themselves in our emotions or imagination. We want to externalize those subtle subjective phenomena so that others can also perceive them, to make shared what has been private.

Music can also be structured to represent in an abstracted form something we experience or perceive in the world around us, a narrative drama or a data set or structure or set of relationships that can be represented in sound. Such sonic captures can run the gamut from dramatic program music to the sonification of astronomical data.

Music can also be a form of soothing for other or self, providing experiences of flow, energy, peace, physicality, emotion or other subjective states.

Compensations for Lacks

The representation of many individual sonic events as a general description of process can be a faster and more efficient method of musical fabric generation than having to specify every aspect of each individual note. This is a time and labor saving innovation. However this method often constrains the music to an overall uniformity, sameness, predictability and flatness of overall form. An algorithm, once written, is outside of oneself and in itself is invariable, not subject to a musician's momentary changes of mood, sensory responses or ideation such as would naturally incorporate themselves into freehand writing or spontaneous improvisation. So the choice of what to automate versus

what to make malleable via means external to the algorithmic, especially by human input, is one a vitally important design consideration.

I have used various methods to attempt to overcome the dramatic flatness toward which most artificially-defined generative processes tend and to impose form on algorithmically computed musical material:

1. Interactivity

One method I have often used is to use make generative algorithms interactive, in effect delegating to logic only those subsets of my decision-making processes I am able to understand sufficiently to be able to encode into logic. I reserve to myself the power to specify interactively in realtime interaction with the sound other aspects of musical creation, those that I can't automate with a sufficient sense of my own aesthetic self to be as musically satisfying as I want the output to be.

The distinction between what I delegate to automated processes and what I reserve to the less-rationally-comprehensible methods of more intuitive specification is perhaps the most important aspect of my algorithmic design process. Doing this has always given me an unparalleled opportunity for introspection and increased self-awareness as to how I compose, of my personal musical preferences and of how my own creative mind works.

The variables reserved to my non-algorithmic control may be any of many kinds, ranging from realtime interactive adjustment of variables used by an algorithm during computation of the music to ex post facto non-algorithmic intuitive orchestration of material that was generated entirely by pre-defined non-interactive logic with no intervention.

2. Entropy

A second method I have frequently used is to employ the concept of informational entropy as per Shannon and Pierce's information theory (Shannon 1948, Pierce 1961). The informational entropy of a musical work can be varied throughout a musical composition and represented as a function of time. This curve can be designed to structure the listener's experience throughout the piece. Such a time function generates and controls the composition's emotional content in that an entropy curve represents the variation over time in the degree to which the listener can predict what will be heard in the next moment. The moment to moment variation of level of predictability that is embodied in an entropy curve arouses in the listener feelings of expectation, anticipation, satisfaction, disappointment, surprise, tension, frustration and other emotions.

3. Inherent structure

Another way to avoid overall dramatic flatness and to create form is to encode as a process description an evolutionary process that unfolds over time automatically until it arrives at some self-terminating culmination. Evolutionary and extrapolative processes tend to be open-ended however, open form, not self-bounding.

Once written, an algorithm is essentially a structure external to its creator. In other words it constitutes a new independent musical instrument or tool. I try to write such procedures in sufficiently general and adaptable form that they can be used to make a variety of different kinds of material and can be used in a variety of compositional works.

Varieties

I have previously listed some of the ways I have used algorithmic process descriptions (Spiegel 1997) and will list them again here:

1. Allusion: to very roughly approximate or simulate a natural occurrence that appears to me somehow inherently musical, capturing more so like an abstract painting than a photograph, a perception of something's process or shape rather than an exact replication (e.g. my piece 'The Expanding Universe').
2. Inverse Analysis: simply rendering into computerized form rules based on analysis of successful music of the past (e.g. 'A Harmonic Algorithm', which resulted from analyses of Bach Chorales).
3. Scientific modeling (designing data for the receptor): implementing, in a set of software-coded rules, generators of data designed to be cognitively meaningful or otherwise comprehensible according to perceptual or other kinds of research, e.g. Shannon and Pierce's information theory that formulates how to optimally encode content for intelligible reception, used in several pieces I made at Bell Labs (Spiegel 1997).
4. Mimicry of Process: coding into a computer program the rules by which some natural phenomenon transpires, unfolds or progresses (e.g. my realization of Kepler's 'Harmonices Mundi' (Kepler 1618-19, Sagan 1978).
5. Mimicry of Process Result : literal mapping of specific non-musical data onto musical variables (e.g. my little piece 'Viroid', in which I mapped the genetic content of a simple organism to a set of pitches).

6. Mixed (combinations of the above): the specification of one or more dimensions of a piece by one generative method while another dimension of the same piece is determined by an unrelated method (e.g. the Knowlton-Spiegel algorithm (Knowlton 1976) for an illusion of perpetual acceleration being used in the rhythmic domain at the same time as realtime interactive control of a corruption process that is being applied to the output of a data generation source in the pitch domain, in my 'Orient Express').

I choose among these approaches on the basis of the aesthetic qualities of both creative process and sonic output, on what I can learn by doing them and for their inherent fascination.

An extra-musical structure can derive from an abstract emotional sequence, a dramatic or documentary scenario, a natural ongoing process or any set of sequential data, each of these either as experienced subjectively by the composer or observed as external to the self. Generative methods can therefore be viewed as having overlap with program music that portrays an extramusical dramatic narrative at one end of the spectrum, and at the other extreme with scientific sonification (auditory data display), both reliant on extra-musical sources of structure.

Further Thoughts

Perhaps one of the reasons that my compositional output appears low is that the output of a generative process is potentially infinite. The overabundance of musical material that

algorithmic generation can produce somehow seems to cheapen the musical result, relative to music created by intentional specification of every minute detail by more traditional means. It can feel almost deceptive to record a mere finite short run of one variant of output from within a long ongoing process that can be altered in any number of ways.

I have most often listened to the output of my logic-based generative algorithms without ever recording it, not feeling it to be “my own music” yet but mere music-like texture. What do I need to do to form the material into my own personal expression, to impress upon it somehow a dramatic form that will infuse it with emotion or to invest it with my own sensuality?

To another way of thinking, that logic-based generative process, rather then any specific subset of its output, may be seen as a musical work in itself. Had computers been ubiquitous when I first did such works, instead of there being only a very few large rare computer installations owned by powerful institutions, would I have distributed the actual computer programs as musical compositions per se, for enjoyment by people at home and for them to vary and play in their own individual ways, instead of making the small number of specific finite form pieces that I did, each being only one mere single instance of a potentially infinite number of musical results that same logic could have generated? Is there any possibility or point of trying to compare the value of my program Music Mouse with that of the music on my “Unseen Worlds” album made almost entirely by use of it? Both are musical works. They are not both music.

Instances of my individual uses of my own various algorithmic logics are very much my own compositions and they constitute musical works for which the process descriptions (computer programs) are mere compositional tools. From a more normal musical perspective, those pieces, those examples of output from individual program runs, the specific musical works themselves, were the ultimate goal: music that others could listen to that was created with the mediation of logical processes I had described to the computer. To a perhaps-greater extent though than even the resultant works, I made them in order to inhabit the state of flow and concentration inherent in all forms of music-making as state of mind. Simply experiencing the process of interacting with the sounds made by sonically responsive computer logic was the highest motivator. This is not really much different from how I loved playing my first guitar.

References:

Spiegel, Laurie. *MDI Protocol Report.* June 1984, (unpublished) as referenced in and its implementation detailed in R. J. Pryor et al, *Music Description Instruction Format Preliminary Proposal: A Study to Define, Design and Develop a Prototype of a Protocol enabling the Encoding, Storage and Transmission of Music compatible with, and to form a subset of, the North American Presentation Level Protocol Syntax (NAPLPS) (TELEDON),* Canadian Dept. of Communications Research Contract 183-00361, April 1, 1985. New Media Technologies. Vancouver, Canada.

Spiegel, Laurie *IMPspeak Interactive Music Processor language*. Unpublished design specification, 1982-4.

Spiegel, Laurie in A*rtists' Statements*, Contemporary Music Review. Special Issue: Generative Music, ed. Nick Collins.Vol. 28, #1, 2009, pp. 127-129. Routledge, London (2009).

Shannon, C. E. 1948. A mathematical theory of communication. *Bell System Technical Journal.* Reprinted in Shannon, C. E. and Weaver, W. *Mathematical Theory of Communication*. University of Illinois Press, 1949.

Pierce, J. R. 1961. *Symbols, Signals and Noise: The Nature and Process of Communications*. New York: Harper and Brothers. Reprinted unabridged and revised as *An Introduction to Information Theory: Symbols Signals and Noise: New York: Dover Publications, 1980.*

Spiegel, Laurie. *An Information Theory Based Compositional* Model. In *Contributors' Notes*. *Leonardo Music Journal*, Vol. 7., pp. 89-90. MIT Press. 1997.