Kerry L. Hagan

Digital Media and Arts Research Centre
CS 2011
University of Limerick
Limerick, Ireland
Kerry.Hagan@ul.ie

Textural Composition: Aesthetics, Techniques, and Spatialization for High-Density Loudspeaker Arrays

Abstract: This article documents a personal journey of compositional practice that led to the necessity for working with high-density loudspeaker arrays (HDLAs). I work with textural composition, an approach to composing real-time computer music arising from acousmatic and stochastic principles in the form of a sound metaobject. Textural composition depends upon highly mobile sounds without the need for trajectory-based spatialization procedures. In this regard, textural composition is an intermediary aesthetic—between "tape music" and real-time computer music, between sound objects and soundscape, and between point-source and trajectory-based, mimetic spatialization.

I begin with the aesthetics of textural composition, including the musical and sonic spaces it needs to inhabit. I then detail the techniques I use to create textures for this purpose. I follow with the spatialization technique I devised that supports the aesthetic requirements. Finally, I finish with an example of an exception to my techniques, one where computational requirements and the HDLA required me to create a textural composition without my real-time strategies.

Textural Composition and Its Aesthetics

Textural composition is the practice of composing primarily with sonic texture over other musical elements to create a particular experience of space and time. Specifically, a textural composition attempts to give the listener a large, immersive experience of sound at slow, environmental movements of time. The philosophy of textural composition combines the aesthetics of the sound object from acousmatic music with large sound masses inspired by Iannis Xenakis. The following sections describe the nature of sound masses in textural composition and the incumbent spatialization strategies to engender the musical and acoustic space. They summarize points and arguments made in greater detail elsewhere (Hagan 2008a, 2008b). I will highlight the conclusions and the main reasons behind them.

Sound Objects

Texture in music is a distinctly useful metaphor for the experience and construction of certain musics, especially when used in opposition to gesture. A gesture is experientially short, with a perceptible

start, duration, and end. It has directionality and impetus. It lasts for a human length of time, to paraphrase Denis Smalley (1997). Conversely, texture becomes about inner detail at the expense of a direction and momentum, existing in what Smalley calls environmental time.

The literal definition of texture can inform and suggest musical treatment. Texture is a characteristic of an object, requiring a substrate on which to exist. In focusing on texture, the object itself becomes irrelevant. The size and shape of the object, for example, no longer provides significant information to the observer. Musically, this suggests that the sound object itself must change to enforce awareness of its texture, rather than its size and shape, edges and boundaries, or the placement of the sound object in the world.

The acousmatic sound object manifests in a work through spectromorphological boundaries, separating it in timbre and in time from other sound objects. As a musical building block, the sound object has become a crucial tool for compositional devices in acousmatic composition. Putting aside the space of a sound object for a moment, the perception, manipulation, and development of sound objects become the form and catalyst of acousmatic works. Imagine, however, that an acousmatic work is frozen in a moment of its life, exposing a single sound object to scrutiny. The sound object is then

substantially magnified so that the listener's entire perspective is of nothing but the sound object in that moment.

The texture metaphor allows for this process: focusing attention on nothing but the inner detail of something that is existing outside of time and larger than the space a listener experiences. In other words, there must be a sound metaobject. It retains its identity by way of its inner details, but the metaobject is so large that it subsumes the listener.

An aside: One approach to focusing on inner details comes from drone music, where long, sustained tones change subtly over time, inducing a listener to target on microscopic changes. Unlike textural composition, however, drone music does not necessarily expand the imagined size of the sound. Textural composition requires multitudinous, smaller sounds that reside in and on the expansive metaobject, like the grit on sandpaper. This facilitates the perception of the metaobject's magnitude. It is made up of many miniature components, which are now heightened to apperception. This idea is congruent with Natasha Barrett's argument that density, texture, and amplitude contribute significantly to "implied spatial occupation" (Barrett 2002, p. 316). The richness of the texture, its long duration, and its stochastic asynchronicity also contribute to its perceived volume (Truax 1998).

When this happens, a new experience emerges, one that approaches what R. Murray Schafer called a "lo-fi" soundscape (Schafer 1994). In this soundscape, sounds become too congested and innumerable to be individuated, and they impart an impression of a circumambient acoustic scene.

But, textural composition does not entirely become a lo-fi soundscape, either. It exists on a fragile edge between countless sound objects and the lo-fi soundscape. This is one reason I have characterized textural composition as an intermediary aesthetic, music living between sound objects and soundscape.

Space and Textural Composition

"Space in sound" (timbre) and "sound in space" (diffusion) are fundamental aspects of acousmatic music (Truax 1998), and perhaps in most music. It follows that the presence of a sound metaobject requires both large musical space (imagined through musical perception) and immersive, acoustic space (created in the world). Spatialization must accommodate the size of the metaobject, assert the complete envelopment of the listener, and convey the complexity of the inner detail. As a result, there are four important, interrelated aspects to the aesthetics of the spatialization: 1) multichannel surrounding sound, 2) a lack of "front" in perspective, 3) an unbounded sense of space (both acoustically and musically), and 4) significant spatial motion.

In enveloping sound, it is clear that only multichannel arrays can create the surround experience. For this reason, textural composition requires a minimum of eight channels to be effective. More speakers, especially elevated speakers, enhance the experience, so high-density loudspeaker arrays (HDLAs) become crucial spaces for textural composition (Harrison 1998; Rolfe 1999; Sazdov, Paine, and Stevens 2007; Otondo 2008; Normandeau 2009).

At the same time, there are arrays with large numbers of speakers that would not suit textural composition. In some arrays, especially those designed for live diffusion, speakers are concentrated at the front of the audience. Smalley (2007) defines a number of acoustic spaces around listeners. In acousmatic music, prospective space, the front view for the listeners, is the main arena for the acousmatic image. Other spaces exist, defined by the relationship of the listener and the imagined "view" of the music. In these spaces and prospective space, the listener is outside the sound, viewing it as it moves and turns in space. Because the listener must be inside the sound metaobject, these spaces defeat perception of the metaobject. There is one space defined by Smalley that is conducive to textural composition: immersive space, a kind of "circumspace." Immersive space is "where the spectral and perspectival space is amply filled, surrounding egocentric space, where the pull of any one direction does not dominate too much, and where the listener gains from adopting, and is encouraged to adopt, different vantage points" (Smalley 2007, p. 52). For this reason, there should be no frontal perspective dominating the environment.

The sound metaobject must be unbounded imaginatively and acoustically. Therefore, textural composition must use unbounded space, even though bounded space is far easier to achieve.

The listener's point of view is one way in which space is bounded. There is little a composer can do to encourage a nondirectional viewpoint. The material can suggest multiple perspectives by saturating the circumspace to create immersive space. Alternatively, the composer can invite the audience to move through the room in the performance. The latter is met with difficulty in many performance spaces, however. Therefore, a consistent, unoriented approach to spatialization unbinds the limited perspectival space, becoming circumspace.

Discernible, individual sounds can also bound space by articulating it through movement and psychoacoustic triggers. The nature of textural composition, that edge between sound objects and soundscape, alleviates this problem, as it does not have individuated objects moving through space.

Too many sounds can also contract a space. Here I draw a subtle distinction between a sound mass and a sound wall. Inspired by much of Xenakis's music, my work in textural composition deals with large sound masses. Masses can be dense, thin, transparent, or dark—they are nonspecific. A sound wall, on the other hand, implies a heavy, impenetrable, and bounding space. A sound mass can surround a listener, appearing to extend beyond the horizon. A sound wall blocks space beyond its perimeter. Though a sound wall can have texture, only a sound metaobject is unbounded. So, textural composition is about sound masses, not sound walls. Spatially, this means that textures extend distally, as well as enveloping the listener.

There is nothing inherently objectionable to static textures in textural composition. I found, however, that when texture becomes the primary component of a work, then it must take on dynamic, changing processes. This may occur in the sonic nature of the texture itself. More importantly, I discovered that mobile textures enhance textural composition. Electroacoustic music is predominantly engaged with mimetic spatialization, that is, the mimicry of actual sounds in space. Sometimes it capitalizes

on surrealism by making identifiable sound objects move faster, higher, longer, etc., than might be their natural circumstance. Imagine an airplane flying around like an insect. This mimetic spatialization dissimulates the loudspeaker to create a perception of space between speakers, e.g., the work of Ambrose Field (Austin 2001). But this mimicry means that there is a gesture in space, which defeats textural composition. An alternative approach is point-source spatialization. Point-source spatialization imbues each loudspeaker with musical agency as performers (Burns 2006). This promotes the circumspace, as each speaker contributes equally to the experience.

This precipitates another aspect of spatialization in textural composition, perhaps the most meaningful to the aesthetics of spatialization. Textural composition needs to have dynamic, active, exaggerated movement without coherent trajectories. There is thus a balance between the hidden speakers of mimetic spatialization and the performer-speaker in point-source spatialization. Colby Leider identifies these extremes as the "dual nature" of loudspeakers (Leider 2007, p. 1891), but I believe that spatialization can exist on a fragile boundary between both, another intermediary aesthetic.

There is an additional, fifth aspect of textural spatialization not mentioned earlier because it is an element of my larger philosophy: real-time spatialization. My concerns regarding real-time generation of music are expounded in the next section. Harrison (1998) and Smalley (Austin 2000) assert that space as a central aspect in music-making is mostly unique to acousmatic practice. I disagree with this, but the way in which they describe sound in space informs my approach, even if as an extension to acousmatic aims. The real-time diffusion in acousmatic music is paralleled in the real-time spatialization of textural composition.

Creating Textures

In this section, I describe the way in which I create textures for two pieces, each using different methods. Both use the same spatialization technique, described in the next section. The first piece,

Real-Time Tape Music III (2008), manipulates sound files for texture sources. The second piece, *Morphons and Bions* (2011), uses noise-based synthesis techniques for its material.

Real-Time Tape Music III

The "Real-Time Tape Music" series evolved from discussions I had with many composers and from trends I saw in festivals and conferences. On the one hand, some computer music composers felt that real-time computer music was the obvious solution to "dead" fixed-medium works. On the other hand, composers working in fixed media—namely, acousmatic composers—similarly felt that working with recorded sounds was the apparent solution to "dead" computer music. Fixed-medium works were considered dead because they were unchangeable regardless of venue or audience. Computer music was considered dead because the sounds were lifeless and uninteresting compared to real-world sounds. I have oversimplified both the arguments and solutions on each side, but this was my initial inspiration. I wanted to create music that used real-world sounds, and I processed them with classical tape techniques in real-time. Initially, I felt I was bridging a gap that was artificially and uselessly wide. Translating acousmatic and fixed-medium aesthetics to real-time computer music practices remains a crucial part of my compositional aims in all of my works.

The pieces in the "Real-Time Tape Music" series sample a number of sound files containing recordings of animals, musical instruments, and quirky sounds that evoked careful listening, e.g., an amplified recording of soda bubbles in a can. The samples were played with different speeds, lengths, directions, and loudness—standard classical tape techniques. Random processes set the values for each of these characteristics. Other random processes selected onset times in the sound files. Timed events triggered different layered montages of sounds, creating the structure of the pieces.

As computer processors became more powerful, more of these layers could be added. Eventually, I layered so many sounds at once that something else emerged: textural composition. This led to my first work in textural composition, *Real-Time Tape Music III*. No longer sounding simply like acousmatic music made in real time (as did the first two works in the series), *Real-Time Tape Music III* demanded that I expand my compositional aim, and my first concept of the sound metaobject arose. Using ten sound files, each sampling of a file led to a texture "stream." Different combinations of the streams, constructed by ear, formed eight possible textures.

With so many parameters and textures, I decided to control the structure of the piece with stochastic processes. In this case, the timed events controlled macroscopic form, while a Markov chain determined which texture would surface at any given moment. I borrowed the transition table from Xenakis's *Analogique A+B* (cf. Xenakis 1992).

The final work consists of two contiguous movements. The first introduces the textures through random processes. The second movement relies on the Markov chain for stochastic movements between textures. Each movement manifests a different sound metaobject. The sound metaobjects that result are large, slowly shifting sound masses, spanning the audible frequency spectrum. The idiosyncratic traits of each sound file establish memorable moments of sounds that recur in various forms. But the montage of these sounds creates a consistency throughout the work. In this way, the metaobjects exist for the duration of the movement, but the ephemeral inner details emerge, shift, and evaporate quickly.

Morphons and Bions

Returning to arguments that digitally synthesized sounds were "flat" or dead, I decided to experiment with creating textures from synthesized sounds that were more "living," or acoustically natural, while exploring textural composition. I was still interested in the sound metaobject and acousmatic approaches, but again, looking for a way to combine them with real-time computer music.

One of the most relevant components to acoustic sounds, I believe, is the inherent noise in the sound.

In some cases, that means white noise. But in most, it means correlated but random fluctuations in the frequency spectrum. Therefore, I combined standard additive synthesis and frequency modulation (FM) synthesis with noise generators. Several philosophies of noise informed this approach, affecting the form and direction of the work (Hagan 2012, 2013).

The wide range of sounds in *Morphons and Bions* comes from two fundamental synthesis methods, which in some cases are combined with other processes. The first is additive synthesis modulated with white noise. The resulting signal is

$$x(t) = \sum_{k=1}^{6} w(t) \sin(2\pi h(t) f_0(t) t + D(t) n(t)),$$

where

$n(t)$ = white noise;
$D(t)$ = depth (amplitude of white noise), changing in time;
$f_0(t)$ = fundamental frequency, changing in time;
k = partial number;
$w(t)$ = Gaussian random variable with mean $1/k$; and
$h(t)$ = Gaussian random variable with mean k.

The noise in this synthesizer does not affect the amplitude of the spectrum directly, but the increase of noise affects the sidebands present as they would in any FM synthesizer. The sidebands are random, however. By changing the depth of noise, the synthesizer can sound like a simple oscillator bank or nearly white noise at the extremes. Because one white-noise generator affects all of the partials, psychoacoustic fusion enforces the perception of a complex sound, rather than six separate sounds. Amplitude envelopes further fuse the synthesizer.

The second basic synthesizer is frequency-modulated noise. Cascaded band-pass filters with the same center frequency filter the noise to result in a strongly colored, almost pitched, noise. An oscillator controls the center frequency of the filters, creating sidebands while retaining some quality of the noise. The depth and frequency of the oscillator determine the sidebands that emerge.

Random processes control the depth and Q of the filters.

The center frequency of the band-pass filters is

$$f_c = D(t) \sin(2\pi f_{filter} t),$$

where f_{filter} is the frequency of the oscillator controlling the band-pass filter's central frequency and $D(t)$ is the depth, or amplitude, of the filter oscillator, controlled by a random process. The sidebands then appear at $D(t)$, $D(t) \pm f_{filter}$, $D(t) \pm 2f_{filter}$, $D(t) \pm 3f_{filter}$, etc. (Hagan 2013).

By changing Q, this synthesizer can also range in noisiness, much like the noise-modulated additive synthesis.

To create sounds with more complexity, I combined these synthesizers. The frequency-modulated noise replaced the white noise in the noise-modulated additive synthesizer, creating partials with sidebands, forming a rich, quasiharmonic spectrum.

Amplitude envelopes articulated these synthesizers. Some envelope durations were fixed, but in other cases, random processes determined the duration. In another synthesis technique, shortened envelopes with a mean duration of 23 msec, chained iteratively, added amplitude modulation to the synthesizer. This created additional sidebands approximately ±43 Hz around the original sidebands generated by FM.

Timed changes in the parameters of random variables determining the timbre of sounds create the overall form of the work. Generally speaking, the timbres begin quite noisily. Over time, they become less noisy and more pitched. By the end of the work, the pitches become stable at integer frequency relationships to each other. The end of the piece complements this process: isolated, digital clicks of instantaneous amplitude changes, familiar to anyone working in digital synthesis.

Morphons and Bions addresses questions of noise from synthesis to form. The definition of noise is fluid, depending mainly on context. Does it mean randomness? Does it mean something that distracts from a signal? Or does it mean a type of sound, qualified by acoustic characteristics and psychological response? First, acoustic noise is

fairly easy to define, and it was the entry point for *Morphons and Bions*. By combining digital synthesis with correlated noise, sounds became more alive, interesting, and rich. It is similar to acoustic noise in traditional instruments, which Cowell identified a long time ago (Cowell [1929] 2004). But the movement from noise to pitch is what John Cage (1958) proposed as the new consonance and dissonance of modern music (Hagan 2012).

When noise is counterposed to signal, it implies that there is message to the signal. But, in *Morphons and Bions*, noise is the signal. The piece relies on the random, from noise itself to the procedures controlling the synthesis methods. For this reason, the piece creates an interesting paradox identified by Kim Cascone's (2000) post-digital aesthetics of failure: that noise, which is signal, is no longer noise (Hagan 2012).

Perhaps the most complicated condition of the aesthetics of noise in *Morphons and Bions* comes from its roots. I was still concerned with textural composition, extending the acousmatic experience to the sound metaobject, and doing so in real-time. As such, noise is not just signal, but meaning and cultural artifact (Hagan 2012). The "nonmessage" as cultural artifact reflects the thoughts of Douglas Kahn (1999).

Ultimately, these processes and philosophies resulted in textures of a sound metaobject. Like *Real-Time Tape Music III*, *Morphons and Bions* represents a frozen, gargantuan metaobject inhabited by the listener, whose inner details mutate over time.

Stochastic Spatialization Technique in HDLAs

I compose and work in Pure Data (Pd), so I created my spatialization algorithm in that software. The controller patches and abstractions can easily be reused in any composition. Using basic amplitude panning, my spatialization algorithm takes the speaker location in the circle (the front arbitrarily assigned to 0°, continuing around the space clockwise until 360°). A texture stream is assigned a virtual angle, controlled in real time by random processes. The stream is also given a width, or signal skirt, in degrees. For each speaker, a Pd abstraction

Figure 1. The virtual angle of the signal (α), speaker angle (φ), and the phased copy of the signal. The cosine curve determines the amplitude of the signal for any speakers located within the signal skirt (k). (Figure from Hagan 2008a.)

first calculates whether the speaker is within the skirt of the virtual angle. It then uses a cosine curve to calculate the amplitude of the texture stream for that speaker. I named this abstraction `weightor~`. The simple calculation is

$$\text{Amplitude of signal in speaker} = \cos\left(\frac{\pi\,|\alpha - \phi|}{k}\right),$$

where

k = signal skirt,
α = virtual angle of the signal (texture stream), and
ϕ = speaker location.

By ramping between virtual angles, the signal moves clearly around the circle. The signal skirt may be widened for arrays with fewer speakers to avoid gaps in the sound between speakers.

Figure 1 shows how the amplitude is calculated for the original signal and its phased copy. I explain the phased copy, a variable delay moving in and out of phase, next.

Maja Trochimczyk (2001) identifies spatialization as either a circle (sound around the audience) or a net (sound throughout the audience). Ideally, textural composition would exist in a net, i.e., speakers around and throughout the audience. Most HDLAs are circles, however. My spatialization technique therefore assumes that speakers exist outside the audience's physical location. Capitalizing on the perception of interaural time delays, a copy of the texture stream is sent to an oscillating variable delay of a few milliseconds. This delayed, phased copy is

Figure 2. The accumulated effects of the spatialization algorithm: First the simple circular movement (a), then the appearance of center crossing with a variable-delayed copy opposing the original (b), and, finally, the effect of reverberation (c). (Figure from Hagan 2008a.)

Figure 3. Result of multiple, independent audio streams. Each curve is a representation of how individual texture streams are moving throughout the listening space, each relying on independent, combined effects of panning, delays, and reverberation. (Figure from Hagan 2008a.)

sent to a virtual angle 180° from the original (see Figure 1). The result is that not only does the signal appear to circle the listener, but it also appears to cross the center along the way. This approximates a "net" when only a "circle" exists. This was combined with the amplitude calculation above in an abstraction called `phweightor~`. Figure 2 depicts the accumulated effects of the spatialization algorithm.

When multiple texture streams, making up a single texture, are spatialized independently, the result is a swarming, motile sound metaobject surrounding the listener. By introducing a simple, cheap reverberation patch, distance and space become even larger. Random processes control the wet/dry mix, varying the distance of the sound, as well. There is no need for sophisticated reverberation because the layers of sound and motion obscure the details of reverberant sound. Figure 3 shows the result of multiple, independent audio streams.

This algorithm was originally designed for any number of horizontal speakers. Scaling azimuth

speakers is easy. Simply adding more `weightor~` and `phweightor~` abstractions, one for each speaker, can adjust the spatialization for any number or placement of speakers. The signal skirt may need to be adjusted, as well, to account for the distance between speakers.

At the same time as I was developing textural composition, I built the Spatialization and Auditory Display Environment (SpADE) at the University of Limerick. This 32.2-speaker system can be configured as needed. I used a configuration of 16 speakers placed symmetrically at eye level with an additional matching speaker placed above: a lower and higher ring of speakers evenly distributed around a circle.

The algorithm did not port well to elevated speakers. One solution would have been to integrate vector base amplitude panning (VBAP; see Pulkki 1997) into the patch, but concerns about processor demands arose. The sophistication and verisimilitude of VBAP was not necessary for the spatialization environment, either, because realistic locations and movements of sounds were not essential.

I experimented with assigning false angles to the upper speakers. One method was to assign the lower speakers values from 0° to 180° and the upper speakers 180° to 360°, as seen in Figure 4. The result was that the sounds moved more quickly through the spaces, but signals did not move from the lower ring to the upper ring except at one point: the front speaker.

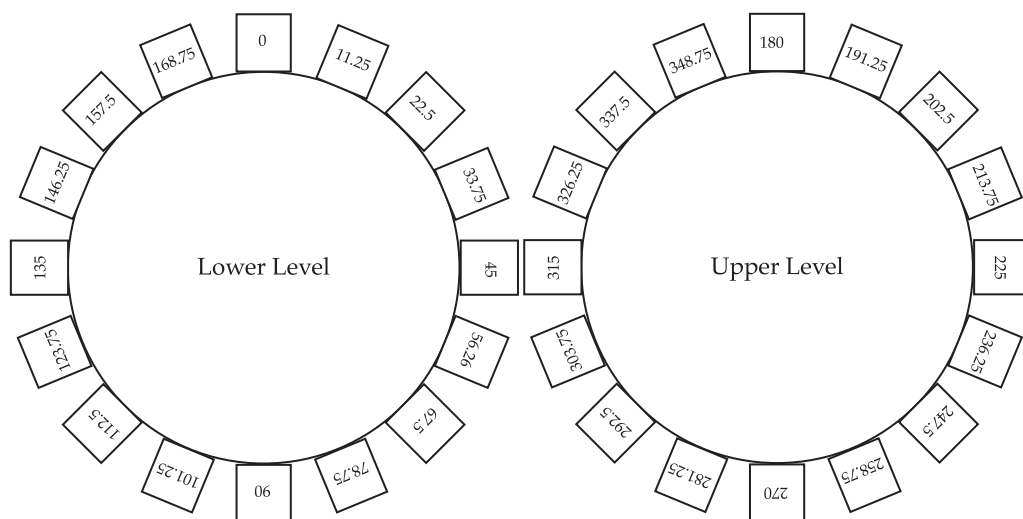*Figure 4. One possible configuration of false speaker angles.*

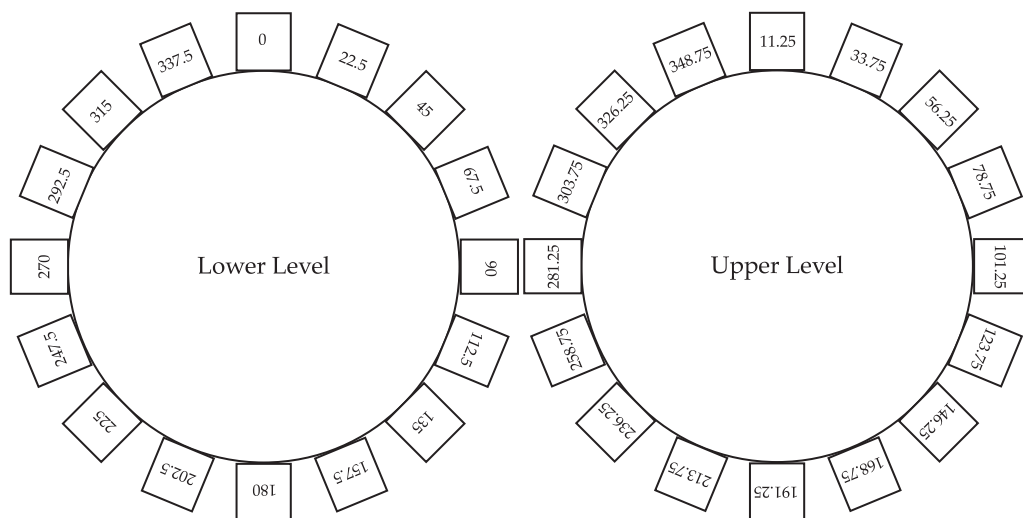*Figure 5. Assigned speaker angles for SpADE.*

Figure 4

Figure 5

Instead, I assigned the upper and lower rings alternating angles, shown in Figure 5.

In this configuration, the rate of motion was the same, but sounds moved up and down as well. I did not need the long, coherent trajectories of path-based spatialization. This adaptation to elevated speakers fared well, creating a sense of turbulent, moving sounds surrounding the listener on all sides.

Adapting Textural Composition to Fixed Medium

The 2015 conference of the Society for Electro-Acoustic Music in the United States (SEAMUS) provided an opportunity to present a work in the Cube, a four-level, 124.4-channel array at the Moss Arts Center at Virginia Tech (see Lyon et al. 2016 for more detail on the Cube). The

technical specifications of the space required a fixed-medium composition. Also, the array has three tiers of speakers and ceiling speakers. This venue thus required me to adjust my process. First, I could not generate the piece in real time. Although real-time composition is important to me, fixed media allowed me to use a much more computationally expensive synthesis algorithm with many more texture streams. Second, my spatialization algorithm, though adjustable as detailed earlier, could not be implemented because the synthesis could not be generated in real time, nor would false angles accommodate the large number of elevated speakers. The situation required me to rethink textural composition spatialization. The result was *Cubic Zirconia*.

Texture Synthesis in *Cubic Zirconia*

The most common stochastic process in my works is the Markov chain. In a discussion with Miller Puckette, he proposed another approach to generating numbers. He suggested that composers who use Markov chains want the results of the process to be more uniform; if there is a 75 percent chance of an occurrence, then all previous results should only have 75 percent of that occurrence at any given time. This is not what I want from my Markov processes, but I found the notion irresistible.

The algorithm that Puckette designed, called z12, takes a set of twelve probabilities for the numbers 0 to 11. It tallies how many of each number has already occurred, then outputs the number that has the greatest deficit. In this way, one can be assured that, at any given point, the percentage of each number closely matches the probability of that number. For simple ratios, this creates a fixed, repeating pattern. When the probabilities are irrational, however (e.g., golden ratios), the patterns that emerge are quite complex. Certain sequences can repeat, or nearly repeat, with minor differences. Sections of sequences can repeat occasionally. There is no small number of iterations, however, for which an exact repetition of sequences occurs. At this point, we cannot prove that the system will or will not exactly repeat at all (Puckette 2015).

At first, I applied this to larger events, much as I use Markov chains. There is room for discovery at this level, but my initial results were underwhelming. Then, Puckette programmed the patch to output a one or zero, based on the sequence of numbers. For example, if three follows two, without a four intervening, generate ones until a five appears, and then generate zeros until a three follows a two again. I called this the "logic patch" because *~ objects act as AND gates, sending out ones and zeros. Then we accelerated the output to audio rate, creating samples of ones and zeros. The result was a complex waveform that transformed both microscopically and macroscopically. Occasionally, the process would go silent, starting up again when the correct sequences returned. Because z12 does not repeat for a significant number of outputs, a single z12 can generate a complex waveform that does not sound the same for the average duration of a work. We have not listened longer than ten minutes, but in that time, no repetitions appeared.

If one seeds z12 with initial values (in other words, tell it that it has already had a certain amount of each number), the output changes dramatically, even when the probabilities remain unchanged. This makes the z12 algorithm incredibly powerful for generating materials. Unfortunately, at this point it still has to be done by ear, because we cannot predict the results of different seeds when irrational probabilities are used. At the same time, if all the inputs and probabilities are the same, z12 will output the same string of numbers. Puckette provides more detail on the mathematics behind z12 in his paper on maximally uniform sequences (Puckette 2015).

Cubic Zirconia uses 19 z12s. Each z12 uses one of ten different seeds and one of four different logic patches, but there is no duplicate combination of seed or logic. This creates nineteen independent texture streams with significantly different timbres.

I used four z12s created unaltered texture streams; 15 others were passed through four high-order band-pass filter banks with center frequencies characterized as low (minimum 20 Hz), mid-low (minimum 160 Hz), mid-high (minimum 640 Hz), and high (minimum 2560 Hz). Because the z12 creates highly complex waveforms, the output of

the filters can be quite different for a single z12. A $1/f$ random walk controls the center frequencies of the filters. At the start and end of the piece, the steps of the walk are very small, fixing the filters at a nearly constant set of frequencies starting at the minimum values. At the peak of the piece, the steps become quite large. Because it is a $1/f$ random walk, extreme frequency jumps are unusual, but because the smallest step is large, the texture streams spread out in the spectrum, opening spectral space in the texture. An exponentially distributed random variable controls the duration of the filter step, starting and ending with longer durations, but moving quickly at the peak. The result is that the timbre at the beginning and end of the work is static, while it moves, jumps, and changes frenetically at the apex. This process suggests that the listener is magnifying the texture and seeing greater detail over the course of the work.

The four filter outputs of each z12 are sent through independent delays starting at 0 sec (no delay), 10 sec, 25 sec, and 37 sec. Each of these, then, is sent through oscillating delays that adjust the delays up to 2.5 sec more. The different outputs of the filters, the long base delays, and the fluctuating smaller delays ensure that each texture stream appears to be an independent voice. This reduces the number of z12 processes needed for the final texture.

The overall length of the work is approximately 8.5 min. The four, unaltered texture streams start and end the piece, playing throughout the work. They act as the substrate on which the rest of the texture resides. The piece accumulates texture streams from the filter banks for the first 3 min. The full texture is then developed through the filter bank settings for approximately 5 minutes. In the last 20 sec, the texture streams dissipate. I took care to ensure the most consistent amount of spectral density, making sure that low, mid-low, mid-high, and high frequency texture streams were as present as possible.

Spatialization in the Cube

Given that the Cube has multiple elevated tiers and that at the time it could only present fixed medium works, my usual spatialization method would have required a few things that were not technically possible. First, I would have to determine the speaker angles necessary to make sure Catwalks 2 and 3 and the ceiling would get equivalent but different materials compared with the 64 speakers on the ground level. Second, all 19 z12s would have to run at the same time. Finally, all 124 channels would have to be recorded to a 124-channel AIFF, something that occasionally failed. Given time constraints, I could not determine which software was failing in reading AIFFs with very high track numbers.

I mapped the four unaltered z12 outputs to four loudspeakers at "compass points" at ground level: front center, left, right, and back center. These anchored the rest of the texture. Then I paired each of the remaining 60 ground speakers with one speaker as diametrically opposed as possible on the higher levels. For example, Speaker 1 at the right-front corner of the ground level was paired with Speaker 75, the left-back corner of Catwalk 2. I worked around the ground-level ring, pairing each speaker with one above and across the room. In this way, each texture stream had its own pair of loudspeakers: 15 z12s with four texture streams each, applied to one of 60 pairs. I then used a uniform random variable determining where between the two speakers a texture stream existed, with a simple cosine curve for amplitude panning. As in my original methods, I also applied a uniform random variable to control the wet/dry balance of a reverberation patch.

I planned the texture so that low, mid-low, mid-high, and high texture streams were as evenly distributed in space as possible, just as I made them evenly distributed in time. I recorded eight texture streams at a time (two z12s with four outputs each) to 16 channels in Pd. In Audacity, I arranged each output according to the channel assignment in the Cube. The full piece is provided as 124 mono AIFFs, a requirement for presenting in the Cube at the time.

Given the spectral consistency, the spatial consistency, and the overall form of the work, the piece achieved the aims of textural composition: highly mobile space; consistent, slowly changing

metaobject; and overall surround without a front-facing orientation. The high number of speakers makes the work clear and open. Surprisingly, the piece feels less dense but more extensive than in smaller arrays. For auditioning purposes, the 124 channels were mapped to the nearest, congruent speakers in SpADE, set to 28 speakers at the time, and the piece was closer and less penetrable. There is significantly less masking when texture streams have their own channels. This suggests that textural composition requires further thought when mapping from smaller HDLAs to larger HDLAs. In the same way that the portable eight-channel version of *Cubic Zirconia* uses considerably fewer texture streams, other HDLAs may require a different number of texture streams to maintain the experience of the piece. More investigation must be done to determine the consequences of scaling between different HDLAs for both the creation and experience of textural composition.

Conclusion

Textural composition is concerned with the sound metaobject, an idea derived from acousmatic aesthetics but transformed into real-time practice. It requires a few things to be true. First, the sound must be complex, consisting of many smaller sounds but bound into a whole object. Second, the sound must be large, surrounding the listener and extending beyond the perceivable horizon. Finally, the sound must not favor a perspective, and it must not have a frontal orientation.

The spatialization of the metaobject, therefore, must be fully encompassing. It must be quickly mobile but must not gel into coherent trajectories of its internal parts. Additionally, there needs to be an average consistency to the spatialization in all directions. This leads to a practice that balances between point-source spatialization techniques and mimetic aesthetics. My spatialization algorithm allows for many texture streams, which make up a single texture, to be spatialized independently, statistically equal, and in all directions. This method is scalable to any number of speakers with minor patching changes. This algorithm does not account

for elevated speakers, however. I devised a solution to this problem by providing false speaker locations that cause the sounds to move vertically as well as horizontally.

In composing for the Cube, I faced additional constraints. First, the piece had to be fixed. Second, there were three tiers of elevated speakers, including an array on the ceiling. I had to adjust my real-time procedures to survive in a fixed form. This allowed for other advantages, however—namely, processing capacity. Likewise, spatialization could not be done in real time, and my schemes for elevated speakers, i.e., false speaker locations, failed. Therefore, I had to more carefully conceive of spatialization that borrowed from point-source and mimetic procedures.

My methods of working in textural composition resulted in fairly equivalent impressions in 4- to 32-channel arrays. Nevertheless, the piece for 124 channels is drastically different when reduced to smaller HDLAs. This suggests that, although the spatialization may be scalable, material differences are also necessary when moving between significantly different speaker numbers. This is clearly obvious in other compositional practice, but not necessarily so in textural composition.

Further work in the Cube or similar places, as well as auditioning in SpADE (a smaller HDLA), could lead to basic compositional practices that would ensure a more consistent experience of material, even if spatialization were enhanced by greater speaker arrays. As more HDLA environments become available, additional work will be possible in different spaces.

References

Austin, L. 2000. "Sound Diffusion in Composition and Performance: An Interview with Denis Smalley." *Computer Music Journal* 24(2):10–21.

Austin, L. 2001. "Sound Diffusion in Composition and Performance Practice II: An Interview with Ambrose Field." *Computer Music Journal* 25(4):21–30.

Barrett, N. 2002. "Spatio-Musical Composition Strategies." *Organised Sound* 7(3):313–323.

Burns, C. 2006. "Compositional Strategies for Point-Source Spatialization." *eContact! 8.3.* Available online

at cec.concordia.ca/econtact/8_3/burns.html. Accessed 16 September 2016.

Cage, J. 1958. "The Future of Music: Credo." *Silence: Lectures and Writings*. London: Marion Boyars, pp. 3–6.

Cascone, K. 2000. "The Aesthetics of Failure: 'Post-Digital' Tendencies in Contemporary Computer Music." *Computer Music Journal* 24(4):12–18.

Cowell, H. (1929) 2004. "The Joys of Noise." *The New Republic*, 31 July, p. 287. Reprint C. Cox and D. Warner, eds. *Audio Culture: Readings in Modern Music*. New York City: Continuum, pp. 22–24.

Hagan, K. 2008a. "Textural Composition: Implementation of an Intermediary Aesthetic." In *Proceedings of the International Computer Music Conference*, pp. 509–514.

Hagan, K. 2008b. "Textural Composition and its Space." In *Proceedings of the Sound and Music Computing Conference*, pp. 100–106.

Hagan, K. 2012. "Aesthetic and Philosophical Implications of Noise in *Morphons and Bions* (2011)." In *Proceedings of the Music, Mind, and Invention Workshop*. Available online at www.tcnj.edu/~mmi/papers/Paper20.pdf. Accessed November 2016.

Hagan, K. 2013. "The Noise of *Morphons and Bions*." In *Proceedings of the International Computer Music Conference*, pp. 376–379.

Harrison, J. 1998. "Sound, Space, Sculpture: Some Thoughts on the 'What', 'How' and 'Why' of Sound Diffusion." *Organised Sound* 3(2):117–127.

Kahn, D. 1999. *Noise Water Meat*. Cambridge, Massachusetts: MIT Press.

Leider, C. 2007. "Multichannel Audio in Electroacoustic Music." In *Proceedings of the IEEE International Conference on Multimedia and Expo*, pp. 1890–1893.

Lyon, E., et al. 2016. "Genesis of the Cube: The Design and Deployment of an HDLA-Based Performance and Research Facility." *Computer Music Journal* 40(4):62–78.

Normandeau, R. 2009. "Timbre Spatialisation: The Medium Is the Space." *Organised Sound* 14(3):277–285.

Otondo, F. 2008. "Contemporary Trends in the Use of Space in Electroacoustic Music." *Organised Sound* 13(1):77–81.

Puckette, M. 2015. "Maximally Uniform Sequences from Stochastic Processes." Paper presented at the Conference of the Society for Electro-Acoustic Music in the United States, 26–28 March, Blacksburg, Virginia.

Pulkki, V. 1997. "Virtual Sound Source Positioning Using Vector Base Amplitude Panning." *Journal of the Audio Engineering Society* 5(6):456–466.

Rolfe, C. 1999. "A Practical Guide to Diffusion." *eContact! 2.4*. Available online at econtact.ca/2_4/pracdiff.htm. Accessed 29 September 2016.

Sazdov, R., G. Paine, and K. Stevens. 2007. "Perceptual Investigation into Envelopment, Spatial Clarity, and Engulfment in Reproduced Multi-Channel Audio." In *Proceedings of the 31st Audio Engineering Society International Conference*. Available online at www.aes.org/e-lib/browse.cfm?elib=13961 (subscription required). Accessed 9 November 2016.

Schafer, R. M. 1994. *Our Sonic Environment and the Soundscape: The Tuning of the World*. Rochester, Vermont: Destiny.

Smalley, D. 1997. "Spectromorphology: Explaining Sound-Shapes." *Organised Sound* 2(2):107–126.

Smalley, D. 2007. "Space-Form and the Acousmatic Image." *Organised Sound* 12(1):35–58.

Trochimczyk, M. 2001. "From Circles to Nets: On the Signification of Spatial Sound Imagery in New Music." *Computer Music Journal* 25(4):39–56.

Truax, B. 1998. "Composition and Diffusion: Space in Sound in Space." *Organised Sound* 3(2):141–146.

Xenakis, I. 1992. *Formalized Music: Thought and Mathematics in Composition*. Revised ed. Stuyvesant, New York: Pendragon Press.