Stochastic Synthesis: Origins and Extensions

Sergio Luque

2006

Institute of Sonology

Royal Conservatory

The Netherlands

I would like to thank Paul Berg for helping me so much with this paper

and

Kees Tazelaar for providing me with such great recordings

Contents

I Introduction

In the mid 1950s, Iannis Xenakis introduced the use of stochastic functions in musical composition. In the 1960s, he started using a computer to accelerate and automate the numerous operations that these methods require. At the same time, he theorized about the possibility of using stochastic functions to synthesize sounds.

It was during his tenure at Indiana University in Bloomington, from 1967 to 1972, that Xenakis first used a computer to synthesize sounds using probability functions. He continued with this research in Paris, from 1972 to 1977. Some of the results of the new sound synthesis technique were used in *Polytope de Cluny* (1972) and in *La Légende d'Eer* (1977).

It was not until the late 1980s that Xenakis continued with his research on stochastic synthesis. He developed an extension of the the stochastic synthesis procedure used in *La Légende d'Eer* (both procedures are known as *dynamic stochastic synthesis).* With the new version of the technique Xenakis composed *GENDY3* (1991) and *S.709* (1994).

In this paper, I present:

- some of the earliest microsound synthesis techniques (by Xenakis, Brün and Koenig) and the ideas behind them
- a full description of the two dynamic stochastic synthesis algorithms
- a new extension: the *stochastic concatenation of dynamic stochastic synthesis*

The available information about the dynamic stochastic synthesis algorithms is often incomplete, confusing or wrong. For the present paper, I decided not to elaborate on the conflicting information that is available through books and journal articles; instead, I tried to present a clear and general description of the techniques, as close as possible to the implementations used by Xenakis in his compositions.

My intention was to be able to relate the specific details of the original implementations with the sounds that Xenakis created with them and to facilitate further uses and developments of these techniques. I think that it is important to remember that for Xenakis these techniques were just "arbitrary starting points":

This approach can be compared to current research on dynamic systems, deterministic chaoses [sic] or fractals. Therefore, we can say it bears the seed of future exploration. (Xenakis 1992, p. 293)

II Early Microsound Techniques

1. Composing Sound with Musical Procedures, aka The Nonstandard Synthesis Approach

Since the late 1950s, with the appearance of computers with digital to analog converters, some composers have been interested in synthesizing sound through the manipulation of individual digital samples. Amplitude and duration values are obtained through musical procedures, they are not based on any acoustical model. This approach, often referred to as *nonstandard synthesis* (Holtzman 1978), reflects a willingness to explore the sound synthesis possibilities that are unique to computers.

In this chapter, I present three nonstandard synthesis approaches that appeared during the 1970s:

- **New Proposals in Microsound Structure** by Iannis Xenakis
- **SAWDUST** by Herbert Brün
- **SSP** by Gottfried Michael Koenig

These three approaches have the following goals in common:

- to unify the macrostructure and the microstructure of compositions
- to use synthesis techniques idiomatic to computers
- to start an experimental field in sound synthesis

2. **New Proposals in Microsound Structure.** Iannis Xenakis

The laws of the calculus of probabilities entered composition through musical necessity. (Xenakis 1992)

Stochastic Music

In 1954, Iannis Xenakis introduced the use of probability distributions in musical composition in order to control the orchestral sound masses of *Pithoprakta*. In 1956, he named this music *Stochastic Music* and enthusiastically set about investigating its possibilities.

There were many reasons why Xenakis was interested in the use of probability functions in music. For him, they were:

- a solution to what he called "the impasse of serial music":

  Linear polyphony destroys itself by its very complexity; what one hears is in reality nothing but a mass of notes in various registers. (Xenakis 1955)

  The composers [Stockhausen, Boulez and others] thought they were orthodox serialists but that was only true on paper. In reality they had mass events which they should have listened to in an unbiased manner. On the level of conscious thinking they should have introduced such notions as average density, average duration, colours and so on. (Varga 1996, p. 54)

- a technique to create and articulate sound masses inspired by the musical aspects of natural events, together with the recognition of the stochastic laws which govern them: "collisions of hail or rain with hard surfaces", "murmuring of pine-forests", "the song of cicadas in a summer field", "political crowds of dozens or hundreds of thousands of people". (Xenakis 1992, pp. 9, 237)

- an opportunity to incorporate concepts from modern science into the field of music composition. For example, the use of probability distributions in kinetic theory (Boltzmann and Maxwell) to determine the energy of a given quantity of gas: *I followed Maxwell's approach step by step: what he did with the molecules I did with the sounds.* (Varga 1996, p. 78)

- the problem: *What is the minimum of logical constraints necessary for the construction of a musical process?* (Xenakis 1992, p. 16)

- aesthetic laws useful in the search for the greatest possible asymmetry in all the levels of a composition, in order to negate "traditionally inherited behavioural frameworks" (sets of constraints and choices). (Xenakis 1992, p. 25)

So we have a formal archetype of composition in which the basic aim is to attain the greatest possible asymmetry (in the etymological sense) and the minimum of constraints, causalities, and rules. We think that from this archetype, which is perhaps the most general one, we can redescend the ladder of forms by introducing progressively more numerous constraints, i.e., choices, restrictions, and negations. (Xenakis 1992, pp. 23-24)

The ST Program

In 1962, Xenakis started using computers to accelerate the numerous calculations required by his stochastic approach to composition. He wrote the Stochastic Music Program in the FORTRAN programming language, running on an IBM-7090 at IBM-France. This program uses an algorithm that is an elaboration on the thesis of the minimum of constraints and rules utilized in *Achorripsis;* it employs interlinked probability functions to simultaneously determine the global structure (e.g., length of sections) and the note parameters (e.g., pitch, duration) of a composition. Xenakis considered this algorithm a musical form, like the fugue, general enough to create a large number of compositions for different instrumental ensembles and adaptable to the personality of different composers.

At the very same time, Xenakis speculated about the possibility of using stochastic techniques to synthesize sounds:

Although this program gives a satisfactory solution to the minimal structure, it is, however, necessary to jump to the stage of pure composition by coupling a digital-to-analogue converter to the computer. The numerical calculations would then be changed into sound, whose internal organization had been conceived beforehand. (Xenakis 1992, p. 144)

Xenakis wrote six pieces with the Stochastic Music Program: *ST/48, ST/10, Amorsima-Morsima, ST/4, Morsima-Amorsima* and *Atrées.* The program was also used to write some sections of *Eonta* (Varga 1996, pp. 101-102). After these pieces, he stopped composing with stochastic processes for several years:

I think I came up against the limitations of the method. I could have gone beyond these but more important problems came up which I had to solve . . . (Varga, 1996, p. 80)

General harmony? No, not yet. (Xenakis 1992, p. 182)

Towards a Metamusic?

Following the ST pieces, Xenakis set about searching for a unified axiomatic theory of music that would allow a formalization of all of the world's music: past, present and future. In order to achieve this, he studied Greek and Byzantine music (Xenakis 1992, p. 182); in particular, the musical writings of the peripatetic philosopher Aristoxenus of Tarentum. After this research, Xenakis developed two deterministic compositional procedures:

- Group theory for the construction of musical form, for example: the structure of *Nomos Alpha* was determined by the twenty four rotations of a cube.
- Sieves: logical formulas for the construction and permutation of sequences of integer intervals that can be applied to any set of musical parameters (e.g., pitches, durations, dynamics). In the early 1990s, Xenakis proposed the use of sieves for sound synthesis:

  [T]he amplitude and/or the time of the sound signal can be ruled by sieves. The subtle symmetries thus created should open a new filed for exploration. (Xenakis 1992, p. 276)

Stochastic Synthesis

It was during his tenure at Indiana University in Bloomington, from 1967 to 1972, that Xenakis first used a computer for sound synthesis. There, he started using stochastic processes again, while experimenting with new approaches to digital sound synthesis. In 1972, he continued these experiments at the Centre d'Etudes de Mathématique et Automatique Musicales (CEMAMu) in Paris. But in 1977, with the advent of the UPIC system[1], the stochastic synthesis research was postponed until the late 1980s. (Barthel-Calvet 2002)

Xenakis' first concrete ideas about stochastic synthesis were published in *Formalized Music* (Indiana University Press, Bloomington 1971), in the manifesto-like chapter "New Proposals in Microsound Structure".

He starts by rejecting:

- Fourier analysis as the basis for sound synthesis: *Now, the more the music moves toward complex sonorities close to "noise", the more numerous and complicated the transients become, and the more their synthesis from trigonometric functions becomes a mountain of difficulties, even*

[1] Unité Polyagogique Informatique du CEMAMu: a computer system with a graphic input device that enables the user to create sounds by drawing lines and shapes.

more unacceptable to a computer than the permanent states. It is as though we wanted to express a sinuous mountain silhouette by using portions of circles. (Xenakis 1992, p. 244)

- "pure" electronic sounds: *Any electronic music based on such sounds only, is marked by their simplistic sonority.* (Xenakis 1992, p. 243)
- serialism in electronic music: *The serial system, which has been used so much by electronic music composers, could not by any means improve the result, since it itself is much too elementary.* (Xenakis 1992, p. 243)

Instead, he advocates for:

- mixing "pure" electronic sounds with "concrete" sounds: *[Only then] could electronic music become really powerful.* (Xenakis 1992, pp. 243-244)
- the use of stochastic processes to efficiently produce sonorities with "numerous and complicated" transients: *it seems that the transient part of the sound is far more important than the permanent part in timbre recognition and in music in general. . . . The intelligent ear is infinitely demanding, and its voracity for information is far from having been satisfied. This problem of a considerable amount of calculations is comparable to the 19th-century classical mechanics problem that led to the kinetic gas theory.* (Xenakis 1992, p. 244)
- an approach in which sound synthesis is performed only in the time domain; starting directly from the sound pressure curves, defining them by means of stochastic variations: *We can start from a disorder concept and then introduce means that would increase or reduce it . . . We can imagine the pressure variations produced by a particle capriciously moving around equilibrium positions along the pressure ordinate in a non-deterministic way.* (Xenakis 1992, p. 246)

In the last part of the chapter, Xenakis proposes seven methods for stochastic microsound synthesis (Xenakis 1992, pp. 246-249):

1. Amplitude and/or duration values obtained directly from a probability distribution (e.g., uniform, Gaussian, exponential, Poisson, Cauchy, arc sin, logistic).
2. Combination of a random variable with itself by means of a function (e.g., addition, multiplication).
3. The random variables are functions of other variables (e.g., elastic forces,

centrifugal forces) or of other random variables (e.g., random walks).

4. The random variables move between two elastic barriers.
5. The parameters of a probability function as variables of other probability functions.
6. Combinations of probability functions (e.g., linear, polynomial). Composite functions (e.g., modulation).
7. Categorization of probability functions through at least three kinds of criteria (e.g. stability, curve characteristics).

For the construction of the macroform, he suggests using the methods previously presented in *Formalized Music:* the ST program (Chapter V), Markovian processes (Chapters II and III), symbolic music and group theory (Chapters VI and VIII).

Polytope de Cluny

Xenakis first used the results of his experiments in stochastic synthesis in the *Polytope de Cluny* (1972), a thirty one minute long multimedia work set in the Roman baths of Cluny in Paris, which consisted of music (recorded in seven channels and distributed over twelve loudspeakers), six hundred flashbulbs and three lasers (redirected with four hundred programmable mirrors). And he was proud to be the first in France to use digitally synthesized sounds (Harley 2004, p. 70).

Decorrelated stochastic synthesis opens the work (just after a brief introduction in the third channel, intended for when the audience entered the performance space), and it is present for about six minutes, sometimes in the foreground and sometimes receding to the background, as it is inhabited by the other sounds: ceramic windchimes, processed gongs, thumb pianos, low stringed instruments bowed with extreme overpressure, cardboard, wind and other sounds sources that are hard to identify. All the sonorities have a very rich spectrum and are full of buzzes, rattles and distortion. It is possible that a combination of stochastic synthesis with other sound sources had been used in the last sections of the piece.

Polytope de Cluny

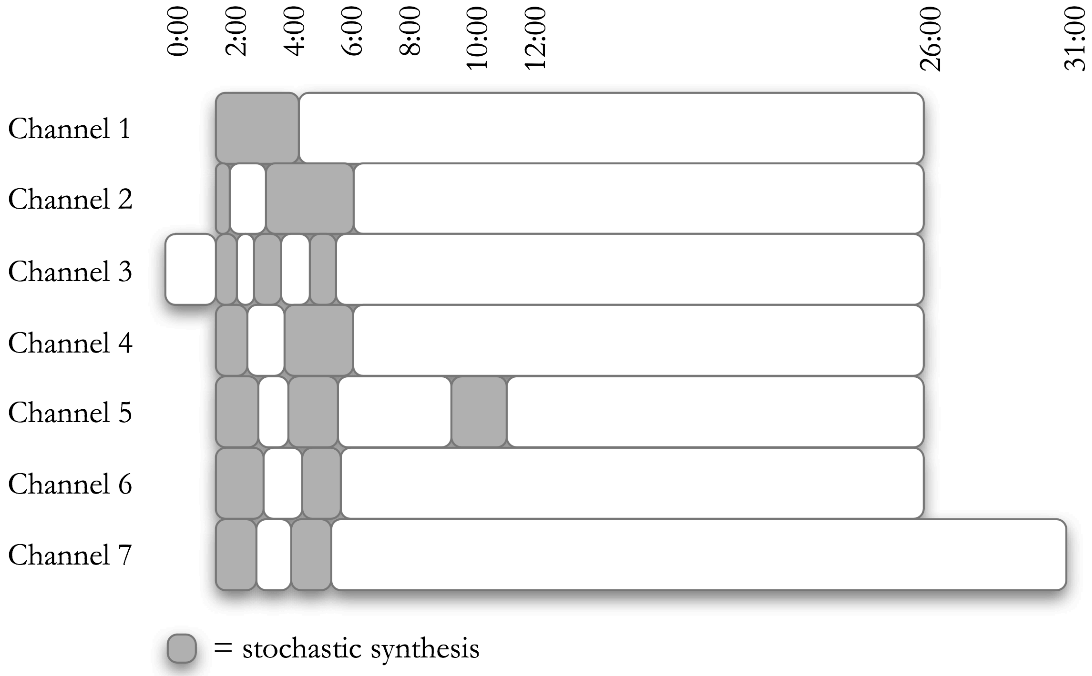

Channel 1:	1:43 - 4:36
Channel 2:	1:43 - 2:12, 3:27 - 6:28
Channel 3:	1:43 - 2:28, 3:03 - 3:58, 4:58 - 5:52
Channel 4:	1:43 - 2:48, 4:05 - 6:25
Channel 5:	1:43 - 3:12, 4:14 - 5:56, 9:51 - 11:45
Channel 6:	1:43 - 3:22, 4:42 - 6:02
Channel 7:	1:43 - 3:07, 4:19 - 5:42

Random Walks in Instrumental Music

Xenakis used the plotted graphs of stochastic synthesis in his instrumental music.

In *Mikka* (1971), *N'Shima* (1975) and *Mikka "S"* (1976), he read the horizontal axis of the graphs as time and mapped the vertical axis onto a grid of quarter-tone pitch values.

From the preface to the score of *N'Shima:*

The melodic patterns of N'Shima are drawn from a computer-plotted graph as a result of Brownian movement (random walk) theory that I introduced into sound synthesis with the computer in the pressure versus time domain. (Xenakis 1975)

Xenakis also used random walks in the following pieces: *Cendrées* (1973), *Phlegra* (1975), *Theraps* (1976), *Retours-Windungen* (1976), *Epeï* (1976), *Akanthos* (1977), *Jonchaies* (1977), *Ikhoor* (1978), *Dikhthas* (1979), *Palimpsest* (1979), *Anémoessa* (1979), *Mists* (1981), *Komboï* (1981), *Chant de soleils* (1983), *Tetras* (1983), and *Thelleïn* (1984). (Solomos 2001)

3. **SAWDUST.** Herbert Brün

If it can be shown that there exist significant musical ideas which require compositional thinking where not the sound but the waveform is the basic element and standard, then it can also be shown how the computer not only helps the composer to the fulfillment of up to now unfulfillable desires, but actually assists the composer in generating desires he never knew before. (Brün, H. and A. Chandra 2001)

SAWDUST is a computer program for composing waveforms developed by Herbert Brün at the University of Illinois in the mid-1970s.

The first version of SAWDUST was finished in 1976. It was written by Gary Grossman in the C programming language under the UNIX operating system, running on a PDP 11/50 at the University of Illinois Digital Computer Lab. With this version of the program Brün wrote: *Dust* (1976), *More Dust* (1977), *Dustiny* (1978) and *A Mere Ripple* (1979).

In 1980, Jody Kravitz added new functionality to the program, with this second version of SAWDUST Brün wrote: *U-TURN-TO* (1980) and *i toLD You so!* (1981). In the late 1980s, Keith Johnson ported SAWDUST to 16-bit personal computers, with this version of the program Brün created the tape parts for: *Aufhören!* (1989) for ensemble and tape, and *on stilts among ducks* (1997) for viola and tape. (Brün, H. and A. Chandra 2001)

The computer program which I called SAWDUST allows me to work with the smallest parts of waveforms, to link them and to mingle or merge them with one another. Once composed, the links and mixtures are treated, by repetition, as periods, or by various degrees of continuous change, as passing moments of orientation in a process of transformations. (Brün, H. and A. Chandra 2001)

SAWDUST is a deterministic and hierarchical approach to sound synthesis and music composition. It consistently uses linear change in its transformational procedures.

User's Task

First, the user specifies a set of *elements* or of lists of *elements*. An *element* has an amplitude and a duration in samples:

```
*    element                 # command to define an element
name:         e0             # set the identifier
amplitude=    100            # its amplitude
samples=      433            # its duration in samples
```

Then, the user defines a set of *links*. A *link* has a sequence of elements and the number of times it is to be played:

```
*    link
name:     l0                 # identifier
0:        e0                 # list of constituent elements
1:        e2
2:        e1
3:        e3
4:        e1
5:        e2
6:                           # end input of list with blank line
statements=  840             # number of iterations when played
```

Also, the user can define a set of transformations between the *links*:

- *mingle* concatenates a sequence of *links*:

  ```
  *    mingle
  name:     mg0          # name of mingle
  0:        l0           # constituent links
  1:        l17
  2:        l2
  4:
  statements=    841     # number of mingle iterations
  ```

  The output will be: *l0, l17, l2, l0, l17, l2, etc.*

  Each *link* will be played only the number of times it was specified in its definition. For example, if *l0* was specified to be played 5 times, *l17* 3 times and *l2* 2 times, then the ***mingle*** will be: *l0, l17, l2, l0, l17, l2, l0, l7, l0, l0.*

- *merge* interlaces the constituent elements of its *links.*
- *vary* gradually transforms one *link* into another by selecting polynomials that connect their elements (the polynomials can be of degrees 3 to 7). The two links do not need to have the same number of elements. The duration of the transformation can be specified in number of steps or in number of samples.

4. **SSP** (Sound Synthesis Program). Gottfried Michael Koenig

Musical sounds may be described as a function of amplitude over time. (Koenig 1971a, p. 93)

SSP is a computer program for sound synthesis designed by G.M. Koenig in the early 1970s. The development of this program was started by Koenig and finished by Paul Berg. The first working version of SSP was available in 1977, it was written in the MACRO-15 assembly language, ran on a PDP 15 at the Institute of Sonology in Utrecht (now in The Hague) and synthesized audio in real time. Three pieces were composed with this program: *Mandolin* by Paul Berg, *Blue Flute* by Robert Rowe and *One Room to Another* by David Theriault.

As opposed to programmes based on stationary spectra or familiar types of sounds, the composer will be able to construct the waveform from amplitude and time-values. The sound will thus be the result of a compositional process, as is otherwise the structure made up of sounds. The composer defines lists of data for amplitudes and time-values; these values will be put together by means of selection principles to form sound segments. Each segment begins and ends with an amplitude of zero. A permutation list, provided by the composer, determines the order of the segments, which may contain any number of repetitions. All known parameters of sound such as duration, dynamics or timbre thus become functions of the constructional principle. (Koenig 1971a, pp. 113-114)

Koenig, a composer of instrumental and electronic music, emerged from the serial school of composition that, during the 1950s, gathered around the electronic studio of the West German Radio (WDR) in Cologne. Some of the properties of SSP stemmed directly from the aesthetic interests shared then by this group: to unify the macrotime and the microtime domains of compositions by using the same governing principles to independently compose all the musical parameters (e.g., pitches, durations, dynamics, timbre, form). Sounds in the electronic studio were not based on acoustical models, but were created with the same compositional procedures used for the other parameters. This approach gave rise to a longing for control over the smallest characteristics of sound:

If the frequency of tuning is 440 cycles per second . . . The individual vibration period thus lasts 1/440th of a second. But the studio has not a device at its disposal which makes it possible to open a generator for this length of time, should one want to use a single period. Even if such a device were available, the tape would still have to be cut off 0.068 of an inch . . . . (Koenig 1959)

The implicit question rather arises as to how instrumental experience in macro-time (rhythmic relationships among parameter values) could be transferred to micro-time (timbre formation laws). (Koenig 1971b)

In SSP, compositional techniques were used to create digital sound from its microstructure: amplitude values (interpolated linearly) and duration values.

Functions (LIST, SELECT, SEGMENT & PERMUTATION)

The user started by defining the source material:

- a LIST of amplitudes: values from 0 to 4095 (usually, all the available values were selected)
- a LIST of durations: expressed in microseconds (with a minimum duration of 38 μs)

Then, values from both lists were SELECTED, generally all of them, and put into the "working area". SEGMENTS were created with amplitude and duration values from the "working area". Finally, sequences of SEGMENTS were arranged (PERMUTATIONS)

Selection Principles

Elements or groups of elements were defined with "selection principles": various types of random decisions inherited from Koenig's algorithmic composition program PR2. These rules were Koenig's generalizations of his own compositional practices and were originally conceived to choose between instrumental music parameters (e.g., harmony, orchestration, duration) (Berg 1979).

- ALEA: random selection
- SERIES: random selection, no element is repeated until all the available elements have been returned once; at that moment, the procedure starts again.
- RATIO: weighted random selection
- TENDENCY: random values chosen between boundaries which change in time
- SEQUENCE: values given in order (successive repetitions of individual values may be specified)
- COPY: all values are selected
- GROUP: values, selected by either ALEA or SERIES, occur in succession the number of times determined by either ALEA or SERIES

In a 1978 interview, Koenig said about SSP:

My intention was to go away from the classical instrumental definitions of sound in terms of loudness, pitch and duration and so on, because then you could refer to musical elements which are not necessarily the elements of the language of today. To explore a new field of sound possibilities I thought it best to close the classical descriptions of sound and open up an experimental field in which you would really have to start *again.* (Roads 1978)

5. Conclusion

When listening to the pieces that Brün composed with SAWDUST, most of the sound materials used in them fall into one of the following categories:

1. fixed pitches with constant amplitudes (sometimes used to create melodies)
2. gradual changes in pitch and/or amplitude (i.e., glissandi, crescendi and decrescendi)
3. fast interpolations between two sets of duration and/or amplitude values

These materials are used thematically: they are clearly established, varied and restated.

In my opinion, the most interesting sounds come from the third category, even though they share the stiffness in timbre that is characteristic of the SAWDUST paradigm: fixed waveforms, repeated, combined or interpolated.

The fact that the SAWDUST software is not capable of producing a tremendous variety of timbres is a restriction that Brün accepted, when he adopted the particular synthesis technique. However, it seems obvious that the composer's major goal was to compose interesting structures and forms rather than new, unheard sounds. (Blum 1979)

While it is true that when composing directly with individual sample points, sounds tend to be limited to a small group of timbres, the quantity and liveliness of these timbres can be substantially increased by:

- employing automatable selection procedures for constructing some or all of the levels of their structure, for example: random (as in SSP), chaotic, logical, combinatorial, etc.
- using stochastic strategies to control the amount of order or disorder at different levels of their structure (e.g. random walks, elastic barriers)

As Xenakis said in an interview in 1980:

The reason composers don't use [probabilities], even if they're interested in them, is because they didn't receive the necessary education. But scientists who also work with music don't use them either, even though they do possess the theoretical background. It's too far from what they mean by acoustics. That's one of the reasons why research in sound synthesis has come to an impasse, *even in the United States, with all its technological advantages. Scientists simply lack imagination in a field which lies outside mathematics or physics.* (Varga 1996, p.76)

Brün's microsound compositions make evident the big conceptual differences between his approach and the one by Xenakis or the one by Koenig: while Brün crafted melodies and variations, Xenakis longed to create automated music from a minimum number of constraints (Xenakis 1992, p. 295) and Koenig wanted to start a field of experimentation independent from traditional musical parameters.

Koenig's SSP provided a large number of possibilities for sound synthesis, as the three pieces that were composed with it can demonstrate. Some of the sound structures produced with this program have a harshness and timbral complexity that, within the somewhat limited timbral space of the non-standard approach, still today sound very attractive and compelling.

The compositional abstractions behind Koenig's algorithmic strategies were fundamental in the creation of a system that could on the one hand, produce the big amount of data required to synthesize interesting sound structures on a sample by sample basis, and, on the other hand, be general enough to accommodate different compositional interests and personalities.

I find it unfortunate that SSP was not used in more compositions and that its development did not continue.

Xenakis' approach differs fundamentally from SSP's and SAWDUST's in that, by using stochastic functions, he wanted to create complex sonorities in the most economical way and to continue investigating the thesis of the minimum of constraints that was behind *Achorripsis* and the ST Program (Xenakis 1992, p. 295).

In his music, Xenakis only used a few combinations of the methods for microsound synthesis that he put forward in the chapter "New Proposals in Microsound Structure" (*Formalized Music,* 1992). I think it would be rewarding to continue exploring new combinations of them; for example, just by reading through these methods, many possible extensions and variations of Xenakis' dynamic stochastic synthesis algorithm come to mind, some of them could yield interesting new sonorities and behaviours.

III Dynamic Stochastic Synthesis

1. Dynamic Stochastic Synthesis (1977): *La Légende d'Eer*

In 1977, Xenakis composed *La Légende d'Eer,* a forty three minute long tape piece commissioned by the West German Radio (WDR) in Cologne. This piece was the music part of *Le Diatope,* a multimedia work also including 1680 flashbulbs, 4 coloured lasers (reflected by 400 programmable mirrors) and a pavilion constructed from red vinyl stretched over a metal frame, somehow related in shape to his Philips Pavilion (Brussels World Exposition, 1958). This spectacle was commissioned as part of the events surrounding the inauguration of the Centre Georges Pompidou in Paris. (Harley 2004, p. 110)

Most of the sound materials used in *La Légende d'Eer* are very similar to the ones used in *Polytope de Cluny,* although a greater prominence is given to synthetic sounds: analog and digital (stochastic).

La Légende d'Eer

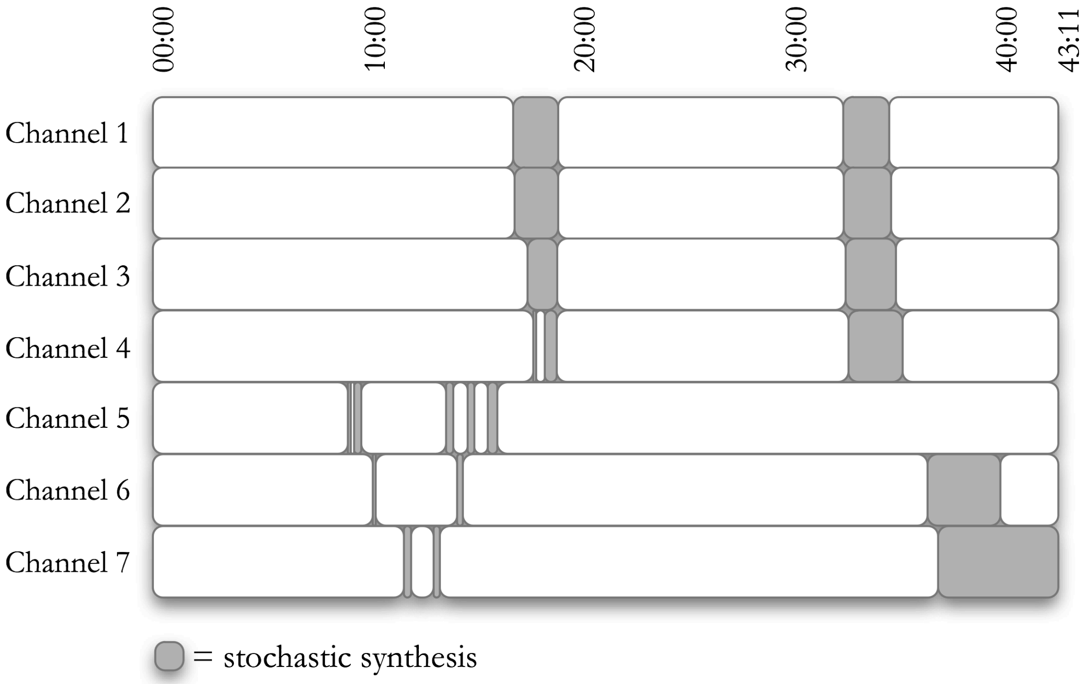

Channel 1: 17:18 - 19:26, 32:58 - 35:09
Channel 2: 17:22 - 19:25, 33:00 - 35:15
Channel 3: 17:58 - 19:24, 33:05 - 35:28
Channel 4: 18:14 - 18:23, 18:48 - 19:22, 33:13 - 35:47
Channel 5: 9:27 - 9:35, 9:45 - 10:05, 14:07 - 14:27, 15:07 - 15:27, 16:06 - 16:33
Channel 6: 10:38 - 10:47, 14:38 - 14:55, 36:58 - 40:27
Channel 7: 12:07 - 12:27, 13:30 - 13:48, 37:29 - 43:11

Reverberation and filtering were applied to some of the stochastic sounds present in the section that starts at 32'58", in the first four channels. (E.g., second channel, from 33'20" to 33'52").

In this piece, Xenakis started using a new technique for stochastic synthesis that he named *Dynamic Stochastic Synthesis.* This technique had its origin in the methods presented in the chapter "New Proposals in Microsound Structure" (*Formalized Music,* 1992) and introduced an important conceptual development: the waveform as the basic unit to be varied stochastically at each iteration.

In this model, waveforms are constructed by linearly interpolating a set of breakpoints. Each breakpoint is defined by a pair of duration and amplitude values. At every repetition of the waveform, these values are varied stochastically using random walks: any probability distribution can be employed to determine the size and direction of the steps (e.g., uniform, Gaussian, exponential, Poisson, Cauchy, arc sin, logistic, nested distributions). There are as many pairs of duration and amplitude random walks as there are breakpoints in the waveform. (Xenakis 1992, pp. 291-293)

The fluctuation speed of a parameter is directly proportional to the step size of its random walks: the smaller the steps, the slower the rate of change in that parameter. Depending on their speed, the perception of these fluctuations in duration and/or amplitude can be located on a continuum ranging from slow glissandi and subtle variations in timbre to noise.

Each random walk is forced to remain within a predefined space by means of two elastic barriers that reflect excessive values back into the barrier range. These barriers provide control over the frequency and amplitude of the waveform, for example:

- the larger the space between a pair of barriers, the bigger the variation that is possible in that parameter (i.e., the bigger the potential sizes of glissandi and the amplitude of the waveform)
- if the two elastic barriers of a parameter are set to the same value, that parameter will be immutable (i.e., fixed pitch, constant amplitude)
- if the two elastic barriers of the duration random walks are set to the same value and the amplitude values fluctuate slowly (i.e., amplitude random walks with small step sizes), then gradual and independent variations in the amplitude of the overtones of a fixed pitch are heard

Previously, Xenakis worked with individual duration and/or amplitude values that were either independent or dependent on the preceding value (e.g. random walks). The new approach evidences Xenakis' interest in having a finer control over the periodicities (duration) and symmetries (amplitude) of stochastic waveforms. According to Xenakis, this control would allow him to modulate from white noise to a square wave, with "melodies, symphonies, natural sounds" in between (Xenakis 1992, p. 289). As usual, his goals were high:

Following these principles, the whole gamut of music past and to come can be approached. (Xenakis 1992, p. 289)

In my opinion, the aforementioned statement is extremely unrealistic. Nonetheless, the Dynamic Stochastic Synthesis model is a very important development and is capable of producing rich and lively sounds that would be difficult to obtain through other sound synthesis techniques.

Dynamic Stochastic Synthesis (1977): Method

1. Select the number of breakpoints for the waveform. For example: 3

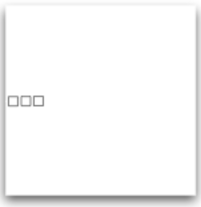

2. Select the waveform's minimum and maximum frequencies and convert them to duration in number of samples. For example:
   368-735 Hertz = 120-60 samples

3. Divide the minimum and maximum number of samples by the number of breakpoints:
   60/3 = 20
   120/3 = 40

   These values are the barriers for the duration random walk of each breakpoint.

4. For the continual generation of steps for all the duration random walks: select a probability distribution, its parameters and the ± number that will be the minimum and maximum size for these steps.

5. An initial duration is given to each breakpoint: values taken from stochastic or trigonometric functions, the minimum or the maximum duration, etc.

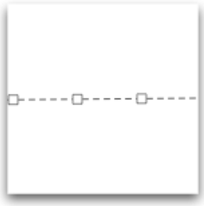

6. A maximum amplitude is selected and this ± value is the barrier for the amplitude random walk of each breakpoint.

7. For the continual generation of steps for all the amplitude random walks: select a probability distribution, its parameters and the ± number that will be the minimum and maximum size for these steps.

8. An initial amplitude is given to the each breakpoint: values taken from stochastic or trigonometric functions, zeroes, etc.

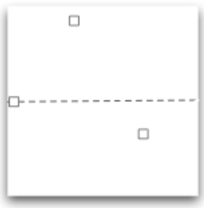

9. Breakpoints are linked by linear interpolation. At each repetition, the last breakpoint of the current waveform is connected with the first breakpoint of the next variation of the waveform.

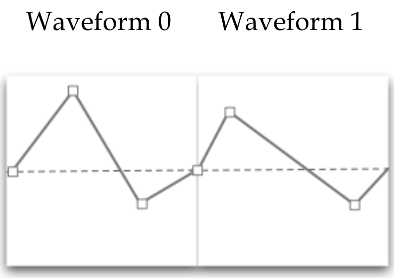

This technique is described in the chapter "Dynamic Stochastic Synthesis" of *Formalized Music* (Xenakis 1992, pp. 289-293) and is often mistaken as the explanation for the dynamic stochastic synthesis algorithm implemented in the beginning of the 1990s, as part of the GENDY program.

Also, it is important to remember that, for Xenakis, this method was just an arbitrary starting point which he used in *La Légende d'Eer*:

This approach can be compared to current research on dynamic systems, deterministic chaoses [sic] or fractals. Therefore, we can say it bears the seed of future exploration. (Xenakis 1992, p. 293)

2. Dynamic Stochastic Synthesis (1991): *More Thorough Stochastic Music*

In an interview that took place in the mid-1990s, Xenakis said:

During my initial tests, I realized that probabilities could yield rich sonic results, but you have to control them – they are like wild horses! I have been working like a laborer to obtain interesting things from the [GENDY] program. I have been obliged to throw away many experimental results and keep only those that interested me. (Robindoré 1996)

It was not until the late 1980s, with the loan of a personal computer from Hewlett Packard, that Xenakis continued with his research on stochastic synthesis (Harley 2004, p. 215). He wrote a program that implemented an extended version of the dynamic stochastic synthesis algorithm used in *La Légende d'Eer* (1977). This program was written in the BASIC programming language, with the assistance of Marie-Hélène Serra, and was called GENDY (a portmanteau constructed from the French words *generation* and *dynamique*).

The only difference between the new implementation of the algorithm and the previous one is the use of second order random walks. A second order random walk consists of three elements: a probability distribution and two random walks. The probability distribution generates the step sizes of the first random walk (the *primary random walk);* the successive positions of the *primary random walk* are the step sizes of the *secondary random walk*. The successive positions of the *secondary random walk* are the values of the second order random walk.

[F]or the duration abscissa a probability distribution and 2 times 2 elastic mirrors; for the amplitude ordinates a probability distribution and 2 times 2 elastic mirrors. (Xenakis 1992, p. 304)

A second order random walk with elastic barriers behaves very differently than a first order one:

- A first order random walk oscillates around an equilibrium position that changes arbitrarily over time. Sudden changes in position happen when several consecutive steps in the same direction are taken or when a step with an atypical size is made.

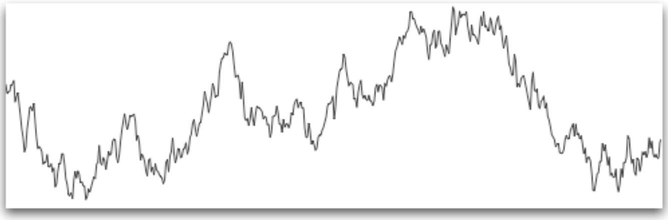

First order random walk with two elastic barriers

- A second order random walk gravitates around one of its two barriers:
    - if the position of its primary random walk is positive, it gravitates around the upper barrier
    - if the position of its primary random walk is negative, it gravitates around the lower barrier

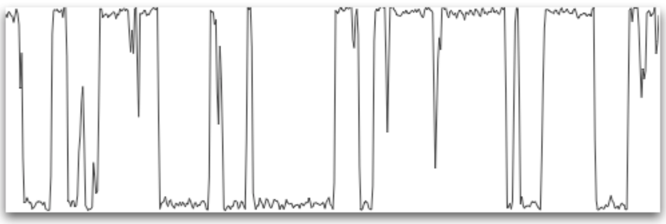

Second order random walk with two pairs of elastic barriers

Dynamic Stochastic Synthesis (1991): Method

1. Select the number of breakpoints for the waveform. For example: 3

2. Select the waveform's minimum and maximum frequencies and convert them to duration in number of samples. For example:
   368-735 Hertz = 120-60 samples

3. Divide the minimum and maximum number of samples by the number of breakpoints:

60/3   = 20

120/3 = 40

These values are the barriers for the secondary duration random walk of each breakpoint.

4. Select a ± number that will be the minimum and maximum size for the primary duration random walks, which will give the step sizes for the secondary duration random walks.

   Usually, this number is equal or smaller than the size of the space between the barriers of the secondary duration random walks. For example: a number between 0.0 and 20.0, if the secondary duration random walks have their barriers at 20.0 and 40.0.

Secondary random walk barriers:	20 – 40
Primary random walk barriers:	-10 – 10

   It is more convenient to think of frequency intervals in terms of proportions rather than in terms of number of samples. So, in order to have a more pertinent control of the size of the primary duration random walk space, I would suggest to calculate it as a ratio to the secondary duration random walk space. For example:

Secondary random walk barriers:	20 – 40
Secondary random walk space:	20
Primary random walk space as a ratio:	**0.5**
Primary random walk barriers:	-10 – 10

5. For the continual generation of steps for the primary duration random walks: select a probability distribution and the ± number that will be the minimum and maximum size for these steps (it is recommended to calculate this value as a ratio to the primary random walk space).

6. A ± maximum amplitude is selected; this ± value is the barrier for the secondary amplitude random walk of each breakpoint.

7. Select a ± number that will be the minimum and maximum size for the primary amplitude random walks, which will give the step sizes for the secondary amplitude random walks.

Usually, this number is equal or smaller than twice the maximum amplitude.

Secondary random walk barriers: -0.5 – 0.5
Primary random walk barriers: -0.1 – 0.1

It is also recommended to calculate this ± size as a ratio to the secondary random walk space.

8. For the continual generation of steps for the primary amplitude random walks: select a probability distribution and the ± number (or ratio) that will be the minimum and maximum size for these steps.

9. Breakpoints are linked by linear interpolation. At each repetition, the last breakpoint of the current waveform is connected with the first breakpoint of the next variation of the waveform.

This technique is described in the chapter "More Thorough Stochastic Music" of *Formalized Music* (Xenakis 1992, pp. 295-322) and is the one used by Xenakis in *GENDY3* (1991).

GENDY3 (1991)

To write a piece with the sounds produced by the GENDY program, Xenakis wrote another program, called PARAG, that treated a number of outputs of the GENDY program as voices and created sections with them.

The number of voices and their parameters (e.g., number of breakpoints, position of the elastic barriers) seem to have been selected by hand by Xenakis and were hardcoded into the PARAG program. The procedure that he used to obtain these values is not known, but it is probable that he determined them intuitively.

In a page from Xenakis' sketchbook at CEMAMu (Hoffmann 2004, p. 139), he listed the parameters of all the voices for a PARAG section and added references about their resulting pitch characteristics: precise pitch or register or direction (e.g., A6 one quarter tone high, "very low", "descending"). This could be an indication of his interest in directly controlling the harmonic content of each section.

The PARAG program also independently broke each of the voices into parts, called *time* *fields* by Xenakis. The duration of each *time field* was determined using an exponential distribution; and a Bernoulli trial[2] decided if a *time field* was silent or not (Xenakis 1992, pp. 296-304; Serra 1993, pp. 252-253). The length of each section was determined by the length of the longest of its voices.

In a 1996 interview, Xenakis elaborated on the difficulty of composing with the GENDY and the PARAG programs:

I am always trying to develop a program that can create the continuity of an entire piece. This is a struggle, because there are always parts that you prefer over others. So you have to change them, to stop the process, start some other one, and then put these two different ones together. This can be taken very far. As I move toward multiple voices, the problem becomes even more complicated. . . . The final result is an edifice realized almost entirely by probability calculations. It takes time to successfully realize the probability calculations on the architectural level. The work is always intuitive. You are lost if you base yourself only on the calculations, unless they themselves are also intuitive. (Robindoré 1996)

Xenakis composed two works with this version of the PARAG and the GENDY programs, each about twenty minutes long: *GENDY301* and *GENDY3* (Serra 1993, p. 239).

GENDY301 was premiered in October 1991, at the International Computer Music Conference in Montreal, Canada. In spite of having been analyzed and cited in many articles (e.g. Di Scipio 1998; Hoffmann 2000), this work seems to have been withdrawn from Xenakis' catalog: it is not available through *Editions Salabert,* Xenakis' publishers at that moment; a recording of it was never released; and the information about it is scarce and contradictory. Unfortunately, I have not been able to find a recording of this piece.

GENDY3 was premiered in November 1991, at the Journées de Musique Contemporaine in Metz, France. This work was released on compact disc in 1994. This is a different piece than GENDY301: "a newly generated work derived from similar control data" (Harley 2004, p. 216).

In spite of the close relation of the two pieces in terms of their genesis, they in fact sound completely different. . . . [GENDY301] exhibits a wider dynamic range that the other better-known piece, with extremely loud textures entering suddenly on top of narrower-range sonorities. In addition, it contains more breaks of

[2] *The Bernoulli distribution is a discrete distribution having two possible outcomes labelled by* $n = 0$ *and* $n = 1$ *in which* $n = 1$ *("success") occurs with probability* p *and* $n = 0$ *("failure") occurs with probability* $q = 1 - p$, *where* $0 < p < 1$. . . . *The distribution of heads and tails in coin tossing is an example of a Bernoulli distribution with* $p = q = 1/2$. (Weisstein 2006)

silence. Xenakis never stated his dissatisfaction with the piece, but he may have decided not to release a "family" of works as he had produced in 1962 with the data from his ST program. (Harley 2002, p. 55)

GENDY3 consists of a sequence of eleven PARAG sections, of about two minutes in length each.

In *Formalized Music,* Xenakis proposed that an arbitrary chain of PARAG sections could produce an interesting musical composition (Xenakis 1992, p. 296).

In the case of GENDY3, the arrangement of these sections does not give the impression of being arbitrary; adjacent sections are clearly separated from each other by the considered use of contrasting material: changes in register, number of voices, barriers, step sizes, etc.

Each section has a consistent and, to a certain degree, static behaviour, as the settings of all the parameters of a section do not vary over time; the abrupt changes from section to section give a sense, or illusion, of progress to the composition. This approach to form is somehow related to the use of juxtaposed material in some of the later works of Morton Feldman (e.g., *Patterns in a chromatic field).*

Most sections of GENDY3 have a combination of fixed pitches, glissandi and noise. The fixed pitches of a section, due to the insertion of silences, create a texture that could be described as a stochastic ostinato. Also, by using the same type of behaviour in all the voices, Xenakis created some homogeneous sections, for example:

- Section IV (4:58 – 6:28): noise (i.e, very fast fluctuation in pitch)
- Section IX (13:50 – 15:49): glissandi
- Section XI (17:06 – 18:53): a cluster of fixed pitches

In GENDY3, Xenakis made a stereo file by joining the mono outputs of two runs of the same PARAG program (he added a delay of about 100 milliseconds between them).

In *Formalized Music,* Xenakis proposed a similar method to achieve a multichannel output: to compute the same PARAG program, as many times as there are channels, but with a different random seed for each of the amplitude and/or duration random walks (Xenakis 1992, p. 298). I think that it would be interesting to implement this suggestion and to evaluate its various possibilities.

S.709 (1994)

In a 1997 radio interview for the BBC, Andrew Sparling asked Xenakis about what he perceived as the "melodic nature" of *S.709*:

Xenakis: *All my music has something random in it.*

Sparling: *But... to me it was a melody, a very melodic piece.*

Xenakis*: It was a melody, yes, sometimes it makes melod[ies]. . . . [It] is like when you walk on a rocky mountain, there are all rocks around you, you don't remember the rocks that you just passed by. So this is the same thing happening. But if you have nice lines - when I say "nice" that is "longer lines" - then you can remember what happens in between.*

After composing GENDY3, Xenakis extended the GENDY program. He added the possibility of modulating the parameters of the dynamic stochastic synthesis algorithm with simple shapes: sine, sawtooth. With this version of the program, Xenakis created *S.709* (Hoffmann 2000, p. 31). Unfortunately, there is very little information available about this piece.

S.709 was premiered at a concert a La Maison de Radio-France in December 1994 (Harley 2004, p. 217). Its title stands for: *Sequence 709*. *Sequence* was the name that Xenakis gave to the sections created by the PARAG program. (Xenakis 1992, p. 296). In the radio program mentioned above, Brigitte Robindoré, head of musical production at Les Ateliers UPIC, said about *S.709*: "It's unedited. It's unrestrained*"*. So, it could be inferred that *S.709* consists of the output of only one PARAG program: it could be a PARAG section of seven minutes in length. This hypothesis could also account for the lack of differentiated sections in this piece.

The rapid and periodic modulation of the stochastic synthesis parameters creates voices that are constantly and widely fluctuating in pitch, amplitude and timbre. Their sound is extremely different from the sustained and continuous sonorities used in *GENDY3*. In the latter piece, harmonies and stochastic ostinati were only possible because the synthesis parameters remained fixed throughout a section.

As a consequence of the sonic complexity produced by the new modulated version of the stochastic algorithm, in *S.709* there are fewer simultaneous voices than in *GENDY3:* no more than four and, for most of the time, only one.

In the same radio interview, Robindoré mentions that *S.709* "produces quite a polemical reaction in the audience". And this is not surprising, this work is extremely original in its materials and in its construction; it does not resemble any other piece by Xenakis, nor any other piece that I have ever heard.

Erod (1997)

Brigitte Robindoré on *Erod (ibid.):*

We extracted waveforms from GENDY samples, and we generated several UPIC sequences which where then mixed with acoustic sounds which were processed on the Macintosh. It takes the UPIC branch, the GENDY branch and then the electroacoustic, "musique concrète", branch. And they all three coalesce into this one piece.

Erod is a five minute long work that was commissioned by the Bath Festival in England, where it premiered in May 1997. It was composed using the PC version of the UPIC system. Because of health problems, Xenakis could not complete this piece by himself and Brigitte Robindoré produced much of the music for him. In gratitude, Xenakis derived the title of the piece from her surname. After its premiere, *Erod* was withdrawn from Xenakis' catalog (Harley 2004, p. 219).

The contribution of Ms. Robindoré to the production of the piece was so integral that Xenakis could not in good faith put his name on it. (Harley 2002, p. 54)

Even after repeated listenings, it is still not clear to me which parts of *EROD* were derived from GENDY samples. Although some of the glissando gestures in this piece, curiously enough, sound similar to some of the gestures that Brün created with SAWDUST (compare the beginning of *EROD* with the beginning *Dustiny* (1977)). These sounds could have originated from GENDY waveforms; the repetition and linear modification of GENDY waveforms in the UPIC could produce similar results to the ones obtained with SAWDUST.

IV Stochastic Concatenation of Dynamic Stochastic Synthesis

Any theory or solution given on one level can be assigned to the solution of problems of another level. Thus the solutions in macrocomposition (programmed stochastic mechanisms) can engender simpler and more powerful new perspectives in the shaping of microsounds. (Xenakis 1992)

After I wrote my implementation of Xenakis' dynamic stochastic synthesis (1991) algorithm in SuperCollider Server[3], I wanted to find ways of extending this model. Among the techniques that I first tried were: to modulate the stochastic synthesis parameters, à la *S.709,* and to use second order random walks to modify the curvature of each of the interpolations between the breakpoints of a waveform. The initial results of these techniques were interesting, but I have not continued investigating their possibilities because I have been focused on another procedure that, almost immediately after its first implementation, yielded very promising results: the stochastic concatenation of GENDYs (i.e., the dynamic stochastic synthesis algorithm from 1991).

In this technique, a Waveform is constructed by concatenating the waveforms of a set of GENDYs, one iteration at a time. For example, a sequential concatenation of two GENDYs:

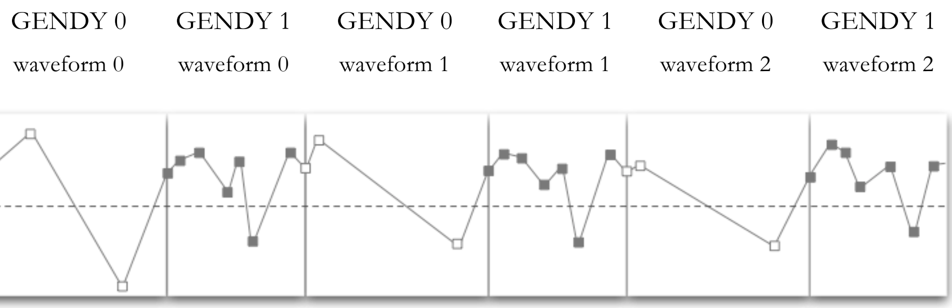

Conceptually, there is no limit to the number of GENDYs in a set, but in my SuperCollider implementation of this technique I have found that 72 is a reasonable limit.

GENDYs are selected from the set through different stochastic procedures (e.g., weighted probabilities, tendency masks, random walks). This approach is very close to

[3] The standard distribution of SuperCollider Server comes with three unit generators that are based on the dynamic stochastic synthesis algorithm: *Gendy1.ar, Gendy2.ar* and *Gendy3.ar*. According to their author, Nick Collins, these implementations offer several generalizations and variations of Xenakis' original algorithms. I did not use these implementations.

SSP's use of *Selection Principles* for the creation of *Permutations*.

The sounds produced by the stochastic concatenation of GENDYs range from continuous textures to differentiated arrangements of microsounds, and are usually lively and rich in harmonics, though limited to the timbral space of the non-standard synthesis approach.

Stochastic Selection

Any stochastic procedure can be used for selecting GENDYs from a set. Each procedure will give its own character or behaviour to the resulting sound.

These are the most fruitful stochastic procedures that I have used:

1.- **Random**

Uniform random selection. The resulting sound is continuous and unstable.

2.- **Series**

Random selection. After a GENDY has been selected, it cannot be selected again until all the available GENDYs have been selected once; at that moment, the procedure starts again.

The resulting sound is continuous, stable and with semi-periodicities.

3.- **Weighted**

Weighted random selection. GENDYs are selected from the set at random using a list of weights that specify their relative frequency of occurrence.

The resulting sound is continuous and unstable.

4.- **Size vs. Probability**

Weighted random selection. The relative frequency of occurrence of each GENDY is inversely proportional to its duration. With this procedure, all the GENDYs, independently of their length, occupy a similar proportion of time.

The resulting sound is continuous and unstable.

5.- **Tendency Masks**

Random selection between boundaries which change in time. This procedure continuously creates differentiated sections that are separated by noisy transitions.

The resulting sound is discontinuous and its degree of stability constantly changes.

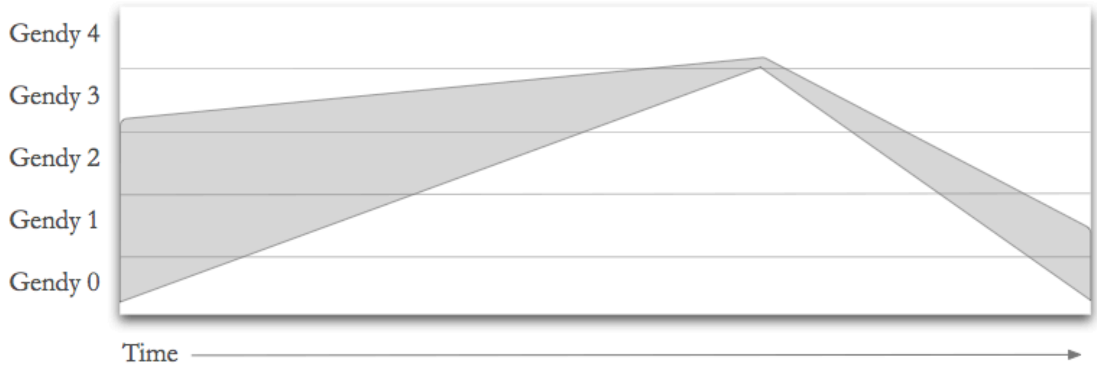

A tendency mask for the random selection of GENDYs

6.- **Markov Chains**

Markov chains of any order can be used to select GENDYs from a set. Usually, as it is very difficult to anticipate the result of this procedure, the possibilities of a set are explored by recurrently filling the transition tables with random values.

The resulting sound is continuous and very unstable.

Next GENDY

Current GENDY		GENDY 0	GENDY 1	GENDY 2	GENDY 3	GENDY 4
	GENDY 0	0.03	0.16	0.15	0.29	0.37
	GENDY 1	0.27	0.01	0.29	0.24	0.19
	GENDY 2	0.13	0.21	0.47	0.03	0.16
	GENDY 3	0.28	0.26	0.13	0.25	0.08
	GENDY 4	0.12	0.2	0.29	0.19	0.2

A first order transition table

7.- **Sequences of Second Order Random Walks**

GENDYs are selected by cycling through a sequence of second order random walks. Each random walk selects one GENDY at a time. This is one of the most successful procedures for the concatenation of GENDYs.

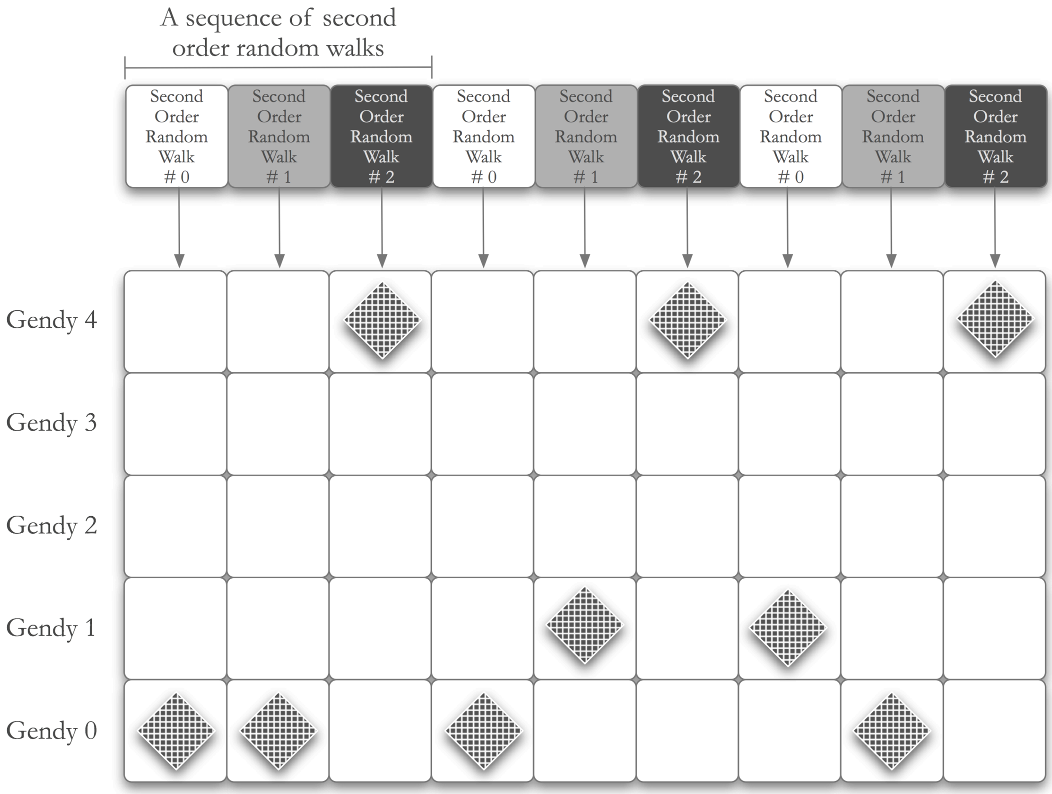

Method:

a) Select the number of second order random walks in the sequence (they will have the same parameter values and will evolve independently).

b) Specify the position of the two pairs of elastic barriers for the second order random walks.

c) For the continual generation of steps for the primary random walks: select a probability distribution and the ± number that will be the minimum and maximum size for these steps.

In this procedure, because of the behaviour of second order random walks, the first and the last GENDYs of a set are more prominent and, in most cases, the balance between the two is constantly changing. The resulting sound is discontinuous and unstable.

If a clearer differentiation between two GENDYs is desired: sort the GENDYs in the set from longest (lowest) to shortest (highest) before each selection.

Conclusion

The stochastic concatenation of dynamic stochastic synthesis is a technique that preserves the timbral quality and liveliness of Xenakis' dynamic stochastic synthesis algorithm. By adding a higher level of organization, the individual characteristics of a group of GENDYs are incorporated into textures and composite sounds.

This sound synthesis technique is *economical* and *automated*, to use Xenakis' own terms; when creating sounds or small sections, parameter values are only needed at the beginning "in order to give the initial impulse and a few premises" (Xenakis 1992, p. 295).

While most combinations of parameters yield continuous sounds, others can create sounds that exhibit interesting behaviours over time and give the illusion of being controlled by more complex mechanisms *(Tendency Masks* and *Sequences of Second Order Random Walks* are the selection procedures that more frequently display this type of behaviours).

Generally, it is difficult to predict precisely what the result of combining a set of GENDYs through a stochastic selection procedure will be. A knowledge of the more prominent characteristics of the technique can be acquired through practice but combinations will still have to be experienced directly and tuned by hand; also, I think that because of the big amount of linked stochastic variables that are used in this technique, unexpected results, good and bad, will always appear.

I am currently searching for selection procedures capable of generating sounds with intriguing behaviours over time, as the ones produced by the selection procedure *Sequences of Second Order Random Walks*. With this kind of behaviours, the timbral characteristics of the continuous sonorities, produced by most of the other selections procedures, can be molded into lively gestures.

References

Barthel-Calvet, A.S. 2001. “Chronologie.” *Portrait(s) de Iannis Xenakis.* Paris: Bibliothèque national de France. 25-80.

Berg, P. 1979. “Background and Foreground.” *SSP. A Bi-parametric Approach to Sound Synthesis.* Sonological Reports. Institute of Sonology, Utrecht. 9-32.

Blum, T. 1979. “Review of ‘Herbert Brün: Project SAWDUST’”. *Computer Music Journal* **3**(1): 6-7

Brün, H. and A. Chandra. 2001. "A Manual for SAWDUST." http://grace.evergreen.edu/~arunc/brun/sawdust

Di Scipio, A. 1998. “Compositional Models in Xenakis’s Electroacoustic Music.” *Perspectives of New Music* 36, no. 2: 201-243

Harley, J. 2002. “The Electroacoustic Music of Iannis Xenakis.” *Computer Music Journal* **26**(1): 33-57

Harley, J. 2004. *Xenakis: his life in music.* London: Taylor & Francis Books

Hoffmann, P. 2000. “The New GENDYN Program.” *Computer Music Journal* **24**(2): 31-38

Hoffmann, P. 2004. “‘Something rich and strange’: Exploring the Pitch Structure of GENDY3.” *Journal of New Music Research* **33**(2): 137-144

Holtzman, S.R. 1978. “A Description of an Automatic Digital Sound Synthesis Instrument.” *D.A.I. Research Report No. 59.* Edinburgh: Department of Artificial Intelligence

Koenig, G.M. 1959. “Studium im Studio.” *die Reihe 5.* English edition 1961. Bryn Mawr: Theodore Presser Company. 30-39.

Koenig, G.M. 1971a. “The Use of Computer Programmes in Creating Music.” *La Revue Musicale,* “Music and Technology”. 93-115.

Koenig, G.M. 1971b. *Summary: Observations on Compositional Theory.* Utrecht: Institute of Sonology

Roads, C. 1978. “Interview with Gottfried Michael Koenig.” *Computer Music Journal* **2**(2): 61-62

Robindoré, B. 1996. “Eskhaté Ereuna: Extending the Limits of Musical Thought-Comments On and By Iannis Xenakis.” *Computer Music Journal* **20**(4): 11-16

Serra, M.H. 1993. “Stochastic Composition and Stochastic Timbre: GENDY3 by Iannis Xenakis.” *Perspectives of New Music* 31, no. 1: 236-257

Solomos, M. 2001. “The Unity of Xenakis’s Instrumental and Electroacoustic Music: The Case for ‘Brownian Movements’”. *Perspectives of New Music* 39, no. 1: 244-254

Varga, Bálint András. 1996. *Conversations with Iannis Xenakis.* London: Faber and Faber.

Weisstein, E.W. 2006. "Bernoulli Distribution." *MathWorld--A Wolfram Web Resource.* http://mathworld.wolfram.com/BernoulliDistribution.html

Xenakis, I. 1955. "La crise de la musique sérielle." *Gravesaner Blätter* 1.

Xenakis, I. 1976. Foreword. *N’Shima* [score]. Paris: Editions Salabert.

Xenakis, I. 1992. *Formalized Music,* rev. ed. Stuyvesant, N.Y: Pendragon Press.