FROM GRAINS TO FORMS

Curtis Roads
Media Arts and Technology
University of California, Santa Barbara
clang@mat.ucsb.edu

In Makis Solomos, (ed.), Proceedings of the international Symposium *Xenakis. La musique électroacoustique / Xenakis. The electroacoustic music* (université Paris 8, May 2012).

ABSTRACT

This paper explores strategies for organizing sound grains into larger structures on the meso and macro time scales. We begin with a look at clouds, streams, sprays, and Xenakis's original proposal based on a sequence of screens. Then we examine tools for montage and micro-montage in the studio, as well as instruments for gestural control. This leads to a discussion of higher-order granulation and per-grain processing. Dictionary-based pursuit is an analytical counterpart to granular synthesis, enabling transformations that derive new sounds from analyzed sounds. Next we look at proposals for generating grains based on physical and biological models. The next section looks at abstract generative models based on algorithms imported from mathematics. This leads into a general discussion on the limitations of formalism in music composition. The core of the paper concerns the problem of multiscale organization, which can

resolve the tension between formal and informal approaches through a combination of heuristic algorithms and direct interventions.

ENCOUNTER WITH XENAKIS

I first met Maestro Iannis Xenakis at his final course Formalized and Automated Music at Indiana University in 1972 (Figure 1).

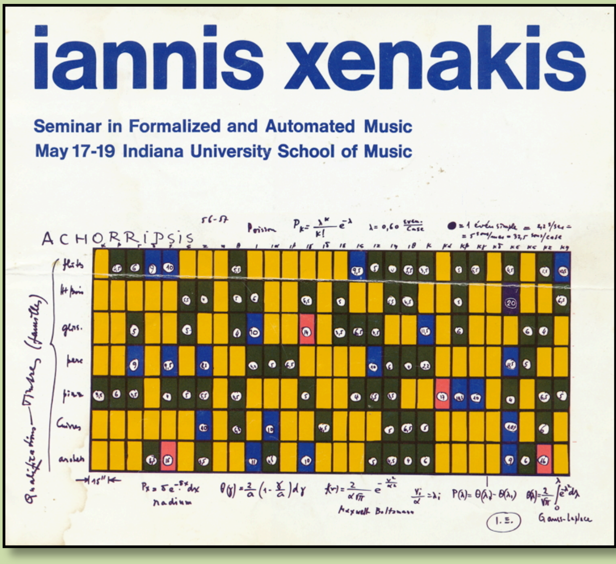

Figure 1. Seminar in Formalized and Automated Music, 1972.

I was deeply inspired by Xenakis's original musical discourse and free imagination. (1) Yet it was evident that Xenakis was not happy in Bloomington (Logan 2000). We heard that the administration was not maintaining the expensive hardware needed for computer sound synthesis experiments. Not surprisingly, soon Xenakis returned permanently to Paris.

In the following months, I studied *Formalized Music* and developed computer programs to test the formulas in the book (Roads 1973). In this period, sound synthesis by computer required specialized facilities. I found these at the University of California, San Diego (UCSD). There I developed a variant of Xenakis's theory of granular synthesis, based on a model of scattering grains into *sound clouds* (Roads 1978). I did not expect that twenty years later I would be working in Paris at Les Ateliers UPIC (later renamed CCMIX and finally CIX) alongside Xenakis and his équipe at the CEMaMu, developing another program for granular synthesis called Cloud Generator (Roads and Alexander 1995) (Figure 2).

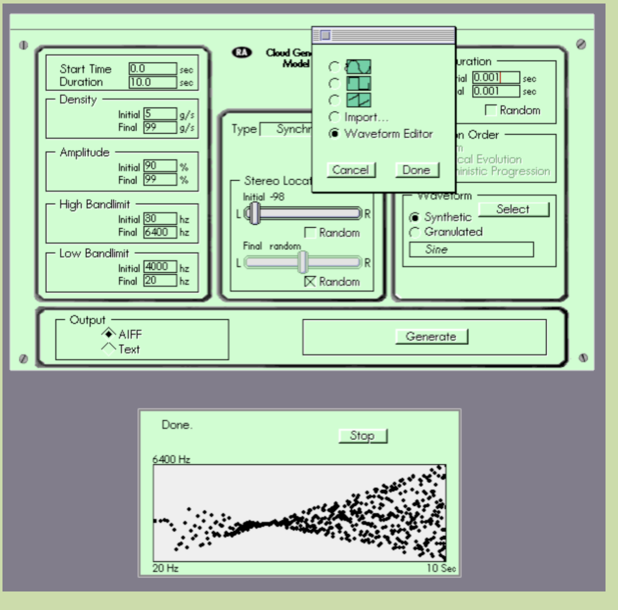

Figure 2. The Cloud Generator app in action 1995.

Since this time, the granular paradigm has proven to be one of the most powerful methods of synthesis and sound transformation, implemented in dozens of incarnations and used by innumerable musicians.

Granular synthesis requires an algorithmic model of grain generation and organization. In this paper, I explore various granular models and how they lead to higher-level musical structures. Of course, as Jean-Claude Risset (2005) observed, a characteristic of the granular paradigm is that it blurs the border between microstructure and macrostructural organization:

> *By bridging gaps between traditionally disconnected spheres like material and structure, or vocabulary and grammar, software creates a continuum between microstructure and macrostructure. It is no longer necessary to maintain traditional distinctions between an area exclusive to sound production and another devoted to structural manipulation on a larger temporal level. The choice of granulation, or of the fragmenting of sound elements, is a way of avoiding mishaps on a slippery continuum: it permits the sorting of elements within a scale while it allows individual elements to be grasped. The formal concern extends right into the microstructure, lodging itself within the sound grain.*

SCREEN, CLOUD, STREAM, AND SPRAY MODELS

Xenakis first demonstrated granular synthesis in his composition *Analogique B* (1959). This was realized by recording sine tones on tape, cutting the tapes into hundreds of tiny pieces, and then recombining them by manual splicing according to a stochastically-generated score.

Xenakis's design for digital granular synthesis, as articulated in the chapter "Markovian Stochastic Music–Theory" was based on generating a sequence of time-frequency screens, akin to frames of a film, running at a fixed rate (Xenakis 1971). He proposed to control the sonic energy distributed on each screen through set theory operations such as "this screen is the set union of two preexisting screens." The sequence of screens was to be controlled by a Markov chain process. In this design, the composer would control the synthesis by means of operations that are not directly tied to acoustical parameters such as pitch, duration, amplitude, spectrum, and spatial position (Xenakis 1971, chapter IX). In this sense, Xenakis's design for granular synthesis was conceptually similar to his Dynamic Stochastic Synthesis or GENDY developed in the early 1990s, in which sound production is an epiphenomenon of an abstract generative process taking place on the subsymbolic time scale of individual samples.

By contrast, software-based granulators today implement a flowing spray jet of sound particles. Spray tools inevitably shape higher levels of a composition; meso and macro forms tend toward patterns of stream and cloud formations on the time-frequency plane (Roads 1978; Truax 1986; Solomos 2006).

This is made visually explicit in programs like Metasynth (Wenger and Spiegel 2004) where the user literally sprays grains onto a time-frequency grid onscreen (Figure 3). In this sense, the cloud, stream, and spray models align with the spirit of Xenakis's UPIC paradigm, in which users directly draw sounds in the time-frequency plane (Xenakis 1992).

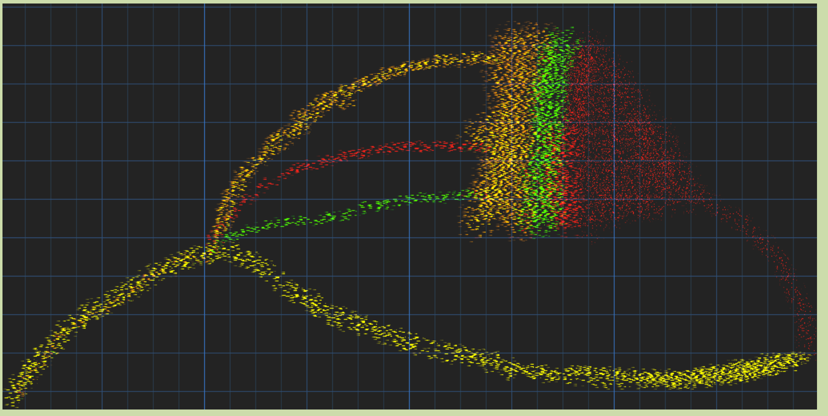

Figure 3. Screen image of MetaSynth granular sprays.

Within the stream and cloud models, the distinction between *synchronous* and *asynchronous* granular synthesis is compositionally pertinent (Roads 2001b). *Synchronous granular synthesis* (SGS) emits one or more streams of grains where the grains follow each other at regular intervals. A prime use for SGS is to generate metric rhythms keeping the grain emissions sparse per unit of time.

One of the most important parameters of granular synthesis is grain *density*–the number of grains per second. In the case of SGS, this corresponds to a regular frequency of grain emission. For example, a density of 2 grain/second indicates that a grain is produced every half second– a repeating “beep.” Synchronous densities in the range of about 0.1 and 20 grains per second generate metrical rhythms. When the densities change over time, we experience precise accelerandi/rallentandi effects. At higher densities, long grains fuse into continuous tones. Here is the sweeter side of granular synthesis, since these tones tend to have a strong fundamental frequency. Depending on the grain envelope and duration, these tones will also manifest sidebands.

Any granulator can generate a huge amount of derived sound material from a given sound file by manipulating the position, speed, and direction of the read pointer. A slow backwards-scan granulation of a sound file changes the identity of the original, especially when this is combined with pitch shifting, filtering, and spatialization, all randomized on a grain-by-grain basis. (See the discussion of per-grain transformations in the next section.) Reducing the duration of the grains has the inevitable effect of churning any source into broadband noise.

Asynchronous granular synthesis (AGS) abandons the concept of perfectly sequential streams of grains. Instead, it scatters the grains over a specified duration within regions inscribed on the time-frequency plane. These regions are clouds –the units with which a composer works. The scattering of the grains is irregular in time, being controlled by a stochastic algorithm (Roads 1991). In the case of AGS, grain density corresponds to the degree of transparency or opacity of the sonic fabric.

STUDIO-BASED MONTAGE AND MICRO-MONTAGE

The output of the Cloud Generator app is a single cloud of grains. The idea is that the composer can generate a number of such clouds in the studio and then organize them using a timeline-based mixing app like Pro Tools.

This approach is extremely free in terms of compositional options, considering that material at any time scale can be processed by plugins and placed anywhere in the time line. Detached from real-time constraints, ideas can be tested, edited, submixed, or deleted at will.

One of the strategies that this enables is a *thematic approach,* i.e., manipulating sound material so as to create repetitions and variations. For example, a copy of a given sonic entity can be pitch-shifted, time-scaled, ring modulated, filtered, reversed, etc. Parts of the piece, if not the entire work, can be organized as a montage of variations of a finite number of elements. This approach forms the basis of my compositions *Never* (2010) and *Always* (in progress) and all of the electroacoustic works of Horacio Vaggione. It appears in his concept of *micro-figures* and was crystallized by the IRIN app developed by Vaggione's student Carlos Caires (Figure 4).

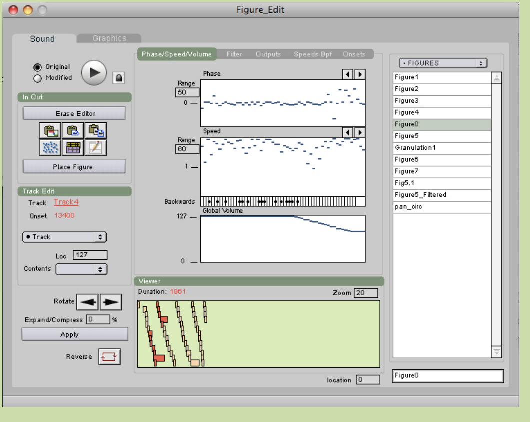

Figure 4. IRIN figure editor.

GESTURAL CONTROL OF GRANULAR INSTRUMENTS

The first generation of granulators operated in non-real time (Roads 1978). Truax (1986) developed real-time implementations of granular synthesis, but these required specialized hardware. By contrast, I conceived the real-time Creatovox instrument (1999-2000) as having a software-based engine written in SuperCollider with common MIDI controllers as the performance interface (Roads 2001b).

In the early 1990s, I began a design notebook containing schemes for a real-time granular synthesis instrument, including both scheduling algorithms as well as protocols for gestural interaction (Roads 1992-1997, 1998). This instrument would look like a typical keyboard instrument, with the additional of a few special controllers. It could be compatible with the MIDI protocol. As a guiding principle, I wanted to convey "an illusion of simplicity" to the performing musician by hiding unnecessary technical details.

By 1998, common microprocessors became fast enough to handle multiple voices of granular synthesis in real time. At the same time, James McCartney released the SuperCollider 2 language (McCartney 1998). SuperCollider 2 combined musical interaction with synthesis and sound processing, so it seemed ideal for prototyping this project. On the basis of these developments, I launched the Creatovox research project at the Center for Research in Electronic Art Technology (CREATE) at UCSB. The goal of the Creatovox project was to invent a prototype instrument optimized for expressive virtuoso performance of granular synthesis based on the cloud/stream paradigm.

The Creatovox (Figure 5) emitted its birth cries on 13 July 1999 in the presence of a project team consisting of Alberto de Campo, Ching-Wei

Chen, and me. As of January 2000, it had gone through another round of development and produced octophonic output.

Figure 5. The Creatovox instrument, January 2000. Here I play two simultaneous clouds: one in the bass register (the bass pedal) and one in the high register (keyboard). The left hand manipulates three parameters via a 3D joystick: grain density, grain duration, and amount of reverberation for example, while the right foot controls volume.

A funny thing happened at this point. Every time I went out in public to demonstrate the Creatovox, the results disappointed me. It was not the fault of our design, it had to do with the fact that I had not invested a great deal of time to practice playing this instrument. What a surprise: a virtuoso instrument requires a virtuoso performer who practices the instrument every day! Ultimately I decided that becoming a virtuoso performer was not central to my interests in composition. One composer who did use the Creatovox was the wonderful Bebe Barron (the soundtrack of *Forbidden Planet*), who made her final piece *Mixed Emotions* in the summer of 2000. However, in this case the Creatovox was used to generate fragments that were ultimately assembled in Pro Tools. (2)

Since then there have been many real-time implementations of granular synthesis, some with more performance potential than others. A current trend favors a graphical "scrubbing" model, after the tape scrubbing methods of the analog era (Figure 6).

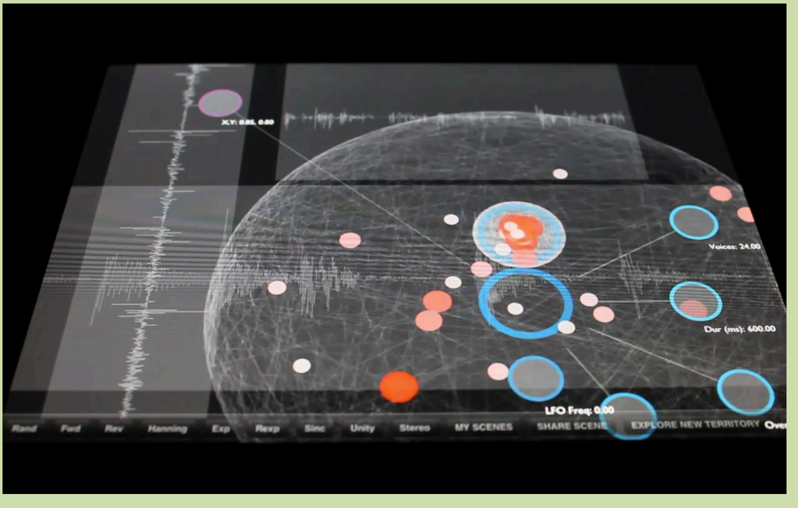

Figure 6. Borderlands app for iPad. The user slides ("scrubs") a circle along the waveform to select segments to granulate. The attached satellite circles function as potentiometers controlling granular synthesis parameters. Another recent scrubbing app is MegaCurtis.

Of course, the challenge of making a satisfying piece with a rich multiscale architecture purely from a live performance on a granular instrument remains a daunting one. It is not impossible, but will require considerable practice…

ENVELOPE CONTROL OF PULSAR PATTERNS

In contrast to the Creatovox, which was designed for virtuoso performance, our PulsarGenerator app (Figure 7), coded by Alberto de Campo, was designed to allow a composer plan detailed pulsar patterns (pulsars are similar to grains) in the studio by means of a set of 12 envelopes with novel envelope-based operations like envelope mixing, inversion, scaling, reversal, and smoothing, (Roads 2001a, 2001b).

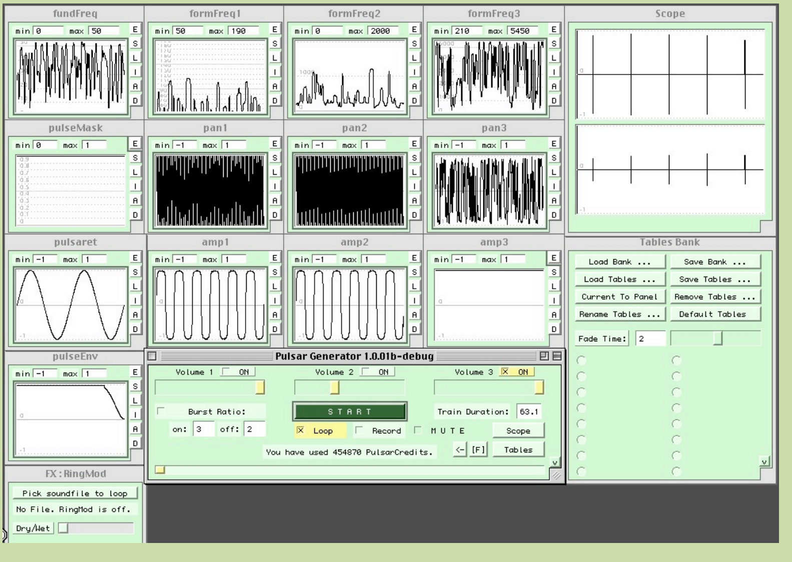

Figure 7. The PulsarGenerator screen focused on envelope control of particle synthesis.

HIGHER-ORDER GRANULATION AND PER-GRAIN PROCESSING

Recycling sounds by means of *higher-order granulation* is a method of spawning new granular mesostructures out of old ones. In effect, we re-granulate one or more existing granulated sound files. Depending on the capabilities of the granulation algorithm, a wide range of variations can be generated. The resulting sounds can be many times the duration of the original input sound. For example, a single stream of granulation using large grains and a sharp attack envelope breaks a continuous stream into discrete chunks. If this granulation stream has a wide range of amplitude variations, each chunk will have its own dynamic articulation, creating articulated phrases.

The technical capability of regranulation was available in my 1988 Granulate program, which could granulate up to 64 sound files at a time to create new hybrid sounds out of existing sounds (Roads 2001b). However, I did not begin to experiment with higher-order granulation until 2003 with the realization of *Now* (2004), a regranulation of my composition *Volt air* (2003). In turn, *Never* (2010) was a third-order granulation of *Now.* I am currently working on a fourth-order granulation of *Now* in the work-in-progress *Always.*

These experiments rely on my Constant-Q Granulator (Roads 1998, 2001b) and the EmissionControl app (Thall 2004a,b; Roads 2006) software. (3). EmissionControl (Figure 8) is particularly interesting as it implements a *matrix modulation* scheme for control of synthesis parameters. In particular, parameters can be controlled by low-frequency oscillators (LFO) and random functions as well as by interactive sliders directly manipulated by the composer. LFO control lends itself to the generation of phrase structures based on cyclical patterns.

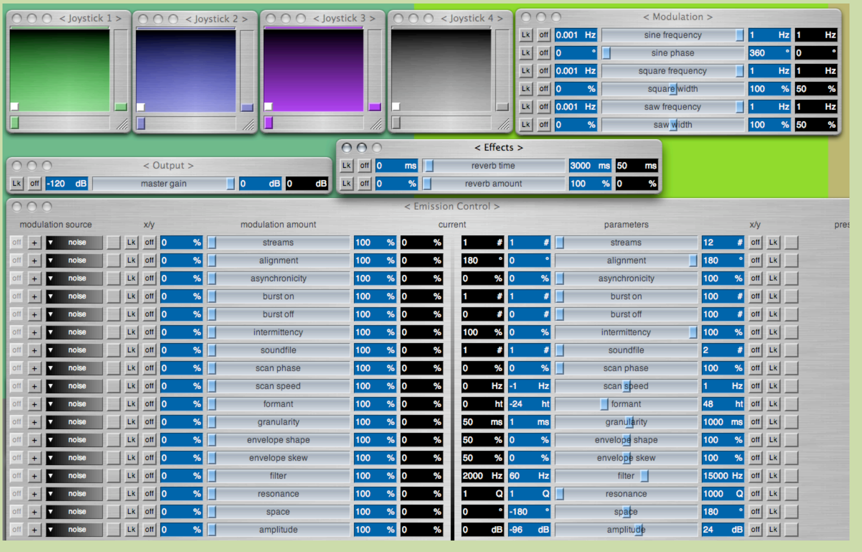

Figure 8. Screen of EmissionControl. The faders on the left determine the amount of modulation of the parameters in the right side. Modulation sources include LFOs and random generators. The joystick controllers at the top let the user manipulate two parameters simultaneously.

I should mention here an essential feature that is characteristic of all my granulators since 1988, and which amazingly seems to be missing in other implementations: *per-grain effects processing*. Starting in 1988 I wrote programs that scattered each grain emitted to an individualized point in virtual space. I later I extended this so that every grain passes through a separate constant-Q filter. Each filter has its own center frequency and bandwidth, selected randomly within limits stipulated by the user. The number of filters in operation corresponds to the density of grains per second, which can be hundreds per second. In a similar per-grain fashion we can pitch-shift, ring-modulate, etc. each grain individually with different parameter settings for each grain. The resulting heterogeneity of sound is the signature of truly granular signal processing. By comparison, many granulators that feed the entire grain stream through the same effects channel tend to sound flat and one-dimensional.

DICTIONARY-BASED PURSUIT: ANALYTICAL COUNTERPART TO GRANULAR SYNTHESIS

The granular paradigm is, of course, not limited to stream, cloud, and spray models. Indeed, it applies to the analysis or synthesis of any sound, which Gabor's pioneering papers demonstrated (1946, 1947, 1952). Indeed, as Xenakis (1960) wrote:

> *All sound is an integration of corpuscules, of elementary acoustic particles, of sound quanta.*

Thus we can extend the granular paradigm to the realm of sound analysis by means of an analytic counterpart to granular synthesis: *dictionary-based pursuit* (DBP), also known as the family of *matching pursuit* algorithms (Sturm, et al. 2009). DBP seeks by iterative search to match the energy in a signal with a vast dictionary of millions of sound atoms or grains, proceeding step-by-step from the strongest unit of time-frequency energy to the weakest.

The result of a DBP analysis is a time-frequency representation of grains that is highly malleable. This opens up a largely unexplored universe of potential sound transformations based on granular processes, including pluriphonic spatial processing at the granular level, where a sound's elemental particles scatter in evanescent spatial patterns, or cavitation processes carve new granular patterns out of existing ones.

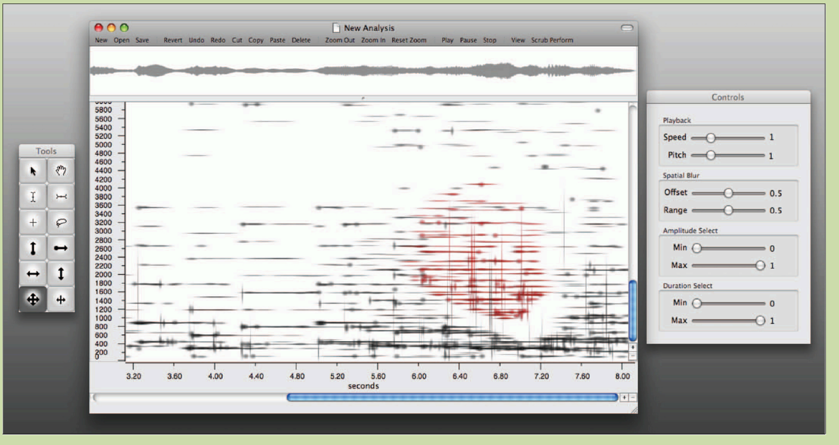

Figure 9. Screen image of Scatter, a graphic interface for manipulating granular time-frequency representations. The vertical axis represents frequency and the horizonal axis represents time. The area in red has been picked up and moved from its original location on the time-frequency plane.

The two primary code libraries for the analysis stage of DBP, Matching Pursuit Toolkit (MPTK) (Krstulovic and Gribonval 2006) and LastWave (Bacry 2008) were not designed for musical purposes. Our Scatter app (Figure 9) was designed precisely to take the analysis data generated by MPTK and make it artistically usable. Scatter provides a graphic user interface for performing dictionary-based pursuit analysis (McLeran et al.

2008). The app displays the Wigner-Ville distribution and energy contribution of each *atom* or grain of sound. (See Preis and Georgopoulos 1999 for more on the Wigner-Ville distribution.) The program offers multiple ways of transforming the decomposition: by direct interaction via mouse, by parametric filtering, and by application of stochastic algorithms. Scatter also enables real-time performance through a variety of playback and scrubbing techniques. For a video demonstration see note (4).

PHYSICAL AND BIOLOGICAL MODELS

A potentially compelling paradigm for granular organization derives from the world of physics. Physical modeling synthesis starts from a mathematical description of a mechanical process (Roads 1996; Fletcher and Rossing 1991). By now, a large body of scientific literature centers on the micromechanics of granular processes, such as grain vibration patterns, mixing, flow, and grain/fluid interactions (Aronson and Tsimring 2009). The basic premise of these simulations is that the rules governing the behavior of one thing are not the same as those governing a million.

An early example of physical model control of granular synthesis was a simulation of the sounds of shaken and scraped percussion: maracas, sekere, cabasa, bamboo windchime, tambourines, sleighbells, and the guiro (Cook 1996, 1997, 2007). Keller and Truax (1998) created granular models of natural physical processes such as liquid streams and a bouncing metallic ball. Natasha Barrett modeled processes of *self-organized critical systems* such as grain avalanches in her work *The Utility of Space* (2000), which was just one facet of a multilayered acousmatic design.

Such research could be taken much further. In recent years, physicists have built sophisticated physical models of a broad range of granular phenomena. These include ordered patterns such as ripples, avalanches, and bands of segregated materials. Another class of patterns emerge out of disordered excitations such as sifting, shaking, and scattering (Bideau and Hansen 1993). Many other phenomena are characterized by *clustering* on a variety of scales (Rivier 1993; Reynolds 1993).

In similar manner, scientists study self-organizing patterns and behavior in the biological world of social insects, schools of fish, and swarms of birds (Camarzine et al. 2001). Today the scientific modeling paradigm seems ripe for exploration by both artists and scientists. These simulations can be arbitrarily complex, and extend into virtual worlds in which events that would be impossible in the real world can be simulated as easily as actual events.

Unfortunately, the problem with physical/biological modeling has always been the same: a model of an instrument/system is worthless without an expert player, whether real or virtual. Virtuosity demands daily practice over a period of years. Obviously, to develop a software model of an expert player poses a daunting challenge. This goes beyond the question of how a performer interprets a score. In physical/biological models there is the need to control dozens or hundreds of low-level parameters in real time in order to simulate a sequence of related gestures in a "natural sounding" way. How does one "play" a physical/biological simulation? Of course, one could give up the "natural sounding" gesture constraint, but then we are left with a robot player whose gestures may not make much sense to human beings.

ABSTRACT GENERATIVE MODELS

Beyond the physical and biological lies the abstract realm of experimental generative algorithms for granular synthesis, not related to a model of a real-world phenomenon. They are sometimes derived from formulae or algorithms that are *en vogue* in a given period, such as recursive substitution, cellular automata, fractal or chaotic functions. For example, Alberto de Campo (1998) proposed a method of grain scattering in time based on a recursive substitution algorithm. Others have employed chaotic functions to scatter the grains in time (Moon 1987). Chaotic functions vacillate between stable and unstable states, between intermittent transients and turbulence (Di Scipio 1990; Gogins 1991, 1995; Miranda 1998; Wolfram 2002).

These generally low-level algorithms emit granular sound as a byproduct of a formal process that is imported from an abstract domain. A fundamental problem with such bottom-up approaches is the fact they do not incorporate a multiscale notion of structure. Like the low-level rules in Wolfram's (2002) experiments, whether larger scale structures emerge out of such rules is mainly a matter of guesswork; in many cases they do not.

As a compositional strategy, reliance on enumeration of a formula is constraining, a classic "inside the box" approach. In effect, the composer invites us to listen to her clockworks. However, the human fascination for monitoring a mechanism is limited. Abstract algorithms are not tethered to human action or perception. To the detached observer who is not invested in the algorithm itself, these processes tend to either sound either predictable or merely random. They lack the meaningful sense of causality and effort that we associate with human gestures, the opposition between tension and resolution to which we are psychologically disposed. Human beings are highly sensitive to the virtuosic play of lively human gestures, whether live on stage or practiced behind the scenes in a studio, and to the unfolding of a sonic narrative played out by recognizable sonic characters or themes.

LIMITATIONS OF FORMALISM IN MUSIC COMPOSITION

We see here a broader philosophical issue opening up: the limitations of formalism in music composition. As Jean-Claude Risset (2007) observed, the opposition between *music-as-formal-concept* versus *music-as-perceived-sound* is as ancient as Pythagoras (the formalist) and Aristoxenus (the listener). Applied at different strata of compositional organization, formal algorithms can be a potent means of invention. For example, granular synthesis proves that the micro layers of sound generation can be controlled effectively by powerful algorithms that spawn masses of fine detail.

Compositional algorithms can also enumerate a collection of variations quickly, offering the composer a wide range of selections from which to choose. Interactive performance systems try to balance preprogrammed automation with spontaneous decisions.

However, at the same time, we see many examples of generative systems that produce unlimited quantities of unremarkable music. (5)

While formal algorithms enable interaction with a machine, absolute formalism in composition means imposing constraints on one's self. The strictly formalist composer follows a predetermined conceptual plan from beginning to end. The plan must ultimately be translated into the real world: acoustics, psychoacoustics, and emotional response. It is in this translation that the game is often lost.

Musical formalism is related to but not identical to the conceptual art tradition. In a famous essay on conceptual art, the artist Sol Lewitt (1967) wrote:

> *In conceptual art the idea or concept is the most important aspect of the work. When an artist uses a conceptual form of art, it means that all of the planning and decisions are made beforehand and the execution is a perfunctory affair. The idea becomes a machine that makes the art.*

Rather than critically evaluating the result on a perceptual basis, another set of aesthetic criteria govern such art: is the concept politically, culturally, and socially meaningful? How does it reflect on, comment on, or critique society? Many such installations take in audience feedback and thus function like polls, surveys, or marketing exercises.

Music has a long conceptual tradition, but composition tends to gravitate toward mathematical and algorithmic strategies rather than cultural probes.

Many algorithmic composers are motivated by formalist ideals of conceptual purity, logical rigor, and algorithmic beauty. Generative strategies are conceptually attractive. They offer the possibility of exploring novel musical processes and the creation of new musical structures from freely invented axioms. Another often stated rationale for generative techniques is that they allow composers to reach beyond themselves. These romantic ideals favor self-transcendence, purity of method, and often incorporate seductive metaphors of natural growth and evolution, artificial life, and holistic ecosystems.

However, elegant rules do not necessarily make elegant music. As a skeptical Debussy observed (quoted in Risset 2004):

> *Works of art make rules, rules do not make works of art.*

In certain masterworks, the composer escapes from the cage of self-imposed rule systems. For example, Xenakis treated the output of his composition programs flexibly. In particular, he edited, rearranged, and refined the raw data emitted by his Free Stochastic Music program (Roads 1973).

> *When I used programs to produce music like ST/4, ST/10, or ST/48, the output sometimes lacked interest. So I had to change [it]. I reserved that freedom for myself. Other composers, like Barbaud, have acted differently. He did some programs using serial principles and declared: "The machine gave me that so I have to respect it." This is totally wrong, because it was he who gave the machine the rule!* – Iannis Xenakis (Varga 1996).

> *"Formalized music" does not sound free, but it is. I wanted to achieve a general musical landscape with many elements, not all of which were formally derived from one another.* – Iannis Xenakis (1996).

In his later years, Xenakis no longer relied on computers for instrumental composition; he had absorbed algorithmic strategies into his intuition (Varga 1996; Harley 2004).

THE PROBLEM OF MULTISCALE ORGANIZATION

We return now to the main theme of this paper: the notion of organizing sound grains into larger structures. Cloud, stream, and spray models have been effective in composition because they agglomerate grains into multiple levels of musical structure, specifically the *object* (100 ms-8 sec) and *meso* (> 8 sec) time scales (Roads 2001b).

By contrast, physical/biological models and abstract algorithms of grain generation pose puzzling challenges. How can one create coherent multiscale structures with these techniques? As Wesley Smith (2011) observed:

> *One of the major challenges in building a system that can increase in complexity as it runs is figuring out how to transfer complex structures in a lower level space into simple structures in a higher level space while still maintaining the essential qualities that the complex lower level structure represents.*

This is one of the great unsolved problems in algorithmic composition. The issue is not merely a question of scale, i.e., of creating larger sound objects out of grains. It is a question of creating coherent *multiscale behavior* extending all the way to the meso and macro time scales. Multiscale behavior means that long-term high-level forces are as powerful as short-term low-level processes. This is why simplistic bottom-up strategies of "emergent self-organization" tend to fall short.

Xenakis speculated on abstract schemes for multiscale organization on several occasions. In *Formalized Music* (1971) he conjectured on a kind of multiscale clustering applicable to granular synthesis, where each point in a cluster contains a sub-cluster (6):

> *Within human limits, using all sorts of manipulations of these grain clusters, we can hope to produce not only the sound of the classical instruments and elastic bodies, and those sounds generally preferred in concrete music but also sonic perturbations with evolutions, unparalleled and unimaginable until now…We can even express a more general supposition. Suppose that each point of these clusters represents not only a pure frequency and its satellite intensity, but an already present structure of elementary grains, ordered a priori. We believe that in this way a sonority of a second, third, or higher order can be produced.*

In "Music composition treks" (1985) Xenakis enumerated nine proposals; six of these involved hierarchical schemes in which one stochastic process determines macroform, while another low-level stochastic process calculates the actual sounds.

> *It seems that a new kind of musician is necessary, that of the artist-conceiver of free and abstract new forms, tending toward complications and generalizations at several levels of sound organization. For example, a form, a construction, an organization built on Markovian chains or on a complex of interlocked probability functions may be carried over simultaneously onto several levels of musical micro-, meso-, and macrocompositions.*

His remaining proposals dealt with logical operations (Cartesian products of sets, sieve theory, logical functions applied to sets of sound parameters), which were previously used by the composer.

Of course, the problem of multiscale granular organization is not merely a matter of building arbitrary clusters. For the music to be convincing, these formations must articulate a compelling process that is essentially narrative in function.

Multiscale composition with physical/biological models poses many challenges. Most physical models describe a single instrument and are gesture-driven. As we have already pointed out, a physical model is worthless without an expert player of gestures. Even then, a player needs to know what gestures to play. The challenge is to build game-like virtual worlds in which a composer can interact with modeled instruments/biosystems in such a way as to construct a meaningful narrative, expressed as a pattern of gestures. For example, one can imagine a kind of musical pinball machine supplied with thousands of grains, where the levers are controlled by the composer. Certain levers could model physical forces acting on the granular flows.

FORMAL/INFORMAL STRATEGIES IN HEURISTIC ALGORITHMS

Music interacts in deep ways with the memory and expectations of listeners (Huron 2006). Human beings respond intuitively to context-dependent cognitive impressions that are difficult to formalize, like wit, irony, tension, surprise, virtuosity, humor, and clever twists and transitions. One sound appears to cause another sound. Sounds converge on points of attraction or scatter at points of repulsion. In general, formal methods do not address these issues. How to codify these essentially narrative functions based on human expectation? This is not obvious, but it is clear that it is unlikely to emerge from a formula borrowed from an arbitrary branch of physics or mathematics.

The effect of music is essentially magical–beyond logical explanation. In her fascinating book, *Music, Science, and Natural Magic in Seventeenth Century England,* Penelope Gouk (1999) quotes the English philosopher Francis Bacon (1561-1626), who defined magic as follows:

> *The science which applies the knowledge of hidden forms to the production of wonderful operations.*

In the context of the 17th century, the “knowledge of hidden forms” involved mastery of esoteric skills–analogous to the knowledge of programming languages today–but also how to apply them “to the production of wondrous operations.” Criteria for the production of wonder have never been formalized. However, we observe that certain highly talented people make inspired choices from myriad possibilities to create fascinating designs. This remains the strong suit of human talent.

For grains to evolve into wondrous forms, it appears that what is needed is a hybrid formal/informal approach, combining the computational power of algorithmic control with the magical influence of heuristics. What is heuristic influence? Heuristics is the art of experience-based strategies for problem solving:

> *Heuristic knowledge is judgmental knowledge, the knowledge that comes from experience–the rules that make up “the art of good guessing.”* (Cohen et al. 1984).

The master chess player uses heuristic insight, in contrast to a computer chess player that must employ brute-force search on millions of possibilities at each move without any knowledge of previous moves. Heuristic methods include rules of thumb, educated guesses, intuitive judgments, and common sense. Heuristic methods are inevitably intertwined with an understanding of context, whether it be the state of a game, or the state of a composition.

Heuristic methods are compatible with formalization. However, in practice they implement tailor-made solutions that are domain-specific and context-dependent, rather than imported whole cloth from one area of study to another. For example, the visual artist Harold Cohen has long applied heuristic algorithms to aesthetic problems. Over a period of forty years, Cohen has been developing a body of highly specific algorithms for drawing and coloring shapes. These algorithms do not attempt to draw like a human being would.

Most importantly, heuristic algorithms are tested by experiments and refined by human perceptual judgments. Xenakis used stochastic processes in a heuristic manner, sometimes modifying and rearranging the results to better suit the piece. Poetic license is the ultimate heuristic.

For granular synthesis, one heuristic approach would be to borrow certain concepts from scientific models of granular processes but then rework them to serve more effectively in a musical context. In this sense the physical model serves as a kind of metaphor for granular organization in music, rather than a strict model. The main point is that we can design heuristic algorithms to implement methods of mesostructural formation and multiscale behavior. These algorithms need to “work” according to testing and expert judgment.

The principle of economy of selection

Hand-in-hand with the use of heuristic algorithms is one of the most important issues in composition: the principle of *economy of selection* (Roads forthcoming). Economy of selection means choosing one or a few perceptually and aesthetically optimal or salient choices from a vast desert of unremarkable possibilities. This choice relies on the powerful aesthetic perception of an expert practitioner.

Making the inspired, intuitive choice from myriad possibilities remains the exclusive domain of human talent. As Stuckenschmidt (1970) observed:

> *Bach was as well versed in the possible uses of the three mirror forms of a melody as any Netherlands polyphonist of the fifteenth or sixteen century. He did not omit to use one or another of them out of forgetfulness or a defective grasp of the full range of possibilities. He knew that a two-part invention can occupy only a limited amount of space. The ability to make the right choice from the million or more possible forms is a creative secret that cannot be uncovered by science or technology. Here, too, is where the astonishing capabilities of computers prove to have limitations.*

Long seen as a gift of the gods, inspired choice seems difficult to teach to human beings and even more so to computers. Indeed, what makes a choice inspired is hard to generalize, as it is particular to its context. Sometimes it is the surprising or atypical choice, but other times it is simply satisfying, optimal, or salient in a way that is not easy to formalize. As Vaggione (2003) observed, reliance on formulas is not adequate; direct action (non-formulaic singularities) are also needed:

> *To articulate a highly stratified musical flux is unthinkable using operations based on statistical means. On the contrary, this requires an approach based on singularities of discontinuity, contrast, and detail. This is why causal formulas are problematic in composition if their automation is not compensated by other modes of articulation, i.e., unique compositional choices–singularities–global as well as local, integrated explicitly in the composition strategy.*

In discussing inspired choice, it is not a question of idealizing either the composer or the selection. Here "optimal" does not imply perfection; it is simply a particularly satisfactory choice given the constraints. Indeed, it would be hard to prove by scientific argument that a specific solution to a musical problem is inspired, satisfying, or optimal. It may simply "satisfice," to use Herbert Simon's (1969) term for "sufficiently satisfactory." Indeed, in many compositional decisions, more than one choice would be equally effective, but the composer simply had to pick one. Caprice is part of the composition process.

Economy of selection is an important concept because it emphasizes the role of intuitive choice in all compositional strategies. Even in generative composition, the algorithms are chosen according to subjective preferences. The rules are inevitably loosely constrained or incorporate randomness so as to allow many possible "correct" solutions. Computer programs can solve for and enumerate many of these solutions, but carefully picking the "best" or "optimal" solution is a human talent.

Moment-to-moment choices create our lives. I would go so far as to say that the talent of a composer lies primarily in his or her ability to listen and understand deeply enough to make optimal choices. This begins with choosing the right compositional problems to solve–a question of strategy, tactics, tools, and materials.

CONCLUSION: THE GRANULAR PARADIGM REFUSES TO DIE

The granular paradigm refuses to die. As Gabor showed, it is a universal representation for sound. The challenge for the composer has always been: what to do with granular materials? How do we build meaningful multiscale musical structures? What is the role of algorithms, and what is the role of gestural control?

The cloud, stream and spray paradigms have served as effective tools when combined with studio-based montage and micro-montage. Live performance with granular instruments is a more open question. We have not yet seen the Steinway of granular instruments nor the Cecil Taylor of granulation. The key to success in performing with granular instruments is the same as any other instrument: the development of virtuosity.

In the studio, graphical envelope control as in PulsarGenerator is an excellent compromise between gestural interaction and the kind of detailed micro control that can only come from scripts or code. Analysis-based granular processing is still in the beginning stages, with much territory to explore. Physical, biological, and abstract models of granular processes have potential, but probably more as metaphors for heuristic algorithms suited to specific compositional problems, rather than as full-fledged scientific models of reality. After all, the power of software is its ability to model not only known realities, but also fantastic imaginary worlds.

ACKNOWLEDGMENTS

I deeply thank Sharon Kanach (Forthcoming) for her detailed comments on this manuscript. Thanks to Makis Solomos and Charles Turner for their views on related historical matters. I would also like to deeply thank my partners in granular research over the years: John Alexander, Alberto de Campo, Garry Kling, Aaron McLeran, Bob Sturm, and David Thall. It has been a privilege to collaborate with them.

NOTES

1. This was only the beginning of my inspiration. A key encounter was experiencing the *Polytope de Cluny* in Paris eight times in the Fall of 1973.

2. Hear Bebe Barron's *Mixed Emotions*:
www.youtube.com/watch?v=Biqz1r2d_xY

3. The Constant-Q Granulator requires Mac OS9. The EmissionControl prototype requires a non-Intel PowerPC processor running MacOSX 10.4 (Tiger).

4. See Scatter: www.youtube.com/watch?v=JgGR6VjTiaA&lr=1

5. I am an advocate of Stephen Wolfram's important monograph *A New Kind of Science* (2002). However, his WolframTones is a classic example of a system that uses four billion cellular automata rules to produce trillions of pieces of unremarkable music. See: tones.wolfram.com.

6. Xenakis GENDY system embodies the notion of chains of interlocked probability functions, but he never applied this paradigm to granular synthesis.

REFERENCES

Aronson, I., and L. Tsimring (2009) *Granular Patterns.* Oxford, Oxford University Press.

Bacry, E. (2008) “Lastwave: the signal processing command language.” Internet: www.cmap.polytechnique.fr/~bacry/LastWave.

Bideau, D., and A. Hansen (1993) *Disorder and Granular Media.* Amsterdam, North-Holland.

Camarzine, S., J.-L. Deneubourg, N. Franks, J. Sneyd, G. Therauaz, and E. Bonabeau (2001) *Self-Organization in Biological Systems.* Princeton, Princeton University Press.

Cohen, H., B. Cohen, and P. Nii (1984) *The First Artificial Intelligence Coloring Book.* Los Altos, William Kaufmann.

Cook, P (1996) "Physically informed sonic modeling (PhISM): percussive synthesis." In L. Ayers and A. Horner, eds. *Proceedings of the 1996 International Computer Music Conference.* pp. 228-231.

Cook, P (1997) "Physically informed sonic modeling (PhISM): synthesis of percussive sounds." *Computer Music Journal* 21(3): 38-49.

Cook, P (2007) *Real Sound Synthesis for Interactive Applications.* Wellesley, A. K. Peters.

de Campo, A (1998) "Using recursive diminution for the synthesis of micro time events." Unpublished manuscript.

Di Scipio, A. (1990) “Composition by exploration of nonlinear dynamical systems.” *Proceedings of the 1990 International Computer Music Conference.* San Francisco, International Computer Music Conference. pp. 324-27;

Di Scipio, A (1999) “Synthesis of environmental sound textures by iterated nonlinear functions.” *Proceedings of the 2nd COST G-6 Workshop on Digital Audio Effects (DAFx99), NTNU, Trondheim, December 9-11, 1999.* Internet: www.iet.ntnu.no/groups/akustikk/meetings/DAFx99/discipio.pdf

Fletcher, N., and T. Rossing (1991) *The Physics of Musical Instruments.* New York, Springer-Verlag.

Gabor, D (1946) "Theory of communication." *Journal of the Institute of Electrical Engineers* Part 3, 93: 429-457.

Gabor, D (1947) "Acoustical quanta and the theory of hearing." *Nature* 159(1044): 591-594.

Gabor, D (1952) "Lectures on communication theory." Technical Report 238, Research Laboratory of Electronics. Cambridge, Massachusetts Institute of Technology.

Gogins, M (1991) "Iterated function systems music." *Computer Music Journal* 15(1): 40-48.

Gogins, M (1995) "Gabor synthesis of recurrent iterated function systems." In R. Bidlack, ed. *Proceedings of the 1995 International Computer Music Conference.* San Francisco, International Computer Music Association. pp. 349-350.

Gouk, P (1999) *Music, Science, and Natural Magic in Seventeenth Century England.* New Haven, Yale University Press.

Harley, J (2004) *Xenakis: His Life in Music.* New York, Routledge.

Huron, D (2006) *Sweet Anticipation: Music and the Psychology of Expectation.* Cambridge, The MIT Press.

Kanach, S. ed. (Forthcoming) *Xenakis Matters.* Pendragon: Hillsdale.

Keller, D., and B. Truax (1998) "Ecologically-based granular synthesis." In M. Simoni, ed. *Proceedings of the 1998 International Computer Music Conference.* San Francisco, International Computer Music Association. pp. 117-120.

Krstulovic, S., and R. Gribonval (2006) “MPTK: Matching pursuit made tractable.” *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, Toulouse, France,* Vol. 3, pp. 496–499. New York, IEEE.

Lewitt, S (1967) "Paragraphs on conceptual art.” *Artforum* 5(10): 79-83.

Logan, G (2000) *The Indiana University School of Music: A History.* Bloomington, Indiana University.

McLeran, A., C. Roads, B. Sturm, and J. Shynk (2008) "Granular methods of sound spatialization using dictionary-based methods." In *Proceedings of the Sound and Music Computing Conference, Berlin, Germany.* Internet: smcnetwork.org/files/proceedings/2008/session2_number4_paper20.pdfhttp://smcnetwork.org/files/proceedings/2008/session2_number4_paper20.pdf.

Miranda, E (1998) "Chaosynth: a cellular automata-based granular synthesiser." Internet: website.lineone.net/~edandalex/chaosynt.htm.

Moon, F (1987) *Chaotic Vibrations.* New York, Wiley-Interscience.

Preis, D. and V. Georgopoulos (1999) "Wigner distribution representation and analysis of audio signals: an illustrated tutorial review." *Journal of the Audio Engineering Society* 47(12): 1043-1053.

Reynolds, P., ed. (1993) *On Clusters and Clustering: From Atoms to Fractals.* Amsterdam, North Holland.

Risset, J.-C (2004) "The liberation of sound, art-science and the digital domain: contacts with Edgard Varèse." *Contemporary Music Review* 23(2): 27-54.

Risset, J.C (2005) "Horacio Vaggione: toward a syntax of sound." *Contemporary Music Review* 24(4-5): 287-294.

Risset, J.C (2007) Comments accompanying the lecture “Fifty years of digital sound for music.” In C. Spyridis, et al. *Proceedings of the SMC’07 4th Sound and Music Computing Conference.* Athens, National Kapodistrian University of Athens. pp. 3-8.

Rivier, N (1993) “Order and disorder in packings and froths.” In D. Bideau and A. Hansen, eds. *Disorder and Granular Media.* Amsterdam, North-Holland.

Roads, C (1973) “Analysis of the composition *ST/10* and the computer program Free Stochastic Music by Iannis Xenakis.” Unpublished manuscript.

Roads, C (1978) “Automated granular synthesis of sound.” *Computer Music Journal* 2(2): 61-62. Revised and updated version printed as "Granular synthesis of sound" in C. Roads and J. Strawn, eds. 1985. *Foundations of Computer Music.* Cambridge, The MIT Press. pp. 145-159.

Roads, C (1991) “Asynchronous Granular Synthesis.” In G. DePoli, A. Piccialli, and C. Roads, eds. 1991. *Representations of Musical Signals.* Cambridge, The MIT Press. pp. 143-186.

Roads, C (1992) "Composition with machines." In J. Paynter, T. Howell, R. Orton, and P. Seymour, eds. *Companion to Contemporary Musical Thought*. London, Routledge. pp. 399-425.

Roads, C (1992-1997) "Design of a granular synthesizer." Unpublished design documents.

Roads, C (1996) *The Computer Music Tutorial.* Cambridge, The MIT Press.

Roads, C (1998) "The Creatovox synthesizer project." Unpublished manuscript.

Roads, C (2001a) “Sound composition with pulsars.” *Journal of the Audio Engineering Society* 49(3): 134-147.

Roads, C (2001b) Microsound. Cambridge, The MIT Press.

Roads, C. (2006) “The evolution of granular synthesis: an overview of current research.” Presented at the International Symposium on The Creative and Scientific Legacies of Iannis Xenakis. Toronto: University of Guelph.

Roads, C (Forthcoming) *Composing Electronic Music: A New Aesthetic.* New York, Oxford University Press.

Roads, C., and J. Alexander (1995) *Cloud Generator Manual.* Distributed with the program Cloud Generator. Internet: www.create.ucsb.edu.

Simon, H (1969) *The Sciences of the Artificial.* Cambridge, The MIT Press.

Smith, W (2011) “Hierarchies of spaces: building from the bottom up.” Internet: moniker.name/worldmaking.

Solomos, M. (2006) “The granular connection (Xenakis, Vaggione, DiScipio…).” Presented at *The Creative and Scientific Legacies of Iannis Xenakis Symposium, Toronto.* Unpublished.

Solomos, M., A. Soulez, and H. Vaggione, eds. (2003) *Formel informel.* Paris. L’Harmattan.

Stuckenschmidt, H (1970) *Twentieth Century Music.* New York. McGraw-Hill.

Sturm, B., C. Roads, A. McLeran, and J. J. Shynk (2009) “Analysis, visualization, and transformation of audio signals using dictionary-based methods.” *Journal of New Music Research* 38(4): 325-341.

Thall, D. 2004a. “Experiments in sound granulation and spatialization for immersive environments.” Presented at the Interactive Digital Multimedia IGERT Annual Research Review, sponsored by the National Science Foundation. Santa Barbara: University of California.

Thall, D. 2004b. *A Scheme for Generalized Control of Particle Synthesis.* MS thesis. Media Arts and Technology. Santa Barbara: University of California.

Truax, B (1986) "Real-time granular synthesis with the DMX-1000." In P. Berg, ed. 1987. *Proceedings of the 1986 International Computer Music Conference.* San Francisco: Computer Music Association. pp. 138-145.

Vaggione, H (2003) “Composition musicale et moyens informatiques : questions d’approche.” In Solomos, M., A. Soulez, and H. Vaggione, eds. 2003. *Formel informel.* Paris: L’Harmattan. pp. 91-117.

Varga, B (1996) *Conversations with Xenakis.* London, Faber and Faber.

Wenger, E., and E. Spiegel (2004) *MetaSynth 4.0 User Guide and Reference.* Redwood City: U&I Software LLC.

Wolfram, S (2002) *A New Kind of Science.* Champaign, Wolfram Research.

Xenakis, I (1960) "Elements of stochastic music." *Gravesaner Blätter* 18: 84-105.

Xenakis, I (1971) *Formalized Music.* Bloomington, Indiana University Press.

Xenakis, I (1985) “Music composition treks.” In C. Roads, ed. *Composers and the Computer.* Madison, A-R Editions. pp. 169-192.

Xenakis, I (1992) “Concerning time, space, and music.” Chapter X of *Formalized Music,* revised 1992 edition. Stuyvesant, Pendragon Press. pp. 255-267.