

Composing with Real-Time Granular Sound

Author(s): Barry Truax

Source: *Perspectives of New Music*, Vol. 28, No. 2 (Summer, 1990), pp. 120-134

Published by: [Perspectives of New Music](#)

Stable URL: <http://www.jstor.org/stable/833014>

Accessed: 18-11-2015 11:05 UTC

REFERENCES

Linked references are available on JSTOR for this article:

http://www.jstor.org/stable/833014?seq=1&cid=pdf-reference#references_tab_contents

You may need to log in to JSTOR to access the linked references.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Perspectives of New Music is collaborating with JSTOR to digitize, preserve and extend access to *Perspectives of New Music*.

<http://www.jstor.org>




Perspectives of New Music

Vol. 28 n. 2 (1990)

Please note: The original version of this article was accompanied by an audio supplement that is not available through JSTOR at this time. For assistance in locating the supplement, you may wish to contact your librarian.

Please click on "Next Page" (at the top of the screen) to begin viewing this article.

COMPOSING WITH REAL-TIME GRANULAR SOUND



BARRY TRUAX

I. INTRODUCTION: SOUND AND STRUCTURE

DESPITE THE MANY CHANGES which the introduction of the computer into the compositional process has brought about, it is remarkable how frequently the compositional models implemented in both software and hardware retain the concepts and techniques of instrumental music. Perhaps these are the only ones which are familiar enough to computer music system designers, or expressed systematically enough to be replicated. In market-driven products, the instrumental music model is the only one considered viable enough commercially.


Central to the instrumental-music approach, and coincidentally to most computer music systems, is the separation of sound and structure. Within the MUSIC V system and its many relatives, this separation is expressed as the “score” and “orchestra”; within the MIDI world it is imbedded in the

difference between the “note on, note off” data and the frequently arbitrary nature of the timbres used in the synthesizer. A second tenet of the instrumental approach, namely constant timbre per voice, is frequently enforced within these same synthesizers. Of course, the thoughtful composer will not be misled by mere nomenclature and the resourceful experimenter will always find a way to overcome limitations, i.e. the influence may be much subtler. But these exceptions still prove the rule. The refined timbral designs of Jean-Claude Risset break free from the instrumental music world within the MUSIC V language; it will be interesting to see where the MIDI world will find its exception.

Several alternative approaches exist that unite sound and structure, the ones interesting me the most being those that stem from the practice of electroacoustic music. Walter Branchi (1983) has described the difference as that between composing “with” sound, as practiced in the instrumental approach, and composing “through” sound. The acousmatic approach, usually associated with the Groupe de Recherches Musicales, provides a particularly rich background of sonic-based thought. Among English-speaking composers, Denis Smalley (1986) has developed the acousmatic approach the farthest with his “spectro-morphological” concepts. Composers working with linguistic or environmental materials are usually less likely to detach sound from structure and meaning, playing instead on the interplay of these aspects. Simon Emmerson (1986) has provided an original insight to the classification of such approaches by his matrix of sound materials and sound syntax, each axis ranging from mimetic to abstract.

Over the past fifteen years of my compositional experience which has centered on computer-assisted composition and electroacoustic studio practice, I have sought a variety of ways of relating sound and structure intimately. A software-based compositional system, named POD (Truax 1977) and later PODX (Truax 1985), has provided the means, in addition to analog studio practice.


Early approaches within POD involved the mapping of timbral “sound objects” onto a syntactic field. Techniques such as the tendency mask (Example 1) allowed timbres to be introduced and gradually transformed into variants or abruptly changed to contrasting families, as in *Sonic Landscape No. 3* (1975, rev. 1977). Later work (Truax 1978) took an approach called “polyphonic timbral construction” to construct complex timbres from simple frequency modulation (FM) sources, as in my work *Androgyny* (1978). The next step involved treating timbre at a form-determining level (Example 2), as in *Arras* (1980), where spectral families implicit in the FM model were used to construct the form of the entire piece (Truax 1982). Mapping timbre onto spatial patterns called trajectories was a further step in *Wave Edge* (1983). These approaches culminated in *Solar Ellipse* (1984–85), where a spinning trajectory pattern informed everything in the work from the smallest base unit (one-quarter second) to the largest cyclic



EXAMPLE 1: TENDENCY MASKS USED TO DEFINE TIMBRE SELECTION (TOP),
FREQUENCY RANGE (MIDDLE),
AND TEMPORAL DENSITY (BOTTOM)
AS USED IN THE POD6 PROGRAM

pattern of the eleven-minute piece (Truax 1986a). Complementary timbral families, both vocal and noise-like, articulated these processes and were intimately entwined with them. It is perhaps significant that as these approaches brought sound and structure closer together, greater emphasis was placed on both the macro and micro levels of the composition—a simultaneous process of moving inwards as well as outwards.

An abrupt break in synthesis technique, though not in conceptual approach, has occurred in my work over the past three years with the development and use of granular synthesis. This technique has been proposed (Xenakis 1971; Roads 1978, 1988) as a unique method of achieving complex sounds by the generation of high densities of small “grains” on the order of magnitude of 10–20 ms duration. My recent work (Truax 1986b, 1988) has shown that this technique may be implemented in real



EXAMPLE 2: STRUCTURE OF OVERLAPPING AMPLITUDE ENVELOPES *Arras*,
 BASED ON THREE FUNDAMENTAL FREQUENCIES f_1 , f_2 , AND f_3 ;
 THIS UNIT STRUCTURE IN TURN OVERLAPS WITH ITS SUCCESSOR
 (DASHED LINES) AFTER TWENTY-FIVE SECONDS

time using the microprogrammable DMX-1000 Digital Signal Processor (Wallraff 1979), thus taking the burden of data specification, storage and calculation off both the user and the machine. Synthesis techniques involving additive synthesis and simple FM pairs for each grain have been implemented, as well as the granulation of stored sampled sound (Truax 1987). A hierarchy of control parameters direct the density, frequency range, and temporal evolution of the synthesized sound textures.

The basis of granular synthesis in the seemingly trivial grain has had a powerful effect on my way of thinking about sound. It clearly juxtaposes the micro and macro levels, as the richness of the latter lies in stark contrast to the insignificance of the former. Moreover, the range of densities obtainable, from the low levels associated with human gestures through those perceived as rapid and virtuosic, culminating with entirely fused textures, suggests a scale of composition ranging from human proportions to abstract. Finally, in terms of the theme of sound and structure, it is clear that the two are inseparable with this technique. The macro level structure is best described in terms of its component sounds, and the resulting sound complex is definable only in terms of the structural levels that characterize its organization.

II. COMPOSITIONAL EXPERIENCE WITH GRANULAR SYNTHESIS

To date, I have composed four works with the granular synthesis technique. The first, *Riverrun* (1986), is based on synthesized grains which are either

simple waveforms or FM signals. However, despite the basis in fixed wavetables, the resulting textures are always dynamically changing and range from swarms of relatively isolated sound events to fused sound masses of great internal complexity, much like environmental sound generally and water sound in particular. The fundamental paradox of granular synthesis—that the enormously rich and powerful textures it produces result from its being based on the most “trivial” grains of sound—suggested a metaphoric relation to the river whose power is based on the accumulation of countless “powerless” droplets of water. The opening section of the work portrays that accumulation, as individual “droplets” of sound gradually multiply into a powerful broad-band texture. The dynamic variation implemented by changing variables continuously allows the piece to create a sound environment in which stasis and flux, solidity and movement coexist in a dynamic balance similar to a river which is always moving yet seemingly permanent. The piece, I find, also captures some of the awe one feels in the presence of the overpowering force of such a body of water, whether in a perturbed or calm state, and as such it seems to create a different mode of listening than does conventional instrumental or electroacoustic music.

Two more recent works, *The Wings of Nike* (1987), for computer images by Theo Goldberg and two soundtracks, and *Tongues of Angels* (1988), for oboe d’amore, English horn, and four soundtracks, are based on the granulation of sampled sound using very short fixed samples of recorded material. In the case of the first work, these samples are male and female phonemes, and in the second piece the samples are derived from the live instruments. Despite the brevity of the source material, very rich textures and complex rhythmic patterns can be obtained from it. The pitch and timbre of the resulting sound are determined by the source material unless the grain duration is too short and a broad-band spectrum results. However, the overlay of up to twenty simultaneous versions of such sound, each with its own variations, produces a “magnification” of the original sound, as well as the possibility of gradual or rapid movement through its micro-level characteristics.

The degree of magnification involved can be appreciated when it is realized that three of the four movements of *The Wings of Nike*, lasting approximately twelve minutes, were derived from only two phonemes, each about 170 ms long! The stereo tape is a mixdown from an eight-track original which includes four stereo pairs of the granular material, and therefore the vertical densities of sound are around eighty at any one moment, and the horizontal densities range from quite sparse through to eight thousand events per second at the very end. The sounds are heard at approximately their original pitch combined with versions an octave up or down, except towards the end of the first movement where a slow downward glissando is heard.

In the case of the male phoneme sequence (“tuh”), timbral changes are possible as the grains move (with changing offset) from the aspiration at the beginning of the sample, through the consonant to the pitch of the voiced vowel and finally past it. Microscopic timbral changes that normally go unnoticed become evident with the repetition of the overlapping grains. Formant regions become especially pronounced. The analogy in the work is between the sampled visual image of the statue (the Winged Victory or Nike of Samothrace) that forms the basis of the computer-generated images and the sampled vocal sound as the source of the synthesized music. Each is based within the human dimension but the transformation techniques extend them towards the epic and the supra-human, as at the end of the first movement (“Album”).

In the second movement, the “Scherzo,” the statue is given a head, though rather androgynously, and the phonemes in the tape part are detached and realistic. However, each image accumulates again into more elaborate patterns that mask the source. In the third movement (“The Illuminated Nike”), FM grains are used. Just as the image of the statue undergoes color variations in illumination, different carrier-to-modulation ratios and modulation indices “color” the carrier frequency of the tape with timbral variations. In the last movement, the “Coda,” both the visual and aural elements are progressively multiplied (the wing feathers and drapery in the visual case, and the syncopated phonemes in the musical accompaniment) until fusion results. A new kind of machine emerges that releases the earth-bound images into upward flight.

Speed, rhythmic patterns, and density changes such as gradual accumulation or evaporation of the sound are other compositional strategies that granulation makes available. Speed and density, for instance, can change from isolated events, through to fast iterations (each with slightly different timbral characteristics), and finally result in fused textures. Stochastic patterning of the grains can occur over the same range of densities as well, giving the sense that the sound is coalescing or evaporating.


In the interaction of the tape with the live performer in *Tongues of Angels*, this progression from speeds and patterns within the performance range of the instrumentalist to those achievable only by the computer creates much of the underlying dramatic tension of the work. We sense the struggle of the performer to match the virtuosity of the tape part and to transcend the limitations of human performance. The transcendence finally materializes at the end of the work when, after a gradual accelerando in both parts from syncopated rhythms through to fused textures and continuous gestures, there is a dramatic downward glissando on the tape. Despite the fact that all material was derived from the instruments heard live in the performance, the effect is suddenly that of a choir of voices, a transcendent image that inspired the Biblically derived title of the work. Perhaps it could also be a

metaphor for the technique of granular synthesis that seems to break into a new sonic domain.

III. COMPOSITIONAL CONTROL STRATEGIES

Given the enormous amount of data involved in specifying thousands of events per second, powerful control strategies are required to make this synthesis technique effective for the composer. The current implementation has developed a hierarchy of control levels (Example 3) that include both improvisational and design-based strategies. At the lowest level are the control parameters already described which may be altered by various keystroke commands. Groups of these control parameters are called presets (which are assigned to specific keys) and these in turn may be incorporated into a higher-level score. Alternatively, a performance of a sequence of presets may be saved as a score. Presets and scores may be performed with two kinds of dynamic articulation. In the first, the overall amplitude level is constant and the preset initiates an abrupt change in the prevailing texture; in the second, each preset initiates a decay in amplitude, thereby articulating it as a separate event.

COMPOSITION HIERARCHY



EXAMPLE 3: COMPOSITIONAL HIERARCHY OF CONTROL LEVELS USED
IN THE AUTHOR'S SYSTEM FOR GRANULAR-SYNTHESIS PROGRAMS

In contrast to the discrete changes implemented with presets, the second control route is via ramps, that is, patterns of continuous change in selected parameters at a specific rate. Ramps may be ascending or descending, as


well as cyclic within minimum and maximum values. With each ramp, the user chooses which parameters are to change (i.e. are “synchronized”), at what speed and whether the change is in the direction of the ramp, or inverse to it. Ramps may be manually activated or a sequence of them may be predetermined and stored in a ramp file. Ramps may be combined with presets and scores in any manner, thereby mixing discrete and continuous patterns of change. When combined with scores, the rule is that any variable singled out for synchronization is no longer controlled by the score, but rather by the ramp or other commands by the user.

The third compositional strategy involves tendency masks. Although they are translated into the equivalent of a combination of presets and ramps, they appear as graphic control shapes to the composer and hence suggest a different compositional approach (Example 4). The tendency mask, being inherently a visual control method, presents an image of the control shape based on the limiting values within which choices are made. The tendency mask suggests gestures, whereas the changing numerical parameters in a ramp suggest ongoing processes. Nevertheless, the two methods are functionally equivalent in terms of the way the computer implements them.

The tendency mask is expressed as an area within which values may be chosen. It is one of the simplest ways to specify time-dependent selection of a variable. Although it appears to the user as a shape, it can also be understood as the combination of a time-dependent average value and a width or range around that average, the same two variables used to control the granular synthesis events. The width of the mask at any point determines the range of choices available, thereby providing a continuum between deterministic and stochastic choices. The average value of the mask, although an imaginary line through the middle, is usually perceptible, particularly when it changes over time. The mask is input numerically by specifying the values of each corner plus duration. Alternatively, masks can be generated as interpolations between presets; that is, each segment of a sequence of masks can be derived as starting with the values of a certain preset and ending with those of another preset. The difference between the two methods is that with interpolation, all segments of each mask have the same duration, whereas with the direct specification method, no synchronization between masks is required. However specified, tendency masks are a convenient high-level control structure that has proved useful to composers working with the system. Operations on masks would provide an additional compositional level.

The current implementation, while just a beginning, has already gone some distance in establishing an appropriate hierarchy of levels of compositional control for granular synthesis, ranging from the control variables at the grain level, through to groups of such variables (presets), rates of change

FILE: NIKE1C



EXAMPLE 4: TENDENCY MASK CONTROL SHAPES USED IN *The Wings of Nike*
 TO DEFINE FREQUENCY RANGE (LIGHT SHADING)
 GRAIN DURATION (MEDIUM SHADING),
 MODULATION INDEX (DARK SHADING),
 AS WELL AS DELAY TIME BETWEEN GRAINS (SOLID LINE SEGMENTS)

(ramps) of the control variables, and macro-level tendency mask and scores to determine large-scale forms.

IV. REDEFINING ACOUSTIC AND MUSICAL MODELS

Technology in general, and electroacoustics in particular, embodies knowledge. Computer systems must embody quite specific models about the tasks they are to perform, hence it is very easy for them to be based on widely accepted, traditional models simply because they are so well known. However, the design of such systems can also require us to think more carefully about what constitutes the basis of our field of knowledge. In my opinion, granular synthesis currently raises the most fundamental questions about contemporary acoustic and musical knowledge.

In taking us back to first principles, granular synthesis questions what those principles really are, both acoustically and musically. It challenges the traditional Fourier model of sound which states that any periodic waveform may be reduced to the sum of a set of harmonically related sine waves each with a different amplitude (and phase). The first problem with the basic Fourier theorem is that it ignores time; it implies that the constituent sine waves have infinite duration. In order to deal with actual musical sounds, each cycle of the sound may be analysed to establish the time variance (or envelope) of each component harmonic (Mathews and Risset 1969). However, such sounds are not periodic as they begin, having instead an attack transient or burst of noise as the instrument is set in motion. It is interesting to note that this short period of time as the sound begins, when the amount of acoustic variation is the greatest, is the most critical for the identification of the sound. It is also the part that contains the greatest amount of information. If that part of the sound is cut off, the rest is barely recognized as coming from the instrument in question.

The second problem with the Fourier model is that it corresponds poorly with the psychoacoustic reality of how our auditory system actually processes musical sound (Roederer 1975). Although we can perform Fourier analysis, i.e. distinguish separate harmonic components, we can do so only for the lower harmonics whose frequencies are sufficiently far apart, and only under conditions where the sound is prolonged such that we have time to focus on those harmonics. In other words, we have to simulate the unchanging conditions implied by the Fourier theorem for it to be applicable. In more realistic situations, with sound densities in music involving several notes per second in a melodic line, possibly combined with chords and various accompanying instruments, the auditory system uses much more efficient means to identify and track particular instruments. The problem is not that the Fourier approach is wrong (within its assumptions it functions quite well), but that it forms the theoretical basis for all oscillator-based, and with digital technology, all wavetable-based synthesis. Its lack of concern for micro-level time variation continues to plague synthesized sound as something “unnatural.”

In 1947 the British physicist Dennis Gabor proposed an ‘acoustical quantum’ as the fundamental unit of sound that incorporates both frequency and time because “it is our most elementary experience that sound has a time pattern as well as a frequency pattern” (Gabor 1947, 591). In other words, the quantum is the shortest duration of sound that will activate the auditory system. It is an “event,” not merely a fixed stimulus. Although the techniques of granular synthesis depart from the ideal Gabor grains, they still involve the principle of building complex events from the seemingly trivial micro-level, enveloped quanta called grains. The parameters of these grains provide the basis for controlling all acoustic properties of the resultant sound.

For instance, if the component grains all have a similar frequency, the result is a musical pitch, but as the frequencies of the grains are spread over a larger range, the result becomes a broad-band noise, much like the natural textures of wind and water. These sounds create a remarkable sense of space and volume. Similarly, if the durations of the grains are all similar, the frequency spectrum and timbre of the resulting sound is much richer because of the phenomenon of amplitude modulation. In addition, when a short delay is placed between grains and then increased, the fused granular texture begins to “pull apart” as isolated events emerge perceptually from the texture and eventually establish rhythmic relations. Once global amplitude contours are added along with spatial deployment, we have replicated all of the basic acoustic properties of sound.

There are also many fascinating parallels between the world of granular synthesis and the subatomic level of physical matter. First, of course, is Gabor’s idea of the quantum of sound, an indivisible unit of information from the psychoacoustic point of view, on the basis of which all macro-level phenomena are based. However, more specifically, Gabor postulated that the parameters of this quantum were frequency and time. The student of acoustics is often taught that these are two distinct and unrelated parameters, but of course, on a deeper level of understanding they prove to be reciprocal to each other, both physically and psychoacoustically. There is also an “uncertainty principle” relating them, analogous to Heisenberg’s uncertainty principle about particles at the subatomic level. He demonstrated that the more precisely you determine the position of an electron the less you know about its velocity, and vice versa. This is because velocity is the rate of change of position. Similarly, the smaller the time window in acoustics, the larger the uncertainty (i.e. bandwidth) in frequency. Frequency, as the rate of change of phase, is the reciprocal of time.

In another analogy to quantum physics, time is reversible at the quantum level. The quantum equations are symmetrical with respect to the direction of time, and similarly, the quantum grain of sound is reversible with no change in perceptual quality. If a granular synthesis texture is played backwards it will sound the same, just as if the direction of the individual grain is reversed (even if derived from natural sound). That is, acoustic time has no preferred direction at the micro level. All of the characteristics that establish the directional experience of time occur at the macro level. Granular synthesis is in a unique position to mediate between the micro and macro time worlds.

This time invariance permits a time shifting of sound derived from the real world, allowing it to be slowed down, up to a thousand times or more, with no change in pitch. Spectra that are normally brief instants in time can now occupy virtually any duration. In my current work (Truax 1989), as used in *Beauty and the Beast* (1989), grains are taken from the sound

sequence at specific points, but the rate at which these points advance through the sequence need not be at the original clock rate. For instance, the current pointer may advance for a certain number of milliseconds, then stop for a certain duration before continuing. When “clock time” is stopped, the grains can move backwards through the sound with no difference in sound quality. Paradoxically, by linking frequency and time at the level of the grain, one makes them independently controlled variables at the macro level.

The complexity and dynamic quality of granular-synthesis sound makes it an attractive alternative to methods based on fixed waveforms or to resynthesis models based on the transposition of sampled sound. Moreover, the basic unit or “quantum” of the grain is a potentially more flexible building block for sound synthesis than the sine wave or the use of longer sequences of sampled sound. It is also a very useful means of manipulating arbitrary groups of sound samples, particularly because the amplitude envelope of the grain avoids transient clicks when extracting and combining sample segments. When granular synthesis is used to produce continuous textures, it has no resemblance to instrumental and other note-based music; instead, its sound world is more closely related to analog electroacoustic music, but with greater precision of control. In certain cases, the acoustic result resembles environmental sounds in terms of their inner complexity and statistical texture. However it is used, granular synthesis is clearly situated in a different psychoacoustic domain than that occupied by most computer music; it creates a unique sound world and suggests new approaches to the way music made with it is formed.

In terms of fundamental musical issues, such as compositional control, granular synthesis also requires a shift in thinking. It makes us rethink what the composing process involves. It is obviously impossible for the composer to specify each individual grain, given that there may be thousands of them per second. It reduces to absurdity the idea of total control by the composer. Hierarchic levels of control are absolutely necessary, and at each level the composer specifies key control variables on the basis of which all of the specific data is calculated by the program. Interestingly enough, the structure of the software that realizes the technique is also different from conventional programming approaches that tend to be linear and deterministic. The granular software, in contrast, more closely resembles an organic system in which parallel processes are working independently, linked by access to common variables representing “the state of the world” and by messages passed between processes. The standard debugging technique of stopping the program and stepping through its instructions one at a time is meaningless in this case; once stopped, the program loses its most essential characteristics, its time behavior and the interaction between levels.

In this kind of system, the composer functions not as an omniscient

arbiter, but as the source of control messages that guide the overall process without directly determining it. True, other composers have used indeterminacy and emphasized processes more than data, but granular synthesis uniquely combines these ideas into both sound and structure. My experience of using the technique has also convinced me of the potential for computer technology to be a new instrument for musical thought, not just musical production. The crucial difference seems to be between the composer being regarded simply as a 'user' of a piece of software, and the composer-software system being treated as an interplay between two forms of musical knowledge. In the second case, the composer's musical insights that are provoked through use of the system are ideally channeled back into it in the form of new control possibilities. In other words, the computer allows the composer's often unstated musical knowledge to be engaged in musical actions that can lead to new compositional tools as well as to an enriched musical discourse.

Yet there seems to be a trend today toward increased standardization of designs in musical technology that embody very little musical knowledge and do not encourage musical experimentation. New ideas may lead to new technology, but will new technology lead to new ideas? New musical ideas seem particularly elusive. The fact that the computer allows the composer a significant choice of compositional process, not simply materials or structure, seems to have been forgotten in much of the currently available software. From the perspective of the granular approach, it seems that the most promising directions may well be found on the fringes, whether in MIDI-based or synthesis-based systems, not in the mainstream of musical development with the computer.

REFERENCES

- Branchi, Walter. 1983. "The State of Anxiety." *Computer Music Journal* 7, no. 1 (Spring): 8–10.
- Emmerson, Simon. 1986. "The Relation of Language to Materials." In *The Language of Electroacoustic Music*, edited by Simon Emmerson. London: Macmillan.
- Gabor, Dennis. 1947. "Acoustical Quanta and the Theory of Hearing." *Nature* 159 (no. 4044): 591–94.
- Mathews, Max, and Jean-Claude Risset. 1969. "Analysis of Musical-instrument Tones." *Physics Today* 22, no. 2:23–30.
- Roads, Curtis. 1978. "Automated Granular Synthesis of Sound." *Computer Music Journal* 2, no. 2 (Summer): 61–62. Reprinted in *Foundations of Computer Music*, edited by Curtis Roads and John Strawn. Cambridge, MA: MIT Press, 1985.
- . 1988. "Introduction to Granular Synthesis." *Computer Music Journal* 12, no. 2 (Summer): 11–13.
- Roederer, Juan G. 1975. *Introduction to the Physics and Psychophysics of Music*. 2d ed. New York: Springer.
- Smalley, Denis. 1986. "Spectro-morphology and Structuring Processes." In *The Language of Electroacoustic Music*, edited by Simon Emmerson. London: Macmillan.
- Truax, Barry. 1977. "The POD System of Interactive Composition Programs." *Computer Music Journal* 1, no. 3 (Fall): 30–39.
- . 1978. "Computer Music Composition: The Polyphonic POD System." *IEEE Computer* 11, no. 8:40–50.
- . 1982. "Timbral Construction in *Arras* as a Stochastic Process." *Computer Music Journal* 6, no. 3 (Fall): 72–77.
- . 1985. "The PODX System: Interactive Compositional Software for the DMX-1000." *Computer Music Journal* 9, no. 1 (Spring): 29–38.
- . 1986a. "Computer Music Language Design and the Composing Process." In *The Language of Electroacoustic Music*, edited by Simon Emmerson. London: Macmillan.
- . 1986b. "Real-time Granular Synthesis with the DMX-1000." In *Proceedings of the 1986 International Computer Music Conference*, edited by Paul Berg. San Francisco: Computer Music Association.

- . 1987. "Real-time Granulation of Sampled Sound with the DMX-1000." In *Proceedings of the 1987 International Computer Music Conference*, compiled by James Beauchamp. San Francisco: Computer Music Association.
- . 1988. "Real-time Granular Synthesis with a Digital Signal Processor." *Computer Music Journal* 12, no. 2 (Summer): 14–26.
- . 1989. "Time-shifting of Sampled Sound with a Real-time Granulation Technique." In *Proceedings of the 1989 Canadian Electroacoustic Community Conference*. Banff.
- Wallraff, Dean. 1979. "The DMX-1000 Signal-Processing Computer." *Computer Music Journal* 3, no. 4 (Winter): 44–49. Reprinted in *Foundations of Computer Music*, edited by Curtis Roads and John Strawn. Cambridge, MA: MIT Press, 1985.
- Xenakis, Iannis. 1971. *Formalized Music*. Bloomington: Indiana University Press.