

Polyrhythmic Texture Derivations of Time-Lapsed Field Recordings

Scott Smallwood
Rensselaer Polytechnic Institute
110 8th St. DCC 135
Troy, NY 12180 USA
Smalls@rpi.edu
<http://www.arts.rpi.edu/~skot>

Abstract

Building sound compositions usually involves the organization of sound structures using some sort of organizational system. In the case of improvisation, aleatoric composition, and noise/sound art compositions, perhaps the systems are more a process of gesture, timbral exploration, or imitation of sounds occurring in the world/environment in which we live. In this paper I discuss some compositional techniques that I have found useful in working with field recordings. In particular, I will be discussing techniques used to compose a series of pieces for a recent CD project entitled *Tempus Fugit*.

Introduction

In this paper I discuss some compositional techniques that I have found useful in working with field recordings. In particular, I will be discussing techniques used to compose a series of pieces for a recent CD project entitled *Tempus Fugit*. I will begin by discussing the idea of time-lapsed audio recording; first outlining some challenges surrounding the use of field recordings in composition, and then discussing how time manipulations can address some of those challenges, while bringing new and interesting materials to work with. Secondly, I will briefly discuss some technical matters regarding time-manipulation techniques and their artifacts. Thirdly, I will discuss the concept of polyrhythmic layering of sound events into sonic textures. In this section I will discuss the layering and spatial placement of textural cells, and will explain the distinction between a textural cell and more discrete rhythmic/metric cells

I. Time-Lapsed Field Recording: What and Why?

Building sound compositions usually involves the organization of sound structures using some kind of system. In the case of improvisation, aleatoric composition, and noise/sound art composition, perhaps the systems are a process of gesture, timbral exploration, or imitation of sounds occurring in the world/environment in which we live. As a composer, I've been learning how to listen to the environment around me, and how to incorporate what I've learned into my compositional work. Within the last several years, the use of the sounded environment as material for composition has become not so much a point of departure, but a destination in and of itself. Pauline Oliveros has

challenged us to ‘connect with the acoustic environment and all that inhabits it’ by ‘going below the surface of what is heard and also expanding to the whole field of sound whatever one’s usual focus might be.’ (Oliveros 1998: 114-15) Composers and acoustic ecologists such as Christina Kubisch, R. Murray Schafer, and Hildegard Westerkamp have shown us that the sounded environment can itself be composition, and have, in fact, encouraged us to ‘imagine finding or creating the places where music and soundscape enhance one another.’ (Westerkamp 1999: 25)

Recently, I’ve tried to figure out how I can be informed by the materials of the sounded environment without being forced to deal with the narrative associated with any given moment. Time lapse photography and video are able to represent the essence of a scene as an abstraction, blurring the narrative associated with that scene. I began to envision a sonic equivalent to this process by recording several hours of environmental sound and compressing it into several minutes. This process leads to interesting sonic results. Rhythmic patterns which previously were undetectable come to the fore. Noise contours and discrete sonic motives are recontextualized. Most importantly, the overall structure of that large moment in time is retained, but the associated narrative; that is, the specific content of human and animal conversation, the measured sound of automobiles and machines, the rhythms of moving water and wind are blurred, leading to a bed of abstract materials. A kind of sonic tableau is created, leaving an underpinning, and perhaps even a lens for orchestration, from which materials can be drawn, and structure can be derived.

As I began experimenting with this process, I quickly discovered that, for me, there was a delicate balance between the amount of time a field recording lasts, and the percentage by which a recording was compressed. In general, compressing by less than about twenty-five percent obscured the soundscape too much, while more than twenty-five percent did not give me a satisfactory level of abstraction. For this project, I chose recording times of between thirty minutes and one hour. This decision was made partially because of the technical problem of tape-length, and partially because at a compression percentage of twenty percent, this yielded single compositions of between seven and fifteen minutes, allowing me to use the entire compressed field recording as the over-arching structure for each composition.

I chose six environments for recording: three indoor environments and three outdoor environments. Although I do not reveal any information on the CD about the specifics of the locations chosen, I will list them here:

- Light rain shower in my relatively quiet neighborhood in Troy, New York
- Arrival/departure gate in the Las Vegas Airport
- Indoor swimming pool in Oak Harbor, Washington
- Woodland park in the Arboretum in Seattle, Washington
- Bowling alley in Oak Harbor, Washington
- Beach underneath the deception pass bridge on Whidbey Island in Washington

For some of these environments I recorded about an hour’s worth of audio, and in others I recorded between thirty and forty-five minutes. All six of these recordings were done with an Audio-Technica AT-825 stereo condenser microphone to DAT, and were later compressed by twenty-five percent using relatively standard time-compression algorithms in Digidesign’s ProTools.

II. A Brief Survey of Pitch/Time Shifting Methods

At this point it might be worthwhile to outline the various techniques for time-compressing audio, and the resultant artifacts that accompany these processes. The problem involved in changing the time of a recording without changing the pitch (or vice versa) involves two possible types of processes. The first type is the application of a mathematical process to the soundfile. The second type of process requires applying what we know about the human auditory system to the soundfile.

There are a variety of mathematical processes that can be used to change the pitch or the time of a sound-file independently of one another. The primary methods include using the phase vocoder, the wavelet transform, time-granulation, linear predictive coding, and simple sample-rate conversion. While comprehensive explanations of these methods can be found in Curtis Roads's *Computer Music Tutorial* (Roads 1996: 440-48), I will give brief descriptions of each:

The Phase Vocoder: This method employs a frequency-domain analysis of the soundfile (using the fast Fourier Transform, or FFT), and then re-synthesizing the soundfile using additive synthesis. The analysis is done in a series of short, overlapping segments, or frames. Once the pitch information in the soundfile is analyzed and catalogued, it is a simple matter to change the time or pitch of the sound. To alter the time, the frames can simply be read back at different rates by changing the amount of overlap between frames. To alter the pitch, the frequency information in the frame can simply be re-scaled. This method is particularly useful for pitched-oriented sounds, but can be problematic for complex sounds, or sounds with complex attacks. Artifacts include smearing and reverberation effects.

The Wavelet Transform: This method is similar to the phase vocoder. A frequency-domain analysis is made, but unlike the FFT used in the phase vocoder, the segments are not of fixed duration; rather, they are dependent on the frequency information. The higher the frequency content in the analysis, the shorter the duration of the segment, or wavelet. Wavelet transforms tend to be more accurate for sounds with lots of high-frequency content because they are better at determining the timing-resolution of higher frequencies.

Time-Granulation: A time-domain-based process, time-granulation involves a kind of rapid-fire, microscopic sampling of the sound-file into hundreds or even thousands of grains per second. These grains can have their own unique envelope structure, and may overlap each other. To alter the time of the original sound file, grains can be dropped or doubled. To alter the pitch, sample-rate conversion is performed, and grains are doubled to account for the temporal change that would normally result. The obvious problem with this approach is that the transitions between one grain and the next can cause transient artifacts due to changes in level, waveform zero crossings, and other inconsistencies.

Linear Predictive Coding: This method involves analyzing a sound-file and extracting information for modeling the excitation/resonance properties of the instrument that is the subject of the recording. For this reason, linear predictive coding is limited to sounds that have the characteristics of the voice or of musical instruments. The information stored in the analysis, which includes pitch information, filter coefficients, duration, and other kinds of data, can be changed and then re-synthesized in order to change the time or pitch.

Sample-Rate Conversion: This is the simplest method of pitch conversion, and is frequently used in hardware-based pitch shifters. Here, the incoming sampling frequency is changed, and the ratio of the incoming sampling frequency and the outgoing sampling frequency determines the change in pitch. To maintain the original timing, samples are added or dropped. This method is limited to pitch conversion only.

The other type of process that can be used for time and pitch scaling involves understanding how we hear those types of changes, and applying this understanding to a model of the human auditory system. Using such a model, we can create an auditory scene analysis of a given sound-file and alter it according to the way we seem to hear those changes. At least one manufacturer of time/pitch scaling software is now using this technique.

For this project, I simply used the Protools time-shifting plug-in to alter the length of the soundfiles. Although the methodology that it uses to time-compress audio is not the most satisfactory in terms of sophistication and resultant artifacts, it worked just fine for this project.

III. Sonic Textures and Polyrhythms

Once the time-lapsed field recordings were completed, I listened to them many times, and began to extract short cells. These cells became my most basic compositional unit, and were looped and in some cases processed, to become pulsating, oscillating textures. Most of the cells were rhythmic in nature in and of themselves, but rather than treat the motivic units as musical metrics, I instead treated them as features of the overall textural landscape of the piece. Each texture of repeated cells became of kind of independent thread within a complex fabric of sound.

As I began to add textural threads, I placed the basic cell into higher-order polyrhythmic relationships to cells in other threads. For example, three repetitions of a cell in one thread might fill the same amount of time as four repetitions of a cell in another thread. At any given moment there might be from one to four different threads happening simultaneously, and together with the original time-lapsed field recording as an underpinning, they would create an ever-changing sonic landscape. Threads were situated in the stereo field in unique places to help maintain some independence.

The large-scale manifestation of sound that results from this process yields a form that is determined by the unique shape of the moment of time that is encapsulated in the field recording. It's inner, middle-scale form is determined by the selection of cellular units of sound and the spinning of those cells into threads which weave in and out of each other rhythmically and spatially. The level at which rhythm is organized in these compositions is above the local level of the cell, since there is no specific tempo or meter that is unique to all of the cells in adjacent threads. This gives the music a sense of rhythmic inertia while still being metrically ambiguous.

As far as the technical aspects of this process, creating the cells was simply a matter of finding short chunks with distinct characteristics from the time-lapsed field recording and looping them, taking care to make the bounds of the selection line up at zero-crossings. The cells were simply copied as stereo units into empty tracks in Protools, panned appropriately, and repeated for the desired amount of time. Cells for adjacent textural threads were measured by comparing their size to other cells,

determining what the closest length ratio appeared to be, and adjusting their length appropriately to fit that ratio exactly.

IV. Conclusions

The process that I've described has been, for me, a useful and effective method for letting the sounded environment inform compositional choices. It has allowed me to frame a sound piece around the natural rhythm of the moment, while blurring the specific narration of that moment. Morton Feldman once wrote

To think of a music without instruments is, I agree, a little premature, a little to Balzacian. But I, for one, cannot dismiss this thought. In creating this indeterminate situation I began to feel that the sounds were not concerned with my ideas of symmetry and design, that they wanted to sing of other things. They wanted to live, and I was stifling them. (Feldman 1972: 111)

Our environment sings to us everyday. Sometimes we take the time to listen to it; at other times we only hear and don't notice the music. For me, this exercise has been not only a method for creating an impression of a soundscape, but it has also helped me to do a better job listening to the music that the environment makes.

I came home one day in March during a light rainstorm, and I heard some sounds that were familiar to me, almost as if I were hearing a piece of music that I knew coming from someone's open window. Then I realized that I was hearing the sound of rain dripping from the roof onto my driveway in rhythmic patterns. One of the pieces on the CD was made from the soundscape in front of my house during a rainstorm. What I was hearing was reminding me of that piece; that is, I heard the same sounds that were in the music I had created. Then I corrected my thought: the music was already there; I had just borrowed the materials.

To hear examples of the work described in this paper, visit <http://www.arts.rpi.edu/~skot>.

References

- Feldman, M. 1972. A compositional problem. In B.H. Friedman (ed.) *Give My Regards to Eighth Street: Collected Writings of Morton Feldman*. Cambridge, MA: Exact Change.
- Oliveros, P. 1998. *The Roots of the Moment*. New York, NY: Drogue Press.
- Roads, C. 1996. *The Computer Music Tutorial*. Cambridge, MA: MIT Press.
- Westerkamp, H. 1999. Say something about music. In B. LaBelle and S. Roden (eds.) *Site of Sound: Of Architecture and the Ear*. Los Angeles, CA: Errant Bodies Press.