Journal of New Music Research
2004, Vol. 33, No. 2, pp. 161–171

Analysis by Modeling: Xenakis's *ST/10-1 080262*

Damián Keller[1] and Brian Ferneyhough[2]

[1]Center for Computer Research in Music and Acoustics; [2]Department of Music, Stanford University, Stanford, CA 94305, USA

Abstract

This paper proposes analysis by modeling as a complement to current analytical practices. We present an analysis of Xenakis's algorithmic composition *ST/10-1 080262*. Two models were developed in order to study the temporal quantization processes and the streaming-fusion processes in this piece. Our study shows that fused granular textures could have been obtained by constraining the use of timbres to closely related classes and by increasing slightly the density of events.

1. Introduction

Musical analysis is usually carried out on a musical work represented by a score, or by a recording of its performance. This methodology functions sufficiently well with deterministic pieces but is less well suited to the discussion of process-based musical phenomena. These phenomena, often included in performance art, sound installations, algorithmic music, and various forms of sound-environment systems, need a more flexible analytical approach. In this paper we focus on issues raised by one of the first computer-generated compositions, *ST/10-1 080262* (Xenakis, 1967). As an algorithmic piece, *ST/10* can be viewed as one among the many possible realizations of a particular compositional model. Hence, to understand the mechanisms of the model is a prerequisite to understand the piece. We suggest that the specification of a model for salient aspects of the piece – analysis by modeling – may provide a methodologically viable approach. The two studies discussed in the last section of the present paper follow this approach to point out two features of the piece: (1) the underlying quantization processes; (2) the process of fusion in dense textures. In the preceding sections, we deal with inconsistencies among conceptual and technical procedures in stochastic music, and with the perceptual mechanisms brought into play by the sound processes found in *ST/10*.

The paper also discusses the algorithms used by Xenakis in composing stochastic music. Their output was mapped onto traditional notation to obtain a readable instrumental score. This mapping process definitely distanced the model from its final manifestation. The most problematic aspect is in the approach to time within the compositional process. In *ST/10*, this issue comes to the foreground on three levels of the musical structure: the overall distribution of events, the division into completely independent sections, and the use of time-invariant parameters.

The paper presents a short description of two perceptual processes that play an important role in auditory processing of sound textures: sequential and simultaneous stream segregation (Bregman, 1990). It raises the question of whether complex events produced by the interaction of many short sounds, or *grains*, occur within the context of *ST/10*. Two factors seem to define whether granular events could be created with the instrumental palette employed by Xenakis: grain density, and timbre-based auditory streaming. In our study, we ask whether fused textures could have been obtained by using higher densities of events.

2. *ST/10-1 080262*

In some of the chapters of *Formalized Music* (Xenakis, 1992), Xenakis described the compositional methods utilized in *ST/10* and *Achorripsis*. In contrast with traditional methods, his approach required the existence of a pre-compositional model in order to realize the work, so that musical pieces might be understood as instantiations of a

Correspondence: Damián Keller, Center for Computer Research in Music and Acoustics, Stanford University, Stanford, CA 94305, USA. E-mail: dkeller@ccrma.stanford.edu

DOI: 10.1080/0929821042000310630 © 2004 Taylor & Francis Ltd.

general model. Ideally, from the observation of these pieces we should be able to infer the structure of the underlying model. In practice, the direct realization of an algorithmic model as a musical piece is impossible. Even when no players are involved, such as it is the case in computer-generated tape pieces, sounds have to be played through loudspeakers, and circumstances including placement of loudspeakers and room acoustics have a strong influence on the final sonic result. By the same token, issues involved in the auditory processing of musical sound may be difficult to fit within this compositional approach. As it will become clear from our discussion, purely formal compositional methods present several shortcomings in relation to perceptually based sonic organizations (Keller, 1999, 2000b; Keller & Capasso, 2000).

Xenakis employed his ST (Stochastic Music) program to compose a number of works, including *ST/10*. The program was coded by M. F. Génuys and M. J. Barraud in FORTRAN IV on an IBM 7090 (Xenakis, 1992, pp. 145–152). The "tremendous" speed of this computer was on the order of 500000 calculations per second. Xenakis utilized ST to generate data in text format which he later transcribed to musical notation. As he declared, the transcription was a delicate step, and required the making of several compositional decisions.

ST provided a list of parameters such as attack time, "instrument class," "instrument" (i.e., selection of instruments within a given class, and the playing technique), pitch, duration, dynamics, and three parameters for glissandi. The parametric representation aimed to facilitate subsequent mappings onto musical notation. Pitch was expressed as a floating-point number between 0 and 85. Duration was also expressed as a fractional number, and dynamics – or intensities – were represented by integers on a scale of 0–60. The three glissando values were produced either as a function of the section's average density: (1) inversely proportional; (2) directly proportional; or (3) as a random distribution independent from the density (Xenakis, 1971, p. 140). Only one of the three mappings was used per section. A positive value indicated a glissando toward a higher pitch and a negative value suggested a glissando toward a lower pitch (Xenakis, 1971, p. 13).

2.1. Mapping

As he did with several other pieces, Xenakis named *ST/10-1 080262* following a shorthand convention: "ST" stands for "stochastics," "10" is the number of instrumental groups (percussion instruments are counted as a single group), and "1" stands for the version of this piece. (A subsequent piece, bearing the title *Atrees*, corresponds to the version number *ST/10-3*.) The final number, "080262," refers to the day when the data for the piece was generated: February 8, 1962.

ST/10 consists of a sequence of sections, or "movements." The duration of each section was determined by a Gaussian distribution around a mean. The mean is provided as a parameter by the composer. From Xenakis's description, it is not clear whether he utilized the random durations as an ordered set or whether he ordered the sections to suit his own taste. In any case, from a compositional point of view the most relevant parameter was the density of each section. Density was defined as the average number of events (note onsets) occurring within a section. Xenakis established a lower limit at 0.11 onsets per second. The upper limit was less than 50 onsets per second. The distribution of densities for the various sections was controlled by a logarithmic curve with a base equal to 2.71827. The density levels were chosen randomly from this curve. This random choice procedure produces a formal averaging or smoothing effect, where the distribution profile impacts the overall texture of the piece but has no effect on the relationships between the movements of the work. Xenakis made an attempt to counteract this effect by providing short-term "memory" to the algorithm: the current density value was calculated as a random deviation from the previous value.

Xenakis (1992, p. 139) controlled the distribution of instrumental choices by linking them to the mean density of each musical section. The choice of instrumentation provides a timbre palette made of two distinct timbral classes: one including percussive sounds such as pizzicato strings, harp, and percussion; the other including winds and bowed strings. Probability of instrument occurrence is distributed almost uniformly along the density axis. Therefore, no linear relationship between density and global timbral characteristics is implied.

A random function was used to place events on the time axis. The number of events was determined by the density level of the section. Xenakis mapped the values obtained as *delta-times*, that is, the time between note onsets among all voices. This reduced the probability of occurrence of synchronous onsets to the case when *delta-time* was equal to zero. A sample of the output produced by ST illustrates how instrumental classes were distributed along the time axis (Figure 1) (Xenakis, 1971, p. 153). There are three synchronous attacks at 1.09, 2.18 and 4.15 seconds.

Also, although he did not expressly state how duration values were translated to traditional notation, Xenakis evidently applied a quantization procedure. The finest granularity he used was six eighths per half note. Throughout the piece he kept polyrhythmic relationships among voices, with overlapping 4 : 4, 5 : 4, and 6 : 4 rhythms at the beginning, and superimposed 3 : 4, 4 : 4, 5 : 4, and 6 : 4 rhythms in the last sections. Durations were quantized by subdividing a half note into 3, 4, 5, and 6 parts.

As in any mapping procedure, there are always incompatibilities among different domains. Xenakis opted for a compromise. Pitches were established by mapping random integers onto the instrumental range. As with density, the current value was calculated as a deviation from the previous value. He adopted the chromatic scale, but also included some quartertones. Intensities were approximated to traditional notation markings. Xenakis chose not to employ crescendos or diminuendos across several notes. He applied time-varying dynamics only to single notes. Playing tech-

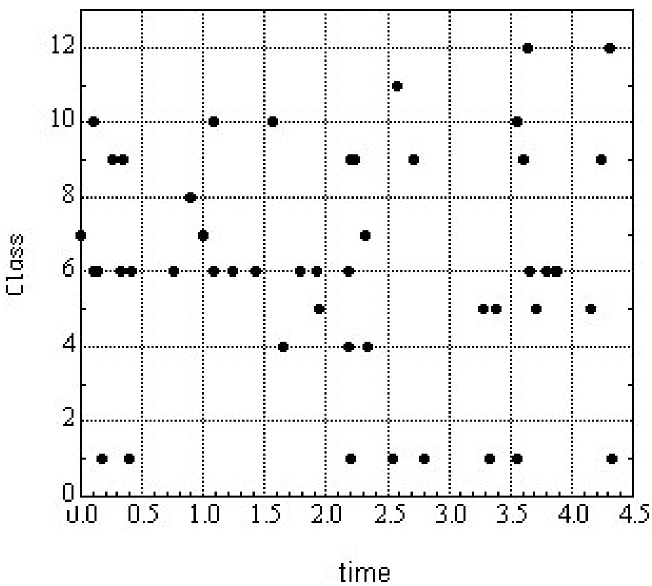

Fig. 1. Plot of distributions of instrumental classes along the time axis generated by ST.

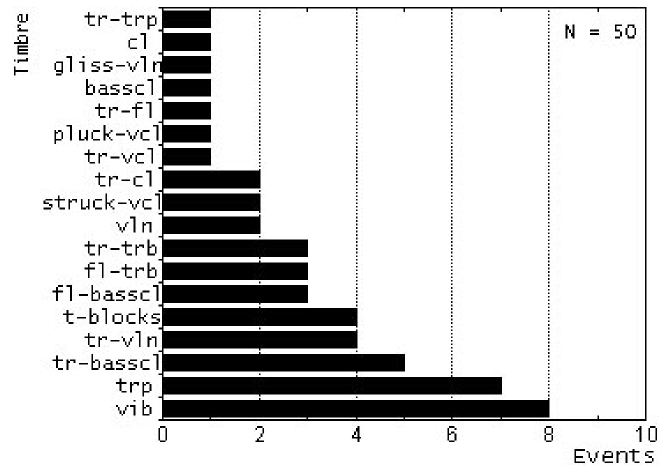

Fig. 2. Analysis of the distribution of timbres generated by ST.

niques included *pizzicato*, *arco*, *sul ponticello*, *col legno*, and *tremolo*, as well as *flutter-tongue* for the winds and rolls for the percussion. No smooth transition among playing techniques was employed.

Instruments were randomly mixed, such that timbre-based grouping of sound events was usually precluded. Even though this seems to be a side effect of utilizing independent processes to generate parameters for onset time and instrument choice, this procedure had a profound effect on the textural surface of the piece (Example 1). A quick analysis of the distribution of timbres generated by ST (Xenakis, 1971, p. 153) yields the following results: (1) the distribution is logarithmic, thus some instrumental classes predominate over others (Figure 2); (2) classes are evenly distributed along the time axis (Figure 1); (3) synchronous attacks are rare, three out of 50 events in this example.

Example 1. An algorithmic model of the quantization processes in *ST/10* (PATCHWORK with the addition of the *Alea Library*).

2.2. Time

In his doctoral dissertation, Xenakis stated that "the faculty of condensation-toward-abstraction is part of music's profound nature (more than any other art's) rather than simply being a function. Consequently, it seems that a new type of musician is necessary, an "artist-conceptor" of new abstract and free forms (. . .). The artist-conceptor will have to be knowledgeable and inventive in such varied domains as mathematics, logic, physics, chemistry, biology, genetics, paleontology (for the evolution of forms), the human sciences and history; in short, a sort of universality, but one based upon, guided by and oriented toward forms and architectures. The time has come to establish a new science of 'general morphology' [to] treat these forms and architectures (. . .)." (Xenakis, 1985, p. 3). He describes this new science as "an abstract approach, free from [the] anecdotes of our senses and habits." The emphasis placed on "abstract forms" should not go unnoticed. Xenakis understood composition as an activity that dealt with a multiplicity of levels organized as a hierarchy. At the bottom of this hierarchy we find the sound wave (Xenakis, 1985, p. 10). As it is clear in the previous quote, the abstract musical form is at the top of the hierarchy. Furthermore, he stated that musical composition should be *free from the anecdotes of our senses and habits*. There can be no doubt that this approach places formal concerns well above social and perceptual issues. This section will analyze how Xenakis's view of music as an abstract entity determines the compositional procedures utilized in ST/10.

The relationship between time and compositional process is the key to understanding the musical results obtained in stochastic music. Xenakis's music is shaped up outside time. As is well known, he proposed two modes of musical operations: outside-time and in-time (Xenakis, 1992, pp. 170–173). First, generation and transformation of material are accomplished by means of operations on elements of a set. These operations remain abstract until they are mapped onto a temporal axis. As a second step, Xenakis proposed operations on temporal classes, namely on temporal distances between events. The relationships between outside-time operations on parameters and operations on distances among events on the time axis define the final compositional result. In other words, compositional practice is based on a network of abstract relationships. These relationships are mapped onto a temporal axis which is treated as a *tabula rasa* where any sonic configuration can take place.

The independence of formal processes from their sonic manifestation creates a conceptual and methodological paradox in *ST/10*. The gap between outside-time processes and time-based sonic evolutions can be tracked through three temporal levels. These levels correspond to the distribution of instrumental events, the temporal proportions of the sections, and the formal structure established by the placement of salient instrumental gestures.

Xenakis's pre-compositional plan set out a sequence of "movements each a_i seconds long" (Xenakis, 1992, p. 134). We see that although all "movements" or sections share a common algorithmic origin, their temporal deployment and their length are not based on sonic relationships. Conceptually, each section is a separate entity linked to the others just by being concatenated. Furthermore, the section length is not determined by the dynamics of its material but by a top-down strategy: all lengths are calculated stochastically before parameter values are generated. Thus, Xenakis obtained a temporal mold which he applies onto his material. As is made clear in the score (Xenakis, 1967), to counterbalance the rigidity of this process, the composer introduced several dynamic tempo markings which range from slow changes through a whole section (e.g., section 1, denoted by the composer as "JW = 1"), to short *rallentandi* and *accelerandi* (e.g., measures 26–29 and 204–206).

In *ST/10*, sound densities are all well under 100 events per second. Measures 132–138 present an average of 15 events every 0.83 seconds. Measures 365–367 do not exceed 50 events per measure, each measure lasting 1.33 seconds. These rather sparse densities – in comparison to current granular synthesis standards – encourage the listener to focus on the characteristics of individual events. Instrumental gestures acquire special salience.

The mapping of successive onsets onto all instrumental voices reduces the occurrence of synchronous attacks to a minimum. This phenomenon can be clearly observed in the low-density passages. For example, there is not a single synchronous onset among instruments in the whole second section. The mapping process produces sparse and consistent textures which effectively differentiate the low-density sections from the rest. Paradoxically, this consistency gives a static quality to the relationships among instruments and a higher degree of predictability to the succession of attacks.

Similar predictability can be found at the highest structural level. Due to the ergodic processes and distributions utilized throughout the piece, the high-level sonic structure averages out local differences. Because parameter generation mechanisms are not time-dependent, changes in the underlying organizational processes are perceptually difficult to track (Keller & Berger, 2001). Although Xenakis (1971) suggests that his compositional system creates transitions from ordered to disordered states, these transitions are not directly reflected in the sonic evolution of the piece.

Xenakis counteracted the ergodic tendency of the formal processes by placing salient events at key moments in the piece. Four processes stand out for their "orderly," improbable behavior, all of which are played by the harp. In measures 77–88, the harp plays a single pitch repeatedly. Interonset times and intensities are subjected to a time-varying correlated process. Dynamics start fortissimo and end piano; onsets begin at eight attacks per measure, slow down to one attack per measure, and accelerate to three attacks per measure. From measure 221 to 248, the harp plays a descending chromatic scale parsed into periodic attacks occurring on the first and second half note of every measure. The only consistent explanation for that event is that the composer mapped a very long glissando on the harp as discrete pitches played at regular intervals. This type of mapping does not occur anywhere else in the piece. The last note of the piece (measure 365) is also played by the harp, with a special type of playing technique. In the score, the composer requires the player to slightly touch the string with his nail – "la corde doit battre sur l'ongle qui l'effleure à son milieu." This produces a quite distinct buzzy timbre that sets this sound apart from its context, thus rhetorically foregrounding this action in a manner essentially alien to the rest of the piece (Example 2).

Example 2. A tool for testing streaming-fusion processes in *ST/10* textures (PATCHWORK).

2.3. Streaming

Roughly speaking, current theories of auditory processing place emphasis either on high-level cognitive mechanisms (see discussion in McAdams & Bigand, 1993; Leman, 1995), usually grouped as central processing, or on the transduction processes that take place in the peripheral auditory system. Presently, there is no coherent set of theories that has enough predictive power to explain most musically significant phenomena. Nevertheless, some advances have been made by working on small subclasses of musical practice, e.g., pitch in European tonal music of the eighteenth and nineteenth centuries (Krumhansl, 1990), and harmony in tonal music (Parncutt, 1989). However, typically the stimuli employed in such studies hamper their generality (see discussion in Trócoli & Keller, 1996; Keller, 2000a). Hierarchic pitch structures and meter-based periodicity are of questionable validity when one considers most recent music.

James Gibson's *ecological psychology* (1966) and Albert Bregman's *auditory scene analysis* (1990) provide perception paradigms that can be effectively applied to computer music analysis and composition. Beside its use in composition and sound synthesis (Keller & Truax, 1998), and timbre perception (Keller, 2000a; Keller & Berger, 2001), the ecological approach may prove fruitful in exploring complex temporal and spectral evolutions from an analytical point of view. However, for the purposes of this paper, auditory-scene-analysis concepts are sufficient. Bregman (1990) proposes two processes for sound segregation and fusion: sequential segregation – i.e., when events are parsed by their temporal proximity with other events – and simultaneous segregation – i.e., when sound complexes are separated in different layers depending on their spectral and temporal characteristics. These two produce the phenomenon commonly known as *streaming*. A typical example of streaming in music can be observed when several instrumental lines are combined. Different perceptual lines – or streams – are created depending on the pitch and the timbre of the notes. A similar effect can be observed in isolated sounds: when modulation is applied to parts of the spectrum while other components remain static, two different sounds emerge (McAdams & Bregman, 1979).

It should be clear why streaming is important in the analysis of instrumental textures composed of multiple events, such as those found in Xenakis's stochastic music: if perceptual grouping is dependent on temporal proximity and on the spectro-temporal characteristics of the events, then the perceptual result of combining various layers of instrumental gestures depends on the interactions among several parameters: interonset delay, pitch, timbre, intensity, and spectral change (for glissandi). The following empirical predictions can be made. Because the formation of streams depends mostly on pitch proximity and on timbral characteristics, high density of instrumental lines in itself is not enough to create a perceptual stream. In sparser textures, sequential and simultaneous segregation mechanisms compete. Thus, depending on their density, events may be grouped either by common spectro-temporal features and pitch proximity, or by temporal contiguity. Finally, sparser textures encourage groupings based on sequential-segregation cues.

2.3.1. Second-order sonorities

On the granular level (Keller, 1999; Keller & Truax, 1998), the parameters utilized in *ST/10* belong to two classes: local and global. Local parameters are the ones that can be defined at the level of the single grain, e.g., onset time, duration, frequency, amplitude. Global parameters are the result of local interactions, i.e., density, overlap, phase-synchronicity. Given that granular sounds require densities in the order of hundreds to thousands of grains per second, it is quite difficult to obtain consistent textures with traditional orchestral instruments played live. Composers such as Paul Dolden have used recordings of instrumental gestures. By overlapping hundreds of tracks, he obtained sounds close to a homogeneous mass of instrumental timbre that can be shaped with post-processing techniques (see Di Scipio's paper in this issue). Since Xenakis did not employ pre-recorded processed sounds, in *ST/10* several factors prevent the formation of he called *second-order sonorities*.

The composer suggested that second-order sonorities are the result of ordered organizations of sonic grains (Xenakis, 1992, p. 47). He argued that these types of sounds cannot be created with traditional instruments due to the limitations of human playing (Xenakis, 1992, p. 103). This is true for the number of instruments involved in *ST/10*, but does not hold for larger instrumental forces. Examples of second-order or mass sonorities can be found in the music of Ligeti and Penderecki, as well as in other Xenakis pieces for large orchestral forces. Nevertheless, it is true that a finer control over the timbral quality and the temporal evolution of these sonic masses can only be realized through digital transformations (Keller & Rolfe, 1998; Roads, 1997; Truax, 1988).

Granular sounds produced by the interaction of lower-level events are hardly present in *ST/10* because of the following reasons: (1) the mechanisms of parameter generation are not directly linked to their temporal evolution; (2) the use of memory-less, stochastic distributions does not take into account local dependencies, thus it precludes pattern-formation processes (Di Scipio, 1997); (3) densities are not sufficiently high; (4) higher-level timbral organizations are fairly static, encouraging perceptual parsing based on local characteristics, i.e., instrumental gestures, changes in dynamics and pitch; (5) durations are inversely proportional to densities, reducing the probability of overlaps among events.

2.3.2. Excitation, resonance and instrumental classes

The perceptual system tends to group sounds in an attempt to assign them to a common source (Bregman, 1990). Generally, common types of excitation appear as separate streams. Because instrumental timbres are uniformly distributed throughout *ST/10*, streaming cues are especially foregrounded. Thus, in the denser sections, we tend to hear several streams which roughly correspond to the characteristics of the excitation applied: plucked and struck sounds, sustained sounds and glissandi.

Undoubtedly, Xenakis realized the importance of playing technique when he determined his instrumental classes. He divided the strings into groups based on different articulations in playing, i.e., *pizzicato*, *col legno*, *tremolo*, and *glissando* (Xenakis, 1992, p. 139). Taking this method one step further, it would have been consistent to apply a tree-based approach to implement this classification. This would ensure a direct relationship between timbral distances and levels of the tree. The branching topology could be made to correspond to the available classes at each level. By distributing the articulations and the instrument classes in a systematic way, streaming and fusion could have been controlled more effectively.

Following previous work (Ferneyhough, 1995), a viable articulation-timbre classification scheme is represented in Figure 3. The highest level includes all articulation and instrumental classes. The next level branches off into three classes: continuous excitation, i.e., arco strings, winds; impulse excitation, pizzicato strings, harp, struck percussion; and iterated excitation, tremolo harp and strings, roll percussion, flutter-tongue winds. The third level parses common spectral characteristics. Thus we obtain a third classification level: within continuous excitation: glissando strings and winds, normal strings and winds; within impulse excitation: plucked harp and strings, struck percussion and strings; within iterated excitation: harp and strings, membrane percussion, wood percussion, and winds. Lower branching levels are possible, such as subdividing the winds into clarinets and horns, or establishing more subtle spectral differences such as bowed strings played *sul tasto* or *sul ponticello*.

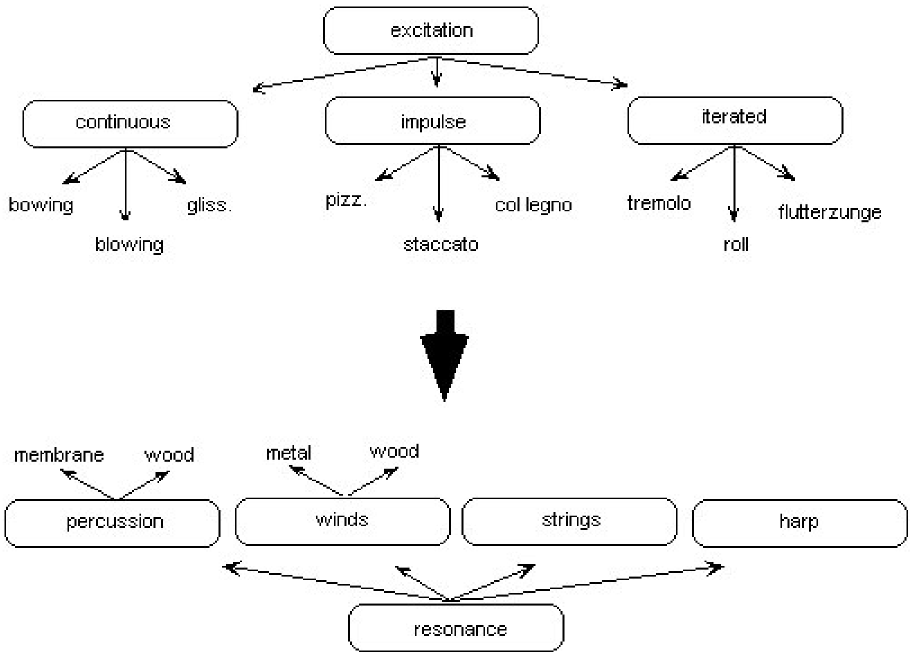

Fig. 3. Simple timbre classification based on the interaction between excitation and resonance properties (as applicable to *ST/10*).

Nevertheless, the scarcity of experimental data on timbre classification based on such attributes and the fact that pitch and intensity interact among them in the perception of timbre (Melara & Marks, 1990; Tróccoli & Keller, 1996), make these distinctions too fine for such complex textures as those found in *ST/10*.

3. Analysis by modeling

Compositional processes may be studied at various degrees of specificity. Depending on the degree chosen, different research strategies need to be employed. Common compositional practices are usually thought to engage planning, notation, and performance. Nevertheless, this view does not take into account the listener's own previous experiences with sound. These experiences shape our perception of musical works and evoke the mental, emotional, and physical sensations determined by our ongoing interaction with the environment (Keller, 2001). In a sense, to describe a musical piece based only on its notation, or even on the recorded waveform of its performance, on graphical representations of its spectral changes in time, on a map of the listener's neural activity, is like studying a rat's behavior by locking it up in a cell. The ecological validity of the analysis, i.e., its ability to reenact the history of perceptual experiences that determine how the piece is contextualized, is precluded. In other words, taking musical phenomena out of their context destroys the social and perceptual network of interactions that is activated when creating, playing, or listening to music. Unfortunately, this limitation is proper to most musical analysis practices. Although algorithmic models do not solve this problem, they may provide a more flexible method to test theoretical hypotheses and, most importantly, offer the possibility to *hear* the model's outcome. Thus, analysis by modeling extends the means to deal with this methodological gap serving as a complement to other analytical approaches.

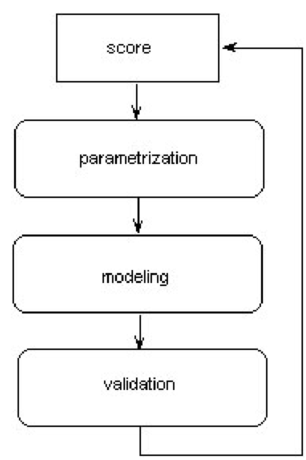

Fig. 4. The four stages of analysis by modeling.

The methodology proposed here consists of four stages, also sketched in Figure 4: (1) study of the score and its sonic realizations; (2) parameterization of the observed process; (3) model development; and (4) validation by comparison of the analytical observations with the model's outcome.

Analysis by modeling is especially useful for works that present widely varying results for each sonic realization of invariant underlying processes. This is the case in Xenakis's algorithmic music. The approach offers the possibility to explore the parameter space of the model under study. Furthermore, by keeping a modular design, various parametric combinations and interactions among models can be tested. The range of possible outcomes of a given procedure provides a qualitative description of the underlying processes. Since the piece is only a particular instance of the algorithmic mechanism, a larger database of parametric combinations proves more informative than the single example presented as the final rendering of the piece.

In the following sections, we present two example studies that address key issues in *ST/10:* the underlying temporal granularity, and the creation of fused granular textures as a function of event density. The first provides an aural example of the effect achieved by the multiple and overlapping quantization procedures used in the piece. The second presents several possible configurations of density and duration distributions. This experiment confirmed the hypothesis that slightly denser distributions would have produced fused granular textures.

3.1. Study 1: granularity

Quantization of random onset values produces an underlying rhythmic structure that determines the temporal granularity of *ST/10*. This temporal grid serves as a measure of the limits of rhythmic complexity and implies a certain regularity which would not be present had Xenakis mapped random durations directly onto the time axis. An obvious question, here, is what are the salient characteristics of this implied rhythmic structure.

Although Xenakis did not leave any documentation on this particular point, from the score it is clear that he used four beat-subdivision patterns: 3 : 4, 4 : 4, 5 : 4, and 6 : 4. The first section of *ST/10* has polyrhythms based on 4, 5, and 6 against 4. The first 3 : 4 pattern only appears in section 11, halfway into the piece. The layout of the subdivisions seems to depend on the onset values. In contrast with *Analogique A* (scored for nine string instruments, 1958) where three different subdivisions were kept in separate voices throughout the piece, *ST/10* presents all possible voice combinations among the given subdivisions.

We developed a simple model of the rhythm combinations produced by the quantization process (see Figure 5 and Example 3). The model, implemented in IRCAM's PATCHWORK environment (Laurson et al., 1996), produces four subdivisions, namely 3 : 4, 4 : 4, 5 : 4, and 6 : 4. In order to provide

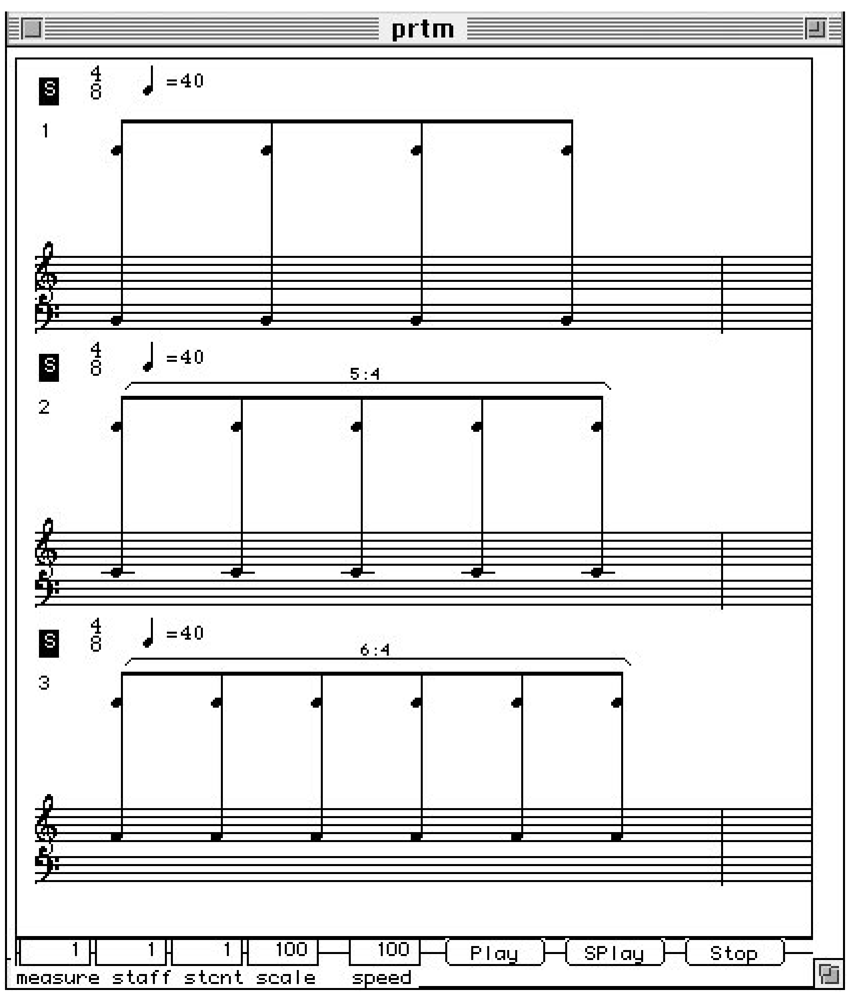

Fig. 5. Output of the model of underlying quantization processes in *ST/10*.

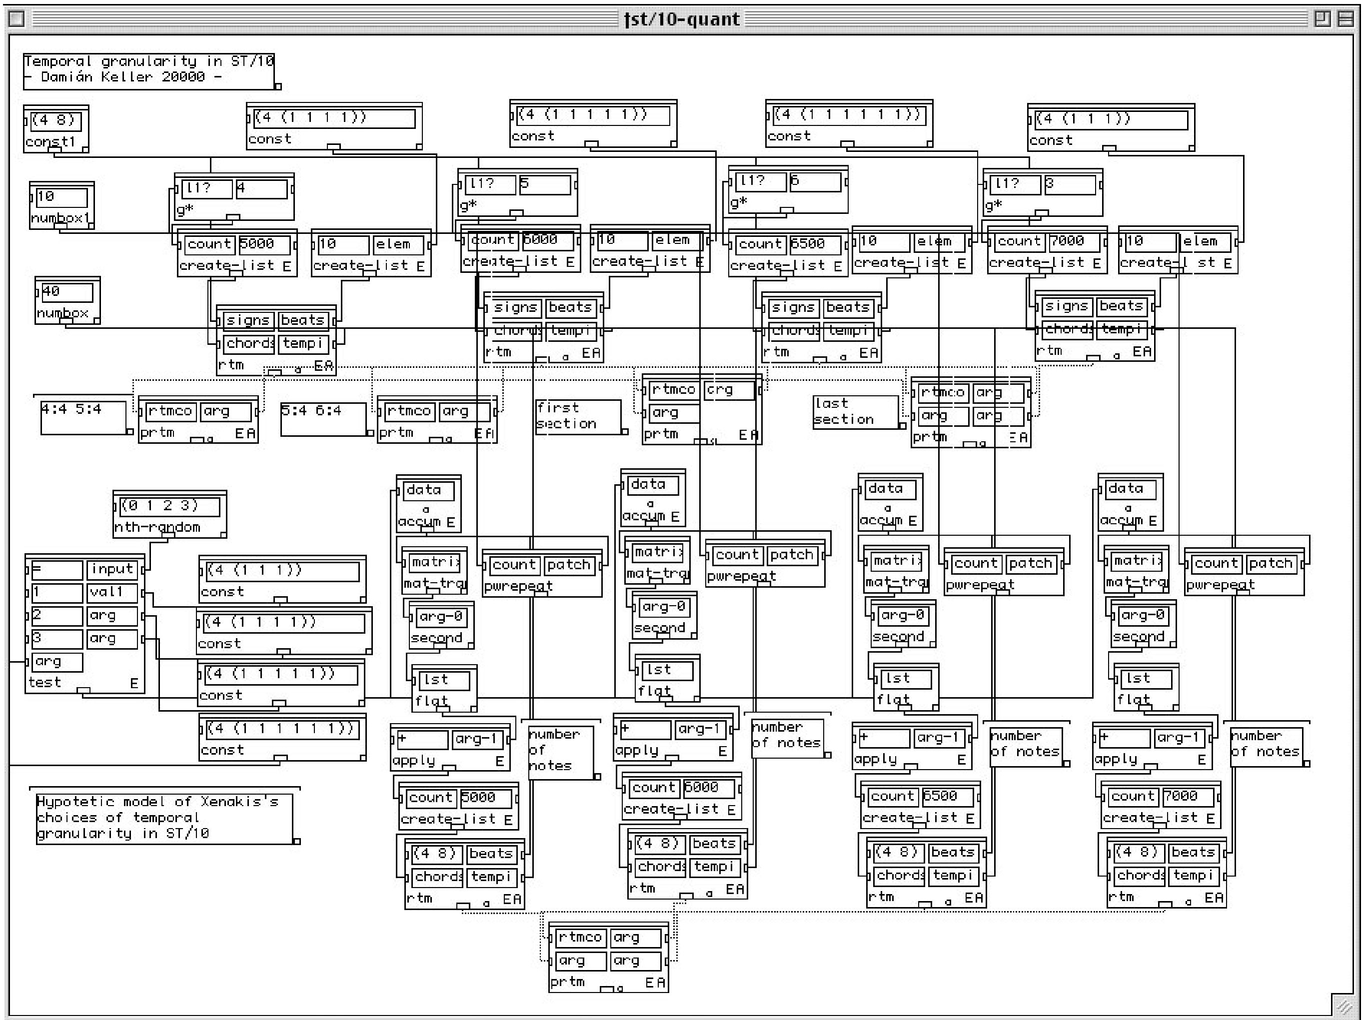

Fig. 6.

a straightforward representation of the phenomena, a simple rhythmic mapping is employed: one onset per subdivision. A single pitch is assigned to each of the voices. As observed in *ST/10*, all subdivisions are combined freely and any pattern can be played by any voice. Tempo and number of measures can be set by the user.

Example 3. The final section of *ST/10,* featuring a distinct harp gesture.

3.2. Study 2: density and fusion

The above discussion on second-order sonorities, or "emergent properties" (Di Scipio, 1997), has brought up several questions that can be answered experimentally. We will not attempt to address all of them here, but will present some preliminary results which illustrate how algorithmic models can help the analyst's work.

Mikhail Malt (1996) implemented a model of the pitch distribution in *Achorripsis*. The model follows the observations provided by Xenakis, so that the algorithm is also applicable to *ST/10* (except for microtones). We extended the *Achorripsis* model by adding a scaling factor for the distribution of durations. This is just a straightforward approximation to what Xenakis did, but the result is musically equivalent: high densities are coupled to short durations; low densities produce longer notes. Density is controlled by interonset differences, i.e., *delta-time* values. On average, duration is directly proportional to delay between events. The number of events is independent from *delta-time*.

Using the algorithm provided in Example 4, we manipulated the parameter controlling *delta-time* between notes within a range from 10 milliseconds to 500 milliseconds. Two types of sound were employed: synthetic piano and synthetic flute. The results show that, similarly to what happens in *ST/10*, average overlap remains constant across different densities, as was expected. The unexpected result was that when using short delays between notes – which are nevertheless very long in terms of granular sound – a single complex event was formed. In other words, notes were fused to create

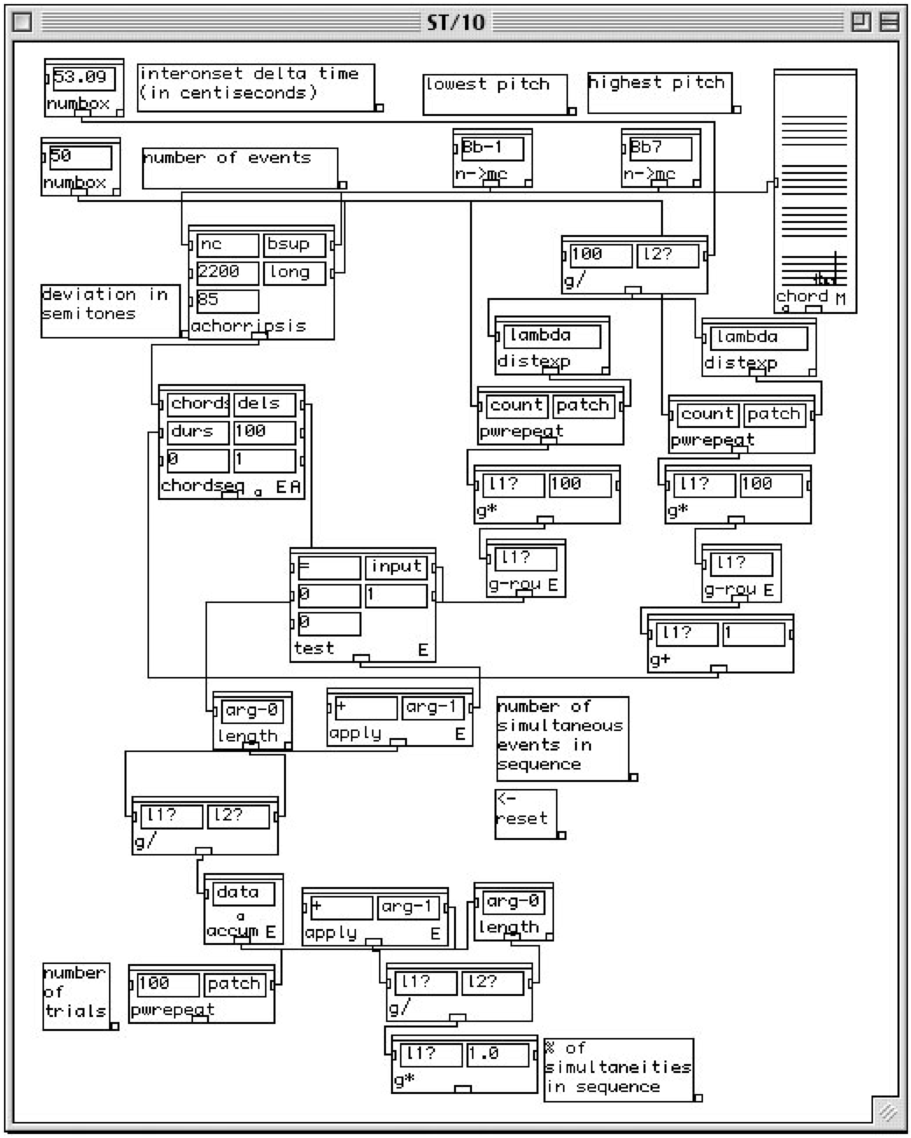

Fig. 7.

a time-varying, grainy texture. This suggests that even with widely scattered pitch material, Xenakis could have obtained his second-order sonorities merely by carefully distributing events among the ten available instrumental classes. Of course, timbre-based streaming mechanisms could come into play. But as long as timbral characteristics are kept closely related, fusion should persist.

Example 4. A dense section of *ST/10*, exemplifying the random distribution of instrumental classes.

3.3. Discussion of results

The *ST/10* model of temporal granularity provided a comprehensible aural representation of the implied rhythmic structures of the piece (Example 1). The goal was to formalize an *ad hoc* compositional procedure that Xenakis omitted from his documentation, namely the mapping of durations onto a metric grid. A further refinement could equip the model with a random generator of durations based on Xenakis's description. With this, durations could be quantized to fit the underlying metric grid.

A simplified model of event distributions produced results that are qualitatively similar to those observed in the *ST/10* score (Example 2). High-density textures featured short duration averages, and low-density ones presented longer durations. The model confirmed the hypothesis that slightly higher densities than those found in *ST/10* produce fused granular textures – at least in those passages where timbre characteristics are kept approximately uniform.

4. Summary

Xenakis's writings have exerted a deep influence on compositional practices of the last 15 years. The conceptual separation between model and implementation of a musical piece has fostered the development of computer-assisted composition and sound synthesis. Algorithmically generated parameters in musical design require mapping procedures between the model and the sound realization. As can be seen in *ST/10*, this methodology constrains and shapes the musical outcome in ways that are not accounted for in the pre-compositional model. As shown in the first study, these *ad hoc* procedures are clearly exemplified by the quantization process which gives rise to an implied metric structure.

As stated in his writings (Xenakis, 1971, 1985), Xenakis's compositional approach gives preeminence to out-of-time formal procedures over time-based and perceptual processes. In *ST/10*, parameters defined out of time are directly mapped onto temporal and spectral evolutions establishing a set of relationships free from references to the music's social environment. The ergodic distribution of events, the division into completely independent sections, and the utilization of time-invariant parameters preclude time-dependent sonic evolutions.

Had Xenakis organized his timbral palette by means of common excitatory and resonant characteristics, the phenomenon of streaming could have been exploited more systematically. Given the lack of a timbre-based sonic organization, the isolated instrumental gestures and a few "orderly" processes spring to the foreground in the sparse sections. In the denser sections, common gestural and timbral articulations are usually grouped into separate streams. Further study is necessary to determine what refinements are required to apply streaming-based organizational procedures. The second study illustrated above suggests that – if timbral features are kept similar – distributions slightly denser than those used in *ST/10* can produce fused granular textures.

Acknowledgements

Thanks to Agostino Di Scipio for his comments on the first version of this paper. All examples were implemented on PATCHWORK with the addition of the *Alea Library* (Malt, 1996). Please check the first author's website for additional information.

References

Bregman, A.S. (1990). *Auditory Scene Analysis: The Perceptual Organization of Sound*. Cambridge, MA: MIT Press.

Di Scipio, A. (1997). The problem of 2nd-order sonorities in Xenakis electroacoustic music. *Organized Sound*, *2*(3), 165–178.

Ferneyhough, B. (1995). Concerning the functional role of timbre in the early works of Anton Webern: a study of some aspects of the five orchestral pieces *Op. Posth.* (1913). In: J.M. Boros, & R. Toop (Eds.), *B. Ferneyhough, Collected* Writings. Amsterdam: Harwood Academic Publishers: pp. 166–182.

Gibson, J.J. (1966). *The Senses Considered as Perceptual Systems*. Boston, MA: Houghton Mifflin.

Keller, D. (1999). *touch'n'go: Ecological Models in Composition.* Master of Fine Arts Thesis, Simon Fraser University, Burnaby, BC. Available: http://www.sfu.ca/sonic-studio/

Keller, D. (2000a). *Beyond timbre space*. Lecture delivered at the Center for Computer Research in Music and Acoustics, Stanford University. Available: http://www-ccrma.stanford.edu/~dkeller

Keller, D. (2000b). Compositional processes from an ecological perspective. *Leonardo Music Journal*, *10*, 55–60.

Keller, D. (2001). Social and perceptual dynamics in ecologically-based composition. *Electronic Musicological Review*, *6*. Available: http://www.cce.ufpr.br/~rem/REMv6/Keller/SPD.html

Keller, D., & Berger, J. (2001). Everyday sounds: synthesis parameters and perceptual correlates. In: *Proceedings of the VIII Brazilian Symposium of Computer Music*. Fortaleza, CE, SBC. Available: http://www-ccrma.stanford.edu/~dkeller

Keller, D., & Capasso, A. (2000). Social and perceptual processes in the installation "The Trade." *Organised Sound*, *5*(2), 85–94.

Keller, D., & Rolfe, C. (1998). The Corner Effect. In: *Proceedings of the 12th Colloquium of Musical Informatics*, AIMI – Univerwity of Gorizia, Italy. Available: http://www-ccrma.stanford.edu/~dkeller

Keller, D., & Truax, B. (1998). Ecologically-based granular synthesis. In: *Proceedings of the International Computer Music Conference*, Ann Arbor, ICMA: pp. 117–120. Available: http://www-ccrma.stanford.edu/~dkeller

Krumhansl, C.L. (1990). *Cognitive Foundations of Musical Pitch*. New York: Oxford University Press.

Laurson, M., Rueda, C., & Duthen, J. (1996). *Patchwork (Version 3.0)* (Computer-assisted composition environment). Paris, IRCAM.

Leman, M. (1995). *Music and Schema Theory: Cognitive Foundations of Systematic Musicology*. Berlin: Springer.

Malt, M. (1996). *Alea Library – 3.0* (PATCHWORK Library). Paris, IRCAM.

McAdams, S., & Bigand, E. (1993). *Thinking in Sound*. Oxford: Oxford University Press.

McAdams, S., & Bregman, A. (1979). Hearing musical streams. *Computer Music Journal*, *3*(4): 26–44.

Melara, R.D., & Marks, L.E. (1990). Perceptual primacy of dimensions: support for a model of dimensional interaction. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 398–414.

Parncutt, R. (1989). *Harmony: a Psychoacoustic Approach*. Berlin: Springer.

Roads, C. (1997). Sound transformation by convolution. In: C. Roads, S.T. Pope, A. Piccialli, & G. De Poli (Eds.), *Musical Signal Processing*. Lisse: Swets & Zeitlinger: pp. 411–438.

Rolfe, C., & Keller, D. (1998). *MacPOD – 1.3* (Real-time granular synthesis software). Vancouver, BC: Third Monk Software. Available: http://www.thirdmonk.com

Smoliar, S. (1995). Parsing, structure, memory and affect. *Journal of New Music Research*, *24*(1): 21–33.

Tróccoli, B.T., & Keller, D. (1996). *A série harmônica: interpretação inter-disciplinar – física, psico-acústica, histórica – e o processo audio-perceptivo* (Technical Report 0222). Brasília, DF: Fundação de Apoio à Pesquisa do D.F.

Truax, B. (1988). Real-time granular synthesis with a digital signal processor. *Computer Music Journal*, *12*(2), 14–26.

Xenakis, I. (1967). *ST/10-1 080262* (Score). London: Boosey & Hawkes.

Xenakis, I. (1971). *Formalized Music: Thought and Mathematics in Composition*. Bloomington, IN: Indiana University Press.

Xenakis, I. (1985). *Arts/Sciences: Alloys.* New York: Pendragon Press.

Xenakis, I. (1992). *Formalized Music: Thought and Mathematics in Composition*. (Revised edition). Stuyvesant, NY: Pendragon Press.