Xenakis' *Psappha* (this is a DRAFT)

Víctor Adán
vga2102@columbia.edu

Working draft: v.2010.8.25

Contents

1 Preamble

1.1 Motivation

My desire to carefully analyze *Psappha* stemmed from my interest in *discrete-time 1-bit music*. In this music there are only two "words" in the musical alphabet. i.e., at any given *time-point* we can find one of two events: '0' or '1'. Thus, the set of all possible events available at any given time-point can be represented with just one bit. In addition, two time-points (and thus two events) can not be arbitrarily close to each other. i.e., there is a minimum time interval between events. This smallest time interval is called the *tatum* (time atom) and all time intervals are integer multiples of it.

The musical thinking behind *Psappha* is very close to this world. The elements on the page (here called *strokes*) have no duration. At any given time-point, there either is a stroke or not. Strokes are sometimes accented or doubly accented, which strictly speaking puts *Psappha* outside of the world of 1-bit music. In this analysis however, I ignore accents for the most part.

1.2 Rambling on Analysis and Modeling

How might one analyze this music? Within the Nattiez/Molino framework,[16] there are three general approaches to his question: the poi-

etic, the neutral and the esthesic. In the first we ask: How is it made? What are the methods and tools employed by the composer to produce the piece? Flint's analysis[11] is mostly of this kind. She gives a pretty complete account of Xenakis' tools and his use of them in constructing *Psappha*. In particular, she discusses Xenakis' use of sieve and group theories to generate two different sets of materials. In the esthesic level we ask: What do I hear? Here we can forget about the composer and the physical data *per se* (the score or the sound), and focus on the perceptions impressed on us by the stimulus. Then there's the neutral level, the most common ground for musical analyzes. This is an analysis of the data itself, be it a printed score or a sonic rendition. While this tripartition may help conceptualize the problem space, it is a bit of a simplification. There are a variety of factors (socio-cultural, historical) that may affect and inform an analysis at any of these three levels. Further, these levels may inform each other.

My main interest here is to reverse engineering *Psappha*, so to speak; not to recover the methods used by Xenakis to construct it, nor to discuss the impressions that the piece imprints on me, but to reconstruct a generative model, or models, from the data found in the score. i.e., to recover the "code" and the syntactic "rules"(what Rowet calls the "synthetic model"[17]) or tendencies that will simultaneously give us a compact representation of the data and a "machine" that will generate sequences with similar properties.

While this is primarily an objective data driven structural analysis (objective in the sense that the data being analyzed is mainly the score, shareable and equally quantifiable by all), it is not unreasonable to ask, however, if the model derived from the data has not been overestimated (overfitted) from the point of view of perception. i.e., if the model encoding the piece is not more complicated than is necessary or more rigid than is desired. In his 1954 article *The Crisis of Serial Music*[23], Xenakis actually makes an implicit invocation of Occam's razor when he criticizes the weak relationship that he perceived to exist between the "complexity" of the serial composition system and the simplicity of the sonic result perceived:

> Linear polyphony is self-destructive in its current complexity. In reality, what one hears is a bunch of notes in various registers. The enormous complexity prevents one from following the tangled lines, and its macroscopic effect is one of unreasonable and gratuitous dispersion of sounds over the whole sound spectrum. Consequently, there is a contradiction between the linear polyphonic system and the audible result, which is a surface, a mass.
>
> This inherent contradiction with polyphony will disappear only once sounds become totally independent. In fact, since these linear combinations and their polyphonic superpositions are no

> longer workable, what will count will be the statistical average of isolated states of the components' transformations at any given moment.

A similar requirement can be set on a model or "synthetic analysis" of a piece of music. What is the simplest structural model that can still render perceptually accurate pieces?

However, if perception is to be factored in a musical analysis/model, what kinds or elements of perception should one consider? One can argue that there are at least as many ways of listening to music as there are listeners. One could also argue, however, that there are families of listening modalities. Expectation based listening, for example, is one of them. Sometimes, expectation plays an important role in the creation and reception of music. Expectation, the prospect of the future given our knowledge of the past, depends on what we know (or think we know) and our prior assumptions about the future. i.e., expectation depends on memory. But the future is uncertain. Most of the time there is more than one thing that can happen and we don't know for sure what will happen. At best we have an accurate idea of the likelihood of each of the possible futures. How much memory do we need for our expectation to be as accurate as possible? How might we introduce expectation into the analysis of a piece?

1.3 Tools and Method

In this analysis I focus on the statistical properties of the data found in the score. I derive localized probabilistic models from this data in the form of ϵ-machines (see Appendix F). Sieves (see Appendix D) are briefly discussed to present Xenakis' composition methods, and as point of reference and comparison for the stochastic ϵ-machines.

Both tools are markedly different in a couple of interesting ways. ϵ-machines are probabilistic, while sieves are strictly deterministic. Further, ϵ-machines are time domain models while sieves define a "period domain".[1]

1.4 Process, Sequence, Perceptual Stream

In general terms, we can say that a *process* is a well defined set of rules (or space) capable of generating certain kinds of sequences. Here we will deal more specifically with stochastic processes. We succinctly define a stochastic process as a sequence or random variables $\overline{S} = \ldots, S_{i-1}, S_i, S_{i+1}, \ldots$ associated with a probability measure $\mathrm{P}(\overline{S})$.[2]

A *sequence* is an ordered collection of things. In *Psappha*, specifically, a sequence is an ordered collection of consecutive *strokes* or points in the

[1] Sieves can be thought of as boolean analogues or Fourier analysis/synthesis (see Appendix D).

[2] See Appendix E.

score. The notation $\mathbf{X}[a : b]$ refers to the sequence of stroke *configurations*, or *multi-sequence*, in layer $\mathbf{X}$ spanning time points a and b, inclusive. A subindex n indicates a *sublayer*, thus $\mathbf{X}_n[a : b]$ refers to the sequence of strokes in sublayer n of layer $\mathbf{X}$.

A *perceptual stream* is a single cohesive percept that unfolds over time. It is well knows that a single sequence of sensory stimuli can give rise to more than one perceptual stream.[2] This depends on the properties of the stimulus and our own faculties of perception (awareness and attention). *Stream segregation* refers to our ability to demultiplex or separate a sequence of stimuli into multiple percepts (or perceptual streams).

1.5 General remarks about the score

1.5.1 Sections

For the purpose of this analysis we have divided the piece into the following eight sections:

Section 1: [0 : 518]	Section 5: [1411 : 1589]
Section 2: [519 : 739]	Section 6: [1590 : 1720]
Section 3: [740 : 989]	Section 7: [1720 : 2175]
Section 4: [990 : 1410]	Section 8: [2176 : 2396]

These sections are clearly delimited by the entrance and/or exit of one or several processes that contrast with those that precede or follow. There are, however, processes that run across section boundaries and some processes that span several sections.

1.5.2 Layers and Timbres

The score is composed of a total of six *layers* labeled A through F. All layers, except E, have up to three sublayers numbered 1, 2 and 3; E has only one. Each layer corresponds to a timbre category and a general register. Layers A, B and C are woods and/or skins, layers D, E and F a variety of metals. Each sublayer corresponds to a finer gradation of the register in each layer. The performer is free to choose the sounds and relative registers he/she considers appropriate.

1.5.3 Accents

Accent marks in *Psappha* are also open to interpretation. Xenakis lists five alternative ways of interpreting the accents:[3]

[3]The notes in the score are originally in french: 1. *intensité plus fôrte.*, 2. *changement brusque de timbre.*, 3. *changement brusque de poids.*, 4. *ajout brusque d'un autre son et le jouer simultanement avec celui du temps non accentué.*, 5. *combinaison simultanée des signification précédentes.*

1. Greater intensity.
2. Sudden change in timbre.
3. Sudden change in weight. (?)
4. Sudden addition of another sound to be played simultaneously with that of the non-accented time.(?)
5. Simultaneous combinations of the preceding interpretations.

1.5.4 Tempi

There are six tempo indications throughout the score: $\geqslant$ 152MM at tick 0, $\geqslant$ 272MM at tick 740 (opening Section 3), $\geqslant$ 110MM at tick 990 (opening Section 4), $\geqslant$ 134MM at tick 1720 (opening Section 7) and $\geqslant$ 152MM at tick 2175 (opening Section 8), plus an *accelérér* indication at tick 1610 going from $\geqslant$ 110MM to $\geqslant$ 134MM. The tempo markings indicate the minimum speed at which the music should be played, suggesting that any speed higher than this might be allowed, limited only by the performers capacity. However, the acceleration indication at tick 1610 suggests also that the relative speeds be preserved.

1.5.5 Hermeneutics

The openness of the tempo indications combined with the openness given to the choice of sounds and execution of the accents inevitably brings up the question of how to choose among the multiple combinations. In the explanation notes, however, Xenakis states quite clearly what the general guiding principle for these choices should be:[4]

> Indicated are only the desired global sonorities that dress the rhythmic structures and architectures of this piece. It is these [the structures and architectures] that must be enhanced by the balance between dynamics and timbres chosen out of ordinary sonorities.

With this in mind, it is interesting to compare various performances of the piece and attempt to deduce the interpretation of each performer, or to see how one's own interpretation might fit with different performances.

During this analysis I listened to three recordings of *Psappha* made by three different performers: Pedro Carneiro[3], Roland Auzet[1] and Steven Schick[18]. Figure 1 shows the waveforms of these recordings with sections marked by shaded blocks. Immediately visible is the different choices of tempo. Steven Schicks' performance is generally slower than that of Carneiro and Auzet. What motivated each performer's choice of tempi?

[4] "*Sont seulement indiquée les sonorités globales souhaitées, qui vêtiront les structures et architectures rhythmiques de cette piéce. Ce sont elles qui doivent être mises en valeur par des équilibres des puissances et des timbres choisis hors des sonorités banales.*"

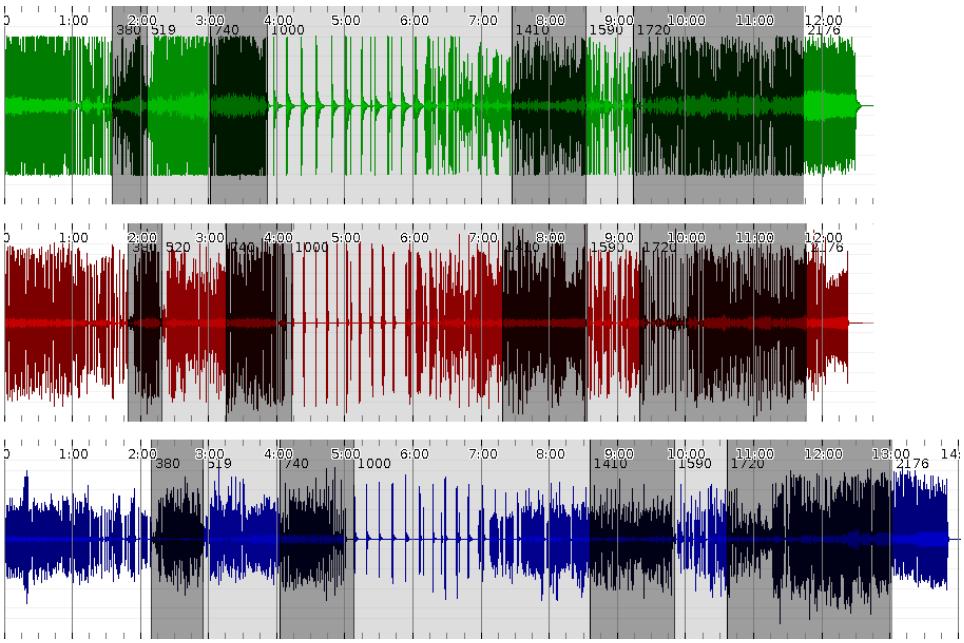

Figure 1: Waveforms of three performances of *Psappha*. The top one is from Pedro Carneiro, the middle one from Roland Auzet and the bottom one from Steven Schick. Sections are demarcated by alternating block of light and dark shades. Section 1 is divided in two subsections at tick 380 to show the beginning of the important **A**[380 : 518] sequence.

There is much that can be said about the hermeneutics of *Psappha* to fit the pages of this structural analysis. Here we would just like to pose the question: what role does *stream segregation* play in the piece and how does it influence the performer's choice of tempi and timbres? In *Psappha*, which combinations of timbres and tempi maximizes the possibility of perceiving multiple streams from a single sequential ordering of strokes? The question is left open.

2 Analysis and Modeling

2.1 The Conception of Time in *Psappha*

Musical time is usually conceptualized in two complementary ways: in one, relatively large time intervals are divided into progressively smaller parts; in the other, small indivisible time units, called *tatums* (time-atoms), are added to create a variety of larger intervals. The first is referred to as *divisive* rhythmic construction, the second as *additive*.[19]

In *Psappha* Xenakis' rhythmic exploration falls squarely in the world of additive rhythms. This is quite clear from the moment we first see the

score. The piece is entirely written on a grid with no standard symbolic rhythmic notation. Except for a very conspicuous quintuplet in $\mathbf{B}[70:74]$ (the only tuplet in the whole score), every stroke falls on a time point that is a multiple of the tatum—equivalently, every ISI (Inter Stroke Interval) is a multiple of the tatum. What is the tatum in *Psappha*? In tick 2023, Xenakis gives the indication: "two or three strokes per point" (*2 ou 3 coups par point*) for the strokes with *tremoli*. In this analysis I interpret these as two strokes per point. Because these tremoli strokes are half a measure apart (the smallest distance between points in the score), the tatum is then 1/4 of a measure.

2.2 Section 1 $[0:518]$

2.2.1 $\mathbf{B}[0:39]$

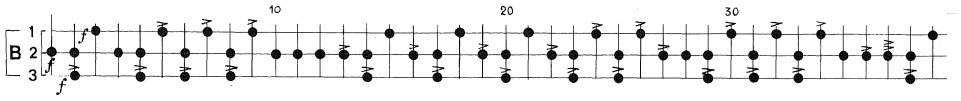

Sieve Decomposition According to Flint[11], Xenakis constructed the sublayer sequence $\mathbf{B_2}[0:39]$ by means of his theory of sieves (see Appendix D). The expression found in Xenakis' sketches is

$$[(8_0 \cup 8_1 \cup 8_7) \cap (5_1 \cup 5_3)] \cup [(8_0 \cup 8_1 \cup 8_2) \cap 5_0] \cup [8_3 \cap (5_0 \cup 5_1 \cup 5_2 \cup 5_3 \cup 5_4)] \cup [8_4 \cap (5_0 \cup 5_1 \cup 5_2 \cup 5_3 \cup 5_4)] \cup [(8_5 \cup 8_6) \cap (5_2 \cup 5_3 \cup 5_4)] \cup (8_1 \cap 5_2) \cup (8_6 \cap 5_1)$$

The expression indeed yields the sequence [0, 1, 3, 4, 6, 8, 10, 11, 12, 13, 14, 16, 17, 19, 20, 22, 23, 25, 27, 28, 29, 31, 33, 35, 36, 37, 38], which corresponds to the time-points in the score occupied by strokes. Taking the measure as an assumed constant delta-time, the corresponding boolean train of the sequence is [1, 1, 0, 1, 1, 0, 1, 0, 1, 0, 1, 1, 1, 1, 1, 0, 1, 1, 0, 1, 1, 0, 1, 1, 0, 1, 0, 1, 1, 1, 0, 1, 0, 1, 0, 1, 1, 1, 1, 0].

The two $(5_0 \cup 5_1 \cup 5_2 \cup 5_3 \cup 5_4)$ sub-expressions defining the sieve are intriguing as they are equivalent to the simpler 1_0, and because the congruent bases of 0 (mod 1) is the set $\mathbb{Z}$ of all integers, it follows that $8_3 \cap 1_0 = 8_3$. Thus, the sieve can more compactly be expressed as

$$[(8_0 \cup 8_1 \cup 8_7) \cap (5_1 \cup 5_3)] \cup [(8_0 \cup 8_1 \cup 8_2) \cap 5_0] \cup [(8_5 \cup 8_6) \cap (5_2 \cup 5_3 \cup 5_4)] \cup 8_3 \cup 8_4 \cup (8_1 \cap 5_2) \cup (8_6 \cap 5_1)$$

Is it possible to recover the original expressions defined by Xenakis to construct the stroke sequences, from the sequences themselves? The answer is no, not in the general case, because there can be multiple logical expressions that yield the same sequence. For example, the sequence

$[\ldots, -3, -2, 0, 2, 3, 4, 6, \ldots]$ can be generated from $rc(2,0) \cup rc(3,0)$ and from $rc(2,0) \oplus rc(6,3)$. We can, however, derive an equivalent expression; one that is different from Xenakis' original expression but that will nonetheless yield the same number sequence. The application of an OR decomposition[5] on $\mathbf{B_2}[0:39]$ yields the following residue classes: 8_3, 8_4, 10_3, 20_0, 20_8, 20_{11}, 20_{16}, 20_{17}, 40_1, 40_6, 40_{10}, 40_{14}, 40_{22}, 40_{25}, 40_{29}, 40_{38}. Thus the following expression also generates $\mathbf{B_2}[0:39]$:

$$8_3 \cup 8_4 \cup 10_3 \cup 20_0 \cup 20_8 \cup 20_{11} \cup 20_{16} \cup 20_{17} \cup 40_1 \cup 40_6 \cup 40_{10} \cup 40_{14} \cup 40_{22} \cup 40_{25} \cup 40_{29} \cup 40_{38}$$

As expected, the moduli of the residue classes resulting from this OR decomposition are multiples or 8 and 5.

Continuing in this fashion, we can OR-decompose the remaining two sublayer sequences to obtain sieve expressions that will characterize them in terms of simple periodic basis. However, looking at $\mathbf{B}[0:39]$ for a moment we see a simple pattern. $\mathbf{B_1}[0:39]$ is the complement of $\mathbf{B_2}[0:39]$, so we can define it simply as such:

$$\mathbf{B_1}[0:39] = \mathbf{B_2}[0:39]^c$$

Flint[11] characterizes $\mathbf{B_3}[0:39]$ as a sieve constructed from elementary residue classes with periods 8 and 5 (i.e., $rc(8,n)$ and $rc(5,m)$) much like Xenakis' own definition of $\mathbf{B_2}[0:39]$. $\mathbf{B_3}[0:39]$ can, however, be directly derived from $\mathbf{B_1}[0:39]$ as its displacement by one:[6]

$$\mathbf{B_3}[0:39] = \mathbf{B_1}[0+1:39+1]$$

The fact that we have been able to, very simply and compactly, define the sublayer sequences in $\mathbf{B}[0:39]$ in terms of each other tells us something about their affinity. Actually, we can recover any two sublayer sequences from any one given sublayer sequence.[7] This tightness of interaction (or coupling) between sequences is an important factor to consider when modeling them. In the "sieve domain", we are better off (have a simpler model) defining sieves in terms of each other, as we have done above, than we are defining each sequence anew with it's own logical expression. A similar principle applies when we model these sequences in the time domain.

[5] See Appendix D for an explanation of the OR decomposition.

[6] Note that the attack at tick 40 in sublayer B_1 is not part of the complement of either $\mathbf{B_2}[0:39]$ or $\mathbf{B_2}[40:79]$, but, according to Flint[11], a stroke added by Xenakis to demarcate a change of sieve.

[7] While this is an interesting property of $\mathbf{B}[0:39]$ as a whole, it does not follow that the sequence itself or the sublayer sequences are necesarily "interesting". Any number of sequences can be constructed with these constraints. e.g. $\mathbf{B_2}[m:n] = noise$, $\mathbf{B_1}[m:n] = \mathbf{B_2}[m:n]^c$, $\mathbf{B_3}[m:n] = \mathbf{B_1}[m+1:n+1]$.

Stochastic Model We have studied the structure $\mathbf{B}[0:39]$ in the period domain[19] using residue classes a la Xenakis. How should we model it in the time domain? What is the temporal behavior of the $\mathbf{B}[0:39]$? We could study and model each sublayer sequence $\mathbf{B_{1,2,3}}[0:39]$ independently and hope that both the individual and collective behaviors of the sequences returned by the model are "satisfactory", i.e., that they present the same statistical characteristics as the originals found in the score. Given the clear coupling that exists between the sublayer sequences, however, it makes more sense to consider $\mathbf{B}[0:39]$ as a whole. We can think of $\mathbf{B}[0:39]$ as a three-dimensional sequence (a multi-sequence) in which each dimension can have one of two values (0 or 1) giving a total of eight (2^3) possible configurations per time-point:

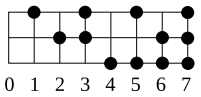

Observe, however, that only configurations $\begin{smallmatrix}\bullet\\ \cdot\\ \cdot\end{smallmatrix}$, $\begin{smallmatrix}\cdot\\ \bullet\\ \cdot\end{smallmatrix}$ and $\begin{smallmatrix}\cdot\\ \bullet\\ \bullet\end{smallmatrix}$ (numbered 1, 2 and 6) appear in $\mathbf{B}[0:39]$. We call each of these a *letter*, and the set of all letters found in the sequence, the alphabet: $\mathcal{A} = \{\begin{smallmatrix}\bullet\\ \cdot\\ \cdot\end{smallmatrix}, \begin{smallmatrix}\cdot\\ \bullet\\ \cdot\end{smallmatrix}, \begin{smallmatrix}\cdot\\ \bullet\\ \bullet\end{smallmatrix}\}$. Thus, the cardinality of the alphabet of $\mathbf{B}[0:39]$ is three. Now we look at the temporal dependencies between consecutive letters. Because there are three letter in the alphabet, there are three possible futures at any given time-point. What is the likelihood of each possible future? Notice, for example, that letter $\begin{smallmatrix}\bullet\\ \cdot\\ \cdot\end{smallmatrix}$ always follows letter $\begin{smallmatrix}\cdot\\ \bullet\\ \bullet\end{smallmatrix}$: $\mathrm{P}(\begin{smallmatrix}\bullet\\ \cdot\\ \cdot\end{smallmatrix}|\begin{smallmatrix}\cdot\\ \bullet\\ \bullet\end{smallmatrix}) = 1$.[8] This means that having heard $\begin{smallmatrix}\cdot\\ \bullet\\ \bullet\end{smallmatrix}$ we can be certain that we will hear $\begin{smallmatrix}\bullet\\ \cdot\\ \cdot\end{smallmatrix}$ immediately after. Thus we say that this process has memory. How much memory? If we retain the last two events instead of just the last one, can we improve our prediction? Table 1 shows the conditional probability of each letter given each of the three possible subsequences (words) of length 1 as found in $\mathbf{B}[0:39]$. Table 2 shows the conditional probabilities of each letter given the subsequences of length 2.[9] It is clear from the Tables that the conditional probabilities do

subsequence	*morph*		
$[\begin{smallmatrix}\bullet\\ \cdot\\ \cdot\end{smallmatrix}]$	$\mathrm{P}(\begin{smallmatrix}\bullet\\ \cdot\\ \cdot\end{smallmatrix}\|\begin{smallmatrix}\bullet\\ \cdot\\ \cdot\end{smallmatrix}) = 0$	$\mathrm{P}(\begin{smallmatrix}\cdot\\ \bullet\\ \cdot\end{smallmatrix}\|\begin{smallmatrix}\bullet\\ \cdot\\ \cdot\end{smallmatrix}) = \frac{7}{12} \approx 0.58$	$\mathrm{P}(\begin{smallmatrix}\cdot\\ \bullet\\ \bullet\end{smallmatrix}\|\begin{smallmatrix}\bullet\\ \cdot\\ \cdot\end{smallmatrix}) = \frac{5}{12} \approx 0.42$
$[\begin{smallmatrix}\cdot\\ \bullet\\ \cdot\end{smallmatrix}]$	$\mathrm{P}(\begin{smallmatrix}\bullet\\ \cdot\\ \cdot\end{smallmatrix}\|\begin{smallmatrix}\cdot\\ \bullet\\ \cdot\end{smallmatrix}) = 0$	$\mathrm{P}(\begin{smallmatrix}\cdot\\ \bullet\\ \cdot\end{smallmatrix}\|\begin{smallmatrix}\cdot\\ \bullet\\ \cdot\end{smallmatrix}) = \frac{6}{14} \approx 0.43$	$\mathrm{P}(\begin{smallmatrix}\cdot\\ \bullet\\ \bullet\end{smallmatrix}\|\begin{smallmatrix}\cdot\\ \bullet\\ \cdot\end{smallmatrix}) = \frac{8}{14} \approx 0.57$
$[\begin{smallmatrix}\cdot\\ \bullet\\ \bullet\end{smallmatrix}]$	$\mathrm{P}(\begin{smallmatrix}\bullet\\ \cdot\\ \cdot\end{smallmatrix}\|\begin{smallmatrix}\cdot\\ \bullet\\ \bullet\end{smallmatrix}) = 1$	$\mathrm{P}(\begin{smallmatrix}\cdot\\ \bullet\\ \cdot\end{smallmatrix}\|\begin{smallmatrix}\cdot\\ \bullet\\ \bullet\end{smallmatrix}) = 0$	$\mathrm{P}(\begin{smallmatrix}\cdot\\ \bullet\\ \bullet\end{smallmatrix}\|\begin{smallmatrix}\cdot\\ \bullet\\ \bullet\end{smallmatrix}) = 0$

Table 1: Conditional probabilities over pasts of length $L = 1$ obtained from $\mathbf{B}[0:39]$.

[8] This is obvious from the definitions of $\mathbf{B_3}[0:39]$ and $\mathbf{B_1}[0:39]$ given above.

[9] In this analysis all the probabilities are estimated simply by counting the relative frequencies of the words found in the sequences. See Appendix E for more details.

subsequence	*morph*		
[⠂⠄]	$P(⠂\|⠂⠄) = 0$	$P(⠄\|⠂⠄) = \frac{3}{7} \approx 0.43$	$P(⠆\|⠂⠄) = \frac{4}{7} \approx 0.57$
[⠄⠄]	$P(⠂\|⠄⠄) = 0$	$P(⠄\|⠄⠄) = \frac{3}{6} = 0.5$	$P(⠆\|⠄⠄) = \frac{3}{6} = 0.5$
[⠆⠂]	$P(⠂\|⠆⠂) = 0$	$P(⠄\|⠆⠂) = \frac{7}{12} \approx 0.58$	$P(⠆\|⠆⠂) = \frac{5}{12} \approx 0.42$
[⠂⠆]	$P(⠂\|⠂⠆) = 1$	$P(⠄\|⠂⠆) = 0$	$P(⠆\|⠂⠆) = 0$
[⠄⠆]	$P(⠂\|⠄⠆) = 1$	$P(⠄\|⠄⠆) = 0$	$P(⠆\|⠄⠆) = 0$

Table 2: Conditional probabilities over pasts of length $L = 2$ obtained from $\mathbf{B}[0:39]$.

not change as we go from words of length 1 to words of length 2. This means that remembering two past letters instead of just one will not improve our prediction of the future. What about retaining three? Will this improve our forecast? Instead of looking at more probability tables, we measure the entropy $H(L)$, entropy gain $\Delta H(L)$ and predictability gain $\Delta^2 H(L)$ of the sequence at growing word lengths L.[10] Looking at the entropy gain in Figure 2 it is clear that increasing our memory to three letters will not improve our estimates of the future. Notice how the *entropy gain* plateaus at lengths 2, 3 and 4, and that the predictability gain at lengths 3 and 4 is 0. This suggests that a stochastic model of $\mathbf{B}[0:39]$ should be quite simple, needing no more memory than the previous letter at any given time. You may wonder why I have chosen to ignore the predictability gain obtained at $L > 4$. The new predictability gained at $L = 5$ would seem to suggest that the process inded has more memory than I have suggested. This is, however, not necessarily the case. The reason is that the *statistical trust boundary* (L_{tb})[11] of $\mathbf{B}[0:39]$ is at $L = 3$, since $3^3 = 27 < length(\mathbf{B}[0:39]) = 40 < 3^4 = 81$. Thus, we cannot assume that any predictability gain beyond this point is a real attribute of the process. More likely it is an attribute of the faulty statistics derived from limited amounts of data.

We can now attempt to derive a stochastic model of the sequence in the form of an ϵ-machine.[12] The two important parameters in Shalizi's ϵ-machine reconstruction algorithm[20] are the maximum subsequence length, L_{max}, and the *significance level*, α. L_{max} will have to be smaller than L_{tb}. L_{max} sets a limit to the ammount of memory the derived model will have. The significance level (α) sets the tolerance in establishing an equivalence between different morphs. Notice that the morphs of the subsequences [⠂⠄], [⠄⠄] and [⠆⠂] in $\mathbf{B}[0:39]$ are quite close, with that of [⠄⠄] being exactly uniform and those of [⠂⠄] and [⠆⠂] fairly close to uniform. Should these three

[10] See Appendix E for an overview of information theory.
[11] See Appendix E.
[12] See Appendix F for an overview of ϵ-machines.

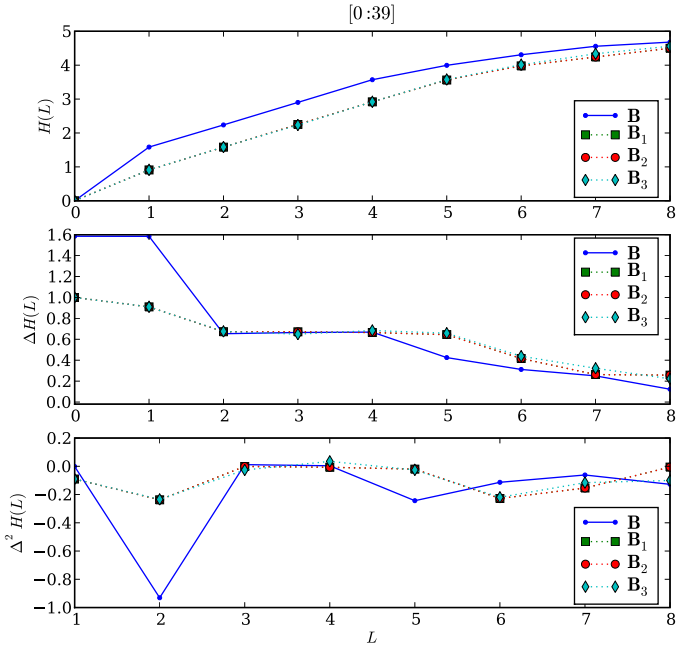

Figure 2: Entropy ($H(L)$), entropy gain ($\Delta H(L)$) and predictability gain ($\Delta^2 H(L)$) for sequences $\mathbf{B}[0:39]$ and $\mathbf{B_{1,2,3}}[0:39]$ at increasing subsequence lengths L.

subsequences be considered equivalent in terms of their predictive contours? What α should we use in our model estimation? The choice is, unfortunately, ultimately arbitrary, but the existence of α can be justified on a couple of grounds:

1. There's always a limited amount of data, which means that the derived probability measures will never be exact.
2. There can be noise in the sequence being studied. This is not applicable here, but is a factor in other kinds of data.
3. Perceptual discrimination. This is usually not a factor in other kinds of implementations, but here we can justitfy a particular choice of α on these grounds too. Does the resulting model sound more random than the original? If so, pump-up α, otherwise, keep or bring it down.

Here we use a significance level of 0.1 and a maximum length L_{max} of 3 in the ϵ-machine reconstruction of $\mathbf{B}[0:39]$. The result is shown in Figure 3). Notice that the graph has only two states, which means that the three close-to-uniform morphs discussed above have been considered equivalent and have been put together in the same causal state.

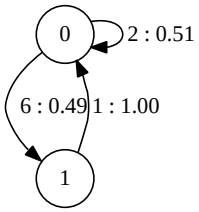

Figure 3: ϵ-machine reconstruction of $\mathbf{B}[0:39]$ obtained with parameters $L_{max} = \{2, 3, 4\}$ and $\alpha = 0.1$. Each numbered circle represents a *causal state*. Edges are labeled $l : p$, where l is the *letter* in the alphabet corresponding to that edge and p is the probability that that letter will follow given the current state. The letters, here represented as numbers, map to the stroke configurations as follows: $\vdots = 1$, $\dot{\bullet} = 2$, $\dot{\vdots} = 6$.

2.2.2 $\mathbf{B}[40:79]$

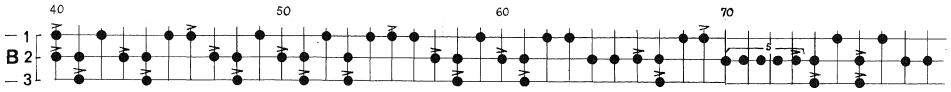

$\mathbf{B}[40:79]$ is quite similar to the opening $\mathbf{B}[0:39]$. The same tight relationship between sublayers that we saw in $\mathbf{B}[0:39]$ is found here. $\mathbf{B_1}[40:79]$ is again the complement of $\mathbf{B_2}[40:79]$. Here $\mathbf{B_3}[40:79]$ is not a simple displacement of $\mathbf{B_1}[40:79]$. However, $\mathbf{B_3}[40:79]$ can still be derived from both $\mathbf{B_2}[40:79]$ and $\mathbf{B_1}[40:79]$ as their intersection:

$$\mathbf{B_3}[40:79] = \mathbf{B_1}[40+1:79+1] \cap \mathbf{B_2}[40:79]$$

As can be inferred by the resemblance between $\mathbf{B}[0:39]$ and $\mathbf{B}[40:79]$, Xenakis used a very similar procedure to construct both sequences. I won't go over Xenakis' sieve expression for $\mathbf{B_2}[40:79]$ here again. However, it is worth pointing out that, according to Flint, the conspicuous quintuplet (the only tuplet in the whole piece) in $\mathbf{B_2}[70:74]$ is explained by Xenakis' desire to fit a sieve with 42 points per cycle (using residue classes with moduli 7 and 6) into the same forty measures occupied by $\mathbf{B}[0:39]$. The reader is referred to Flint[11] for the details.

Stochastic Model As before, we perform a preliminary statistical analysis of multi-sequence $\mathbf{B}[40:79]$ and its sublayer sequences $\mathbf{B_{1,2,3}}[40:79]$. The resulting statistics are quite similar to those of $\mathbf{B}[0:39]$ (see Figure 4). This is expected given the similarity between the two sequences. Notice, however, that in this case we do have a slight increase in predictability gain as we go from 2 to 3 and to 4 word lengths.

The ϵ-machine derived from $\mathbf{B}[40:79]$ is slightly more complex than the one obtained from $\mathbf{B}[0:39]$ (see Figure 5). This is due to the presence of consecutive $\overset{\bullet}{\vdots}$ configurations absent in $\mathbf{B}[0:39]$.

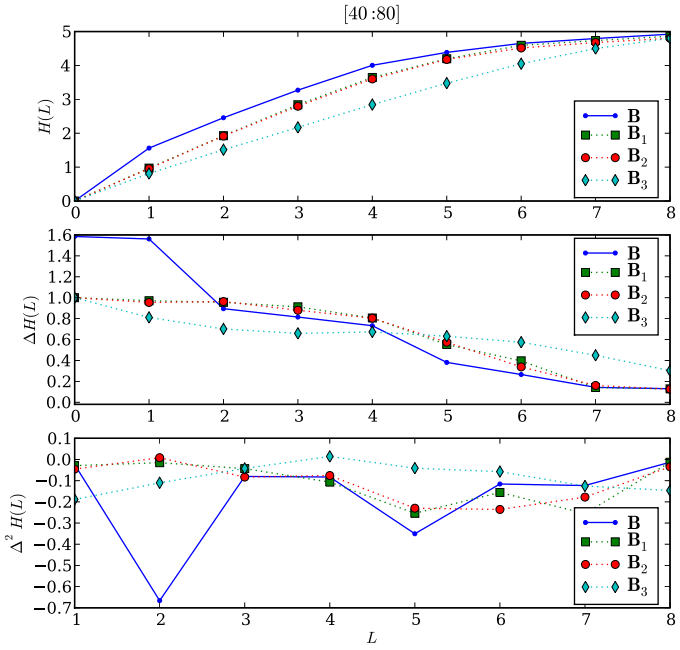

Figure 4: Entropy ($H(L)$), entropy gain ($\Delta H(L)$) and predictability gain ($\Delta^2 H(L)$) for sequences $\mathbf{B_{1,2,3}}[40:79]$ and $\mathbf{B}[40:79]$ at increasing subsequence lengths L.

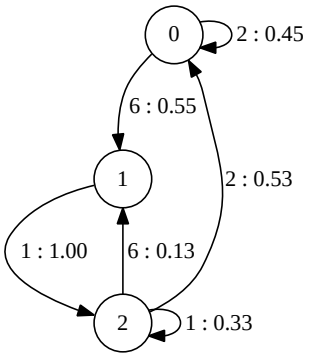

Figure 5: ϵ-machine reconstruction of $\mathbf{B}[40:79]$ obtained with parameters $L_{max} = \{2,3,4\}$ and $\alpha = 0.1$.

2.2.3 $\mathbf{B}[80:167]$

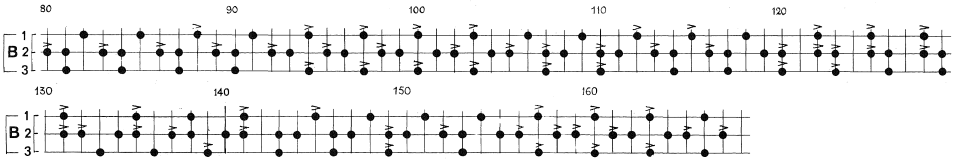

In $\mathbf{B}[80:167]$ we find a treatment of the sublayers that differs from that found in the first 80 measures. $\mathbf{B_2}[80:167]$ is a steady and unchanging alternation between 1 and 2 measure ISIs: $[\ldots,1,2,1,2,\ldots]$, and can be compactly represented as $rc(3,0) \cup rc(3,1)$. A period 3 cyclic pattern is also suggested in both $\mathbf{B_1}[80:167]$ and $\mathbf{B_3}[80:167]$, but every now and then the otherwise constant 3 measure ISIs become 4 measures long, one period at a time—except for the two consecutive 4 measure ISIs in $\mathbf{B_3}[149:157]$. i.e., these simple quasi-periodic sequences seem to be phase shifted (delayed, never anticipated) to alter the relationships between them and the steady $\mathbf{B_2}[80:167]$ that serves as a reference.[13]

The predominance of period 3 is visible also in the entropy measures (Figure 6) of all four sequences $\mathbf{B_{1,2,3}}[80:167]$ and $\mathbf{B}[80:167]$; notice how the entropy gain suddenly plateaus at $L=3$ for all four sequences. Because $\mathbf{B_2}[80:167]$ is the only perfectly periodic sequence, it is the only one that reaches a zero entropy gain.

A mechanical model suggests itself from these observations. Imagine three spinning wheels or gears of equal radii touching each other by the edge. One of the wheels ($\mathbf{B_2}[80:167]$) steadily rotates and drives the other two which, due to faulty contact with the driver, cannot keep up and slip every so often (see Figure 7). What is the level of contact/coupling between the wheels? i.e., What is the degree of interdependence between sublayer sequences? Taking the *normalized mutual information distance* $D(X,Y)$[14] between all three sublayer sequences (Figure 8) we see that they are almost completely independent of each other. i.e., the coupling between sublayers found in $\mathbf{B}[0:39]$ and $\mathbf{B}[40:79]$ is absent. See, for comparison, the normalized mutual information distances between sublayer sequences in $\mathbf{B_{1,2,3}}[0:39]$ and $\mathbf{B_{1,2,3}}[40:79]$ (Figure 9). From the point of view of our mechanical model, this suggests that the wheels have very little coupling, if any. More likely they are just driving themselves independently of each other, with wheels 1 and 3 ($\mathbf{B_1}[80:167]$ and $\mathbf{B_3}[80:167]$) simply stopping randomly every now and then for only one time-step (one measure). Given that the sublayer sequences are independent, it makes more sense to model

[13] $\mathbf{B_3}[80:167]$ is phase-shifted 7 times, at ticks 90, 103, 110, 129, 139, 149 and 153. $\mathbf{B_1}[80:167]$ is phase-shifted 3 times, at ticks 118, 131 and 141.

[14] See Appendix E

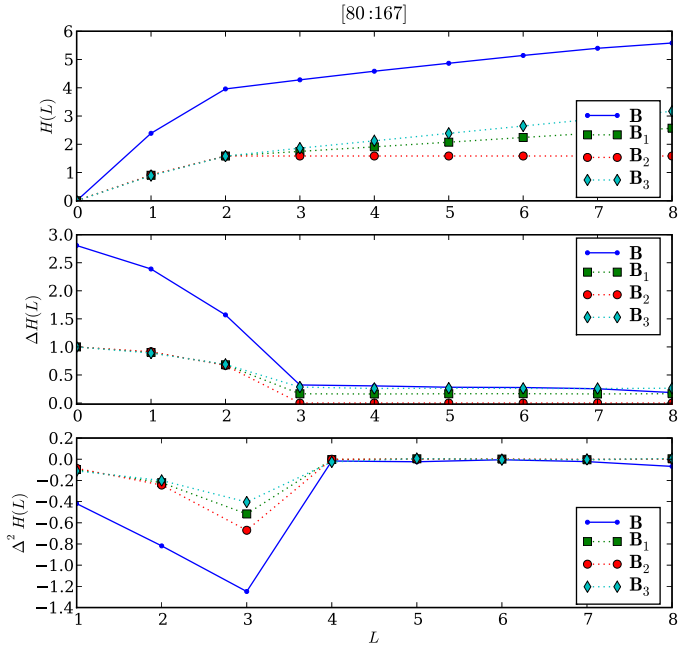

Figure 6: Entropy measures for $\mathbf{B}[80:167]$ and $\mathbf{B_{1,2,3}}[80:167]$ at increasing block lengths L. The entropy gain ($\Delta H(L)$) drops and plateaus at $L=3$ due to the presence of a strong period 3 structure in all sublayer sequences.

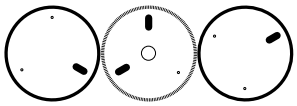

Figure 7: Three-gear model of sequence $\mathbf{B}[80:167]$.

each sequence independently than to take them together as a single multi-sequence as we did with $\mathbf{B}[0:39]$ and $\mathbf{B}[40:79]$. Just as storing more memory than necessary does not give us a better prediction of the future, modeling a collection of independent sequences together as one will not give us a better model i.e., it will not give us a model that will generate multi-dimensional sequences with more accurate statistics than those of thee parallel independent models.

Using sieves, $\mathbf{B_2}[80:167]$ can be compactly expressed as $rc(3,0) \cup rc(3,1)$, as we have already seen. As per $\mathbf{B_{1,3}}[80:167]$, the intuition is to model them as some form of periodic sequences with noise. Within the framework of sieve theory, however, it is not quite possible to represent these compactly.

We thus turn to our stochastic model once more. We compute a stochas-

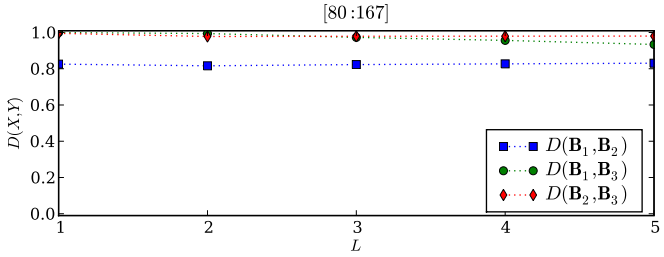

Figure 8: The normalized mutual information distance $D(X, Y)$ between sublayer sequences $\mathbf{B_{1,2,3}}[80 : 167]$. The distance between sublayers 2 and 3, and sublayers 1 and 3 is very close to the maximum distance of 1, showing a complete independence between them. Sublayers 1 and 2 are at a constant distance of slightly above 0.8, with no change at increasing L.

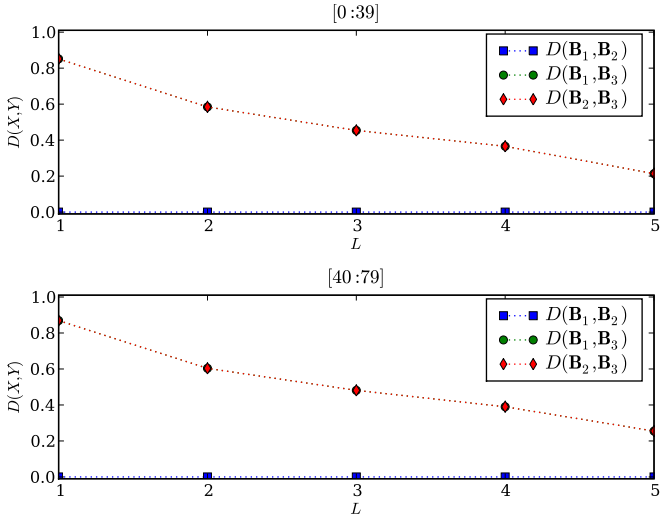

Figure 9: The normalized mutual information distance $D(X, Y)$ between sublayer sequences $\mathbf{B_{1,2,3}}[0 : 39]$ and between sublayer sequences $\mathbf{B_{1,2,3}}[40 : 79]$. For both, sublayers 1 and 2 share the exact same information and are thus redundant to each other. i.e., their distance is 0. The information distance between sublayers 1 and 3, and between 2 and 3 is the same, starting big for $L = 1$ and diminishing gradually as L increases.

tic model of each sublayer sequence independently. The three resulting ϵ-machines are shown in Figure 10.

It is interesting to note that, according to Flint, Xenakis composed $\mathbf{B}[80 : 167]$ "intuitively" i.e., without the formalization found in $\mathbf{B}[0 : 39]$ and $\mathbf{B}[40 : 79]$. This can explain the free play that we have observed between the sublayers. Xenakis may have been thinking precisely about phase shifts to try out the various configurations that result between sublayers as patterns

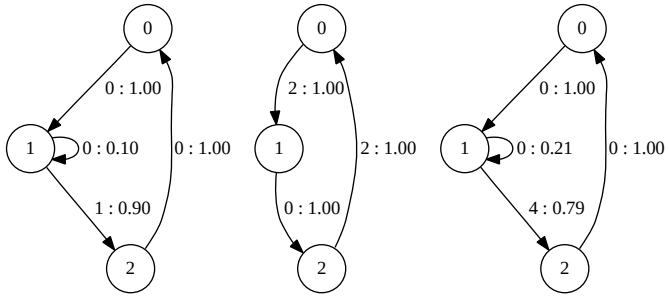

Figure 10: The three ϵ-machines obtained from the independent sublayer sequences $\mathbf{B_{1,2,3}}[80 : 167]$. Parameters used where $L_{max} = \{3, 4, 5\}$ and $\alpha = 0.1$.

are shifted. To what extent are the phase relationships between all three sublayer sequences explored? Since all three sublayer sequences $\mathbf{B_{1,2,3}}[80 : 167]$ have period three (ignoring the noise component), each can be phase-shifted to three indexes: 0, 1, and 2. Thus, both sequences $\mathbf{B_{1,3}}[80 : 167]$ can together occupy nine (3^2) different phase positions relative to $\mathbf{B_2}[80 : 167]$. The displacements found in $\mathbf{B}[80 : 167]$ result in six of the nine possible phase configurations between the three sublayers (see Table 3).

B_1 *phase*	B_3 *phase*	*measures*
0	0	
0	1	[122 : 130]
0	2	[131 : 133]
1	0	
1	1	
1	2	[134 : 139]
2	0	[107 : 112], [143 : 151]
2	1	[80 : 91], [113 : 118], [152 : 154]
2	2	[92 : 103], [155 : 166]

Table 3: Phase relationships between $\mathbf{B_1}[80 : 167]$ and $\mathbf{B_3}[80 : 167]$ relative to $\mathbf{B_2}[80 : 167]$.

2.2.4 $\mathbf{B}[200 : 519]$

The regularly pulsating "machines" found in $\mathbf{B}[0 : 200]$ seems to break down and loose power at tick 200. $\mathbf{B}[200 : 519]$ is a gradual scattering and slowing down of $\mathbf{B}[0 : 200]$-like stroke configurations. The temporal stretching (density thinning) of the layer is accompanied by a gradual general decent in pitch, as noted in the score at tick 204: "The B voices descend

progressively towards the low sound of C3, which they reach at 519."[15] Thus, **B**[200 : 519] bridges two otherwise unrelated processes and sound worlds: Section 1's opening timbre and tempo and Section 2's introduction of layer C and the slower tempi that follow.

2.2.5 Properties of $\mathfrak{B}$-like sequences

The three multi-sequence fragments covered so far (**B**[0 : 39], **B**[40 : 79] and **B**[80 : 167]) share two clear properties:

1. Both single strokes and simultaneous strokes per time-point are possible. Present in **B**[0 : 200] are all but one of the 8 possible three-sublayer stroke configurations: ⁝, ⁝, ⁝, ⁝, ⁝, ⁝ and ⁝, but not ⁝.
2. The two ISIs found are 1 and 2 measures long.

For future reference, I will refer to sequences that have these two properties as $\mathfrak{B}$-like sequences.

2.2.6 **A**[380 : 518] and other layer A sequences

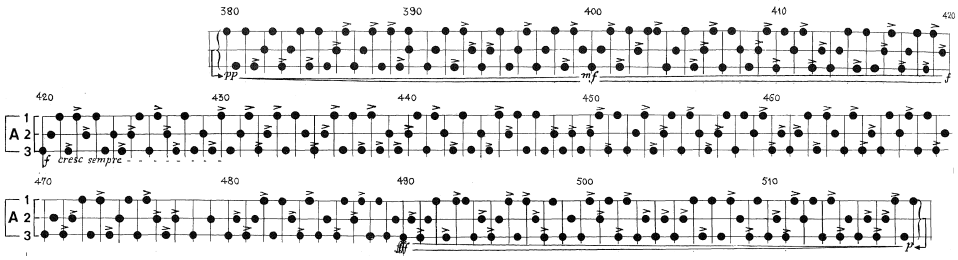

In contrast to layer B, layer A has a fragmentary nature. While layer B feels like a single long sequence with relatively smooth and gradual transformations, layer A is characterized by abrupt alternations between silence and sound. A fragmentary quality can also be observed in the stroke sequences themselves. Some of these sequences are segments of other layer A sequences that closely follow or precede them. e.g., **A**[275.5 : 296] is a fragment of **A**[380 : 518] (**A**[275.5 : 296] = **A**[451.5 : 473]), and **A**[327.5 : 336] is in turn the last 18 strokes of **A**[275.5 : 296], with a glitch: a missing stroke at time-point 331.5. From these observations, and from the fact that a fragment of **A**[380 : 518] also reappears as a kind of coda at the very end of the piece (see Section 2.5), it is clear that **A**[380 : 518] is a well defined and recurrent structure in *Psappha*. Thus, I will here focus on this sequence.

With a few exceptions, **A**[380 : 518] has the following general properties:

[15] "*Les voix B descendent progressivement vers le grave de C3 qu'elles atteignent à 519*".

1. At any given time-point there is at most one stroke. i.e., There are no simultaneities. Exception at tick 390.

2. Two consecutive strokes fall always on different sublayers.
   i.e., if $\mathbf{A}[n] = \overset{\bullet}{\cdot}$ then $\mathbf{A}[n+0.5] = \dot{\bullet}$ or $\dot{\bullet}$. Exceptions at [403 : 403.5] and [493 : 493.5].

3. There is an attack every half a measure. Exceptions at time-points 390.5, 477.5, 478.5 and 479.5.

All three properties make $\mathbf{A}[380 : 518]$ clearly distinct from the layer B sequences seen so far. For future reference, I will call those sequences that exhibit these properties $\mathfrak{A}_1$-like. These properties suggest that $\mathbf{A}[380 : 518]$ was constructed differently from the sequences in layer B covered so far, or at least with markedly different criteria. Indeed, Flint[11] tells us that Xenakis composed it by sequencing and dovetailing the six possible permutations of sequential strokes on the three sublayers, all formalized under a group theoretic framework. I will not go into Xenakis' methods again here. Instead, as suggested in the introduction, I will focus on the data and attempt the reconstruction of a stochastic model from it.

It is clear upon hearing and seeing the score that $\mathbf{A}[380 : 518]$ has a lot of internal repetition. What is the extent of this repetition? We find that the longest sequence that repeats is $\mathbf{A}[412 : 430]$ with a length of 37 strokes (18.5 measures), reappearing at $\mathbf{A}[448.5 : 466.5]$. This is clearly visible in the boolean auto-correlogram computed from the sequence (Figure 11).[16] The glitches or pulse discontinuities at $\mathbf{A}[390.5]$ and $\mathbf{A}[477 : 480]$ can also be clearly seen. As we increase the subsequence length L in the autocorrelation analysis, we begin to see a trend that roughly divides the sequence in two. At $L = 8$, we see that some subsequences that appear at the very beginning of $\mathbf{A}[380 : 518]$ reappear a few times, only up to about two thirds of the sequence, while some subsequences that appear for the first time after one third of the sequence repeat up to the very end.

Stochastic Model Property 1 mentioned above is evidence enough of sublayer interaction. It is clear then that $\mathbf{A}[380 : 518]$ is best treated as a single multi-sequence for our modeling purposes. The internal repetition in the multi-sequence suggests a process with more memory than that found in the layer B sequences. The predictability gain plot of $\mathbf{A}[380 : 518]$ shows a strong "connection" between consecutive pairs of strokes (Figure 12). This is also expected from property 2. After the relatively large negative peak at $L = 2$, the predictability gain drops markedly at $L = 3$, but stays relatively constant around -0.1 and does not reach 0. While the statistical trust

[16] For a given sequence $\mathbf{A}$, the auto-correlogram is computed by taking every subsequence of length L in $\mathbf{A}$ and comparing it to every other subsequence in $\mathbf{A}$. If the segments are identical we have a match, otherwise we don't.

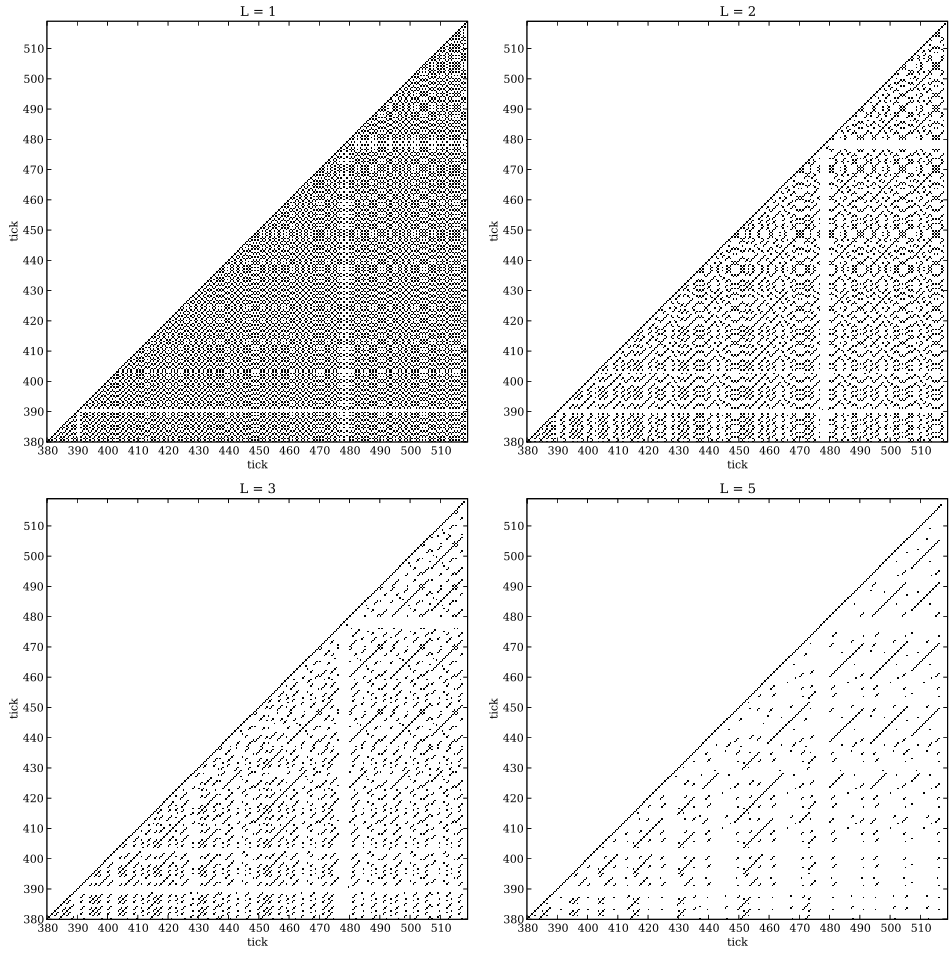

Figure 11: Boolean auto-correlogram of $\mathbf{A}[380:518]$ at increasing lengths L. Black dots indicate exact match (identical subsequences), white space indicates differing subsequences.

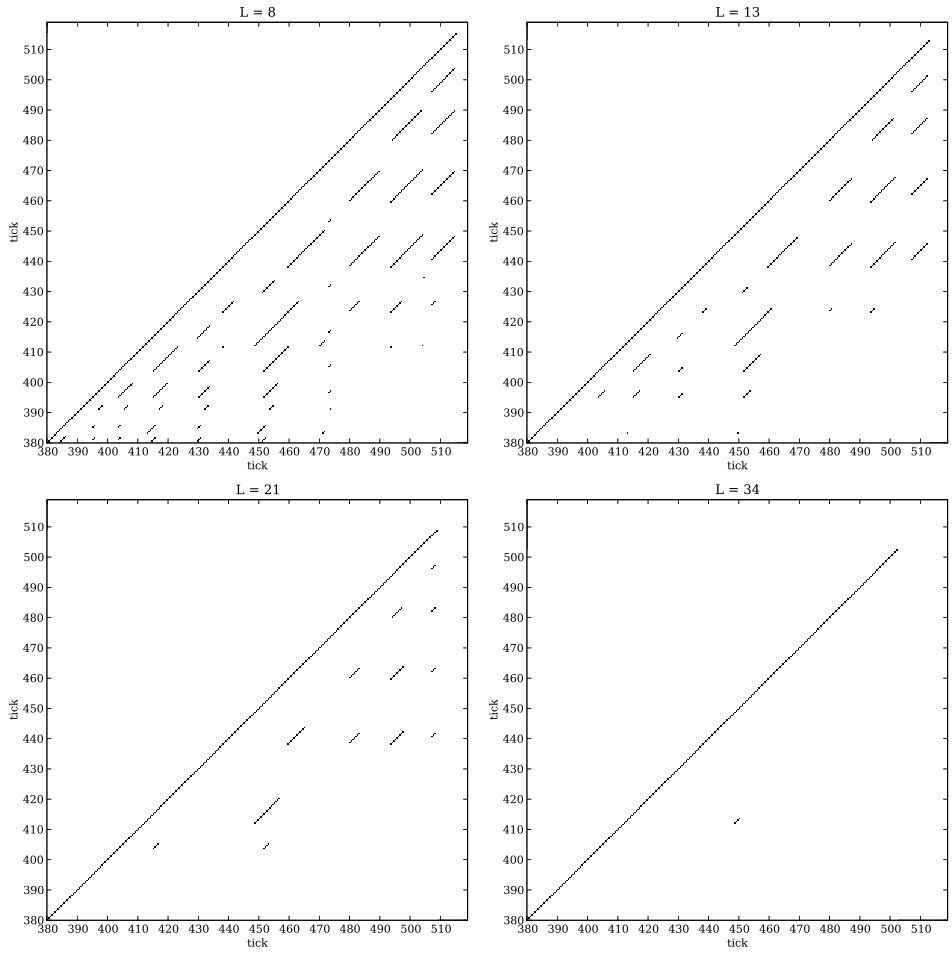

Figure 11: Continued...

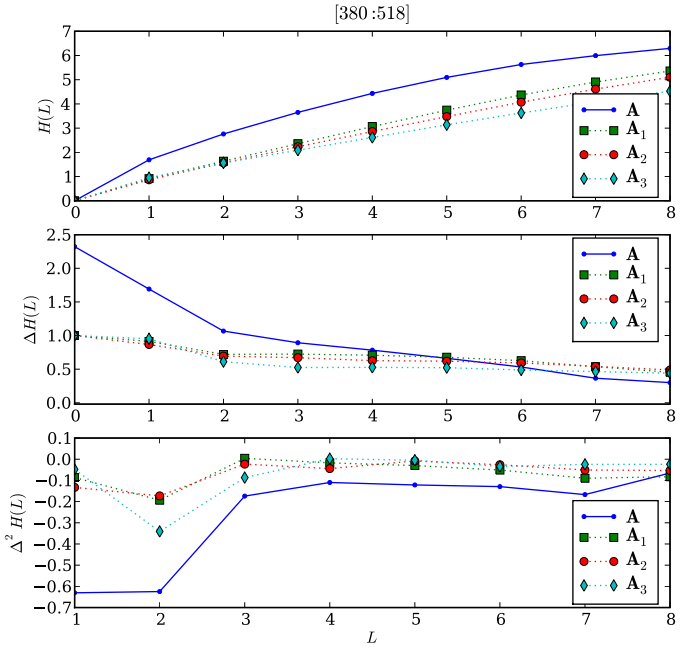

Figure 12: Entropy ($H(L)$), entropy gain ($\Delta H(L)$) and predictability gain ($\Delta^2 H(L)$) for sequences $\mathbf{A}[380:518]$ and its sublayer sequences at increasing block-lengths L.

boundary of $\mathbf{A}[380:518]$ is $L_{tb}=3$, the *soft statistical trust boundary* is 4. This is due to the statistically insignificant presence of stroke configurations ⁝ and ⋮. Thus, we cannot freely interpret the entropy measures for $L>4$.

As it appears in the score, sequence $\mathbf{A}[380:518]$ has the following 5 letter alphabet:

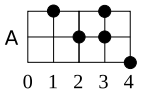

However, as noted in property 1, the stroke configuration ⁝ appears only once in the whole sequence, making it negligible for our modeling purposes; thus, we remove it from the alphabet.[17]

We now attempt an ϵ-machine reconstruction from $\mathbf{A}[380:518]$ using parameters $L_{max}=4$ and $\alpha=0.1$. Figure 13 shows the reconstructed ϵ-machine. One important thing to note is that, due to the short L_{max} used in the reconstruction, there is no hope that this model will capture the

[17]Remember that the larger the alphabet, the longer the sequence must be in order to recover from it accurate statistics (see Appendix E).

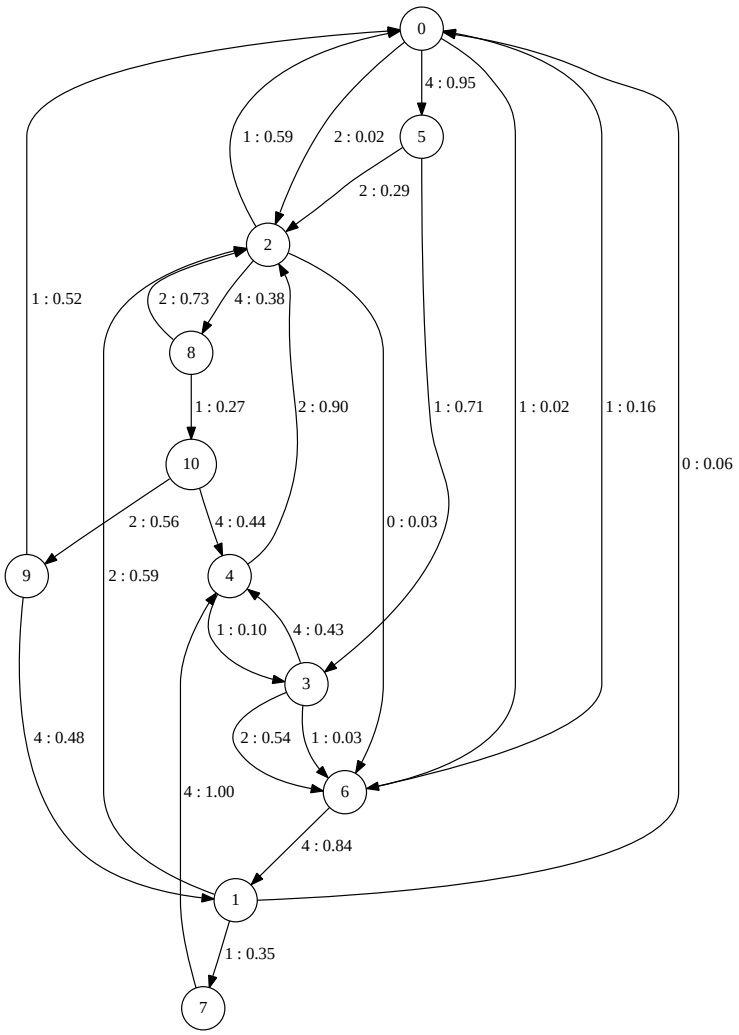

Figure 13: A single ϵ-machines reconstruction of $\mathbf{A}[380:518]$. Parameters used where $L_{max} = 4$ and $\alpha = 0.1$.

subtle trend that we have observed in the auto-correlogram for subsequences of length $L > 7$. To overcome this difficulty we might think of dividing the sequence in two and model each half independently. However, shorter sequences mean smaller statistical trust boundaries, which in turn imply smaller L_{max}, and since the trend is visible for sequences longer than 7 strokes, this will not help. Thus we are left with a model that has less memory than what the original sequence appears to have. There's nothing to prevent us from setting $L_{max} > L_{stb}$ of course, but the model will begin to exhibit probability 1 transitions caused in part by the limited amount of data in the sequence relative to the number of possible L_{max}-length words.

2.3 Section 2 [519 : 739]

2.3.1 Meta-processes

In Section 1, layers A and B are completely independent. They do not have structural interactions; they do not seem to "know about each other". Section 2 is more complex at this higher level. In this section, layers A, B and C do interact. Layers B and C seem to be coupled, and together form a process BC; this is especially clear after tick 594. Further, layer groups A and BC alternate and almost never sound together. Exceptions of this can be seen in the range [700 : 703] and tick 643.[18]

We can think of this alternation as layers communicating to "take turns". This pattern or structure can be modeled as a linked list[19] of ϵ-machines, where each process holds a reference to its alternating partner; process A has a reference to process BC and vice versa. Another way of thinking about this is by considering the existence of a third process (a *meta-process*) governing these two other processes taking place in layer groups A and BC. In this model, layers A and BC don't have to "know about each other". They just start and stop when they are told to by the meta-process orchestrating them. This is a hierarchical model, and it resembles the structure of nested lists. Yet a third model is that of a single (more complex) process consisting of a single ϵ-machine graph characterized by two local and highly connected clusters that are in turn connected to each other at least at one node.[20] I will keep the hierarchic meta-process model in mind for what follows, as it is easier to visualize.

What are the characteristics of this meta-process driving the two processes in layers A and BC? What is its behavior? Very generally, what we see is a stepped, irregular acceleration. Xenakis starts with a slow alternation between A and BC and then, in about three steps, accelerates the rate of alternation until the subprocesses last only one stroke. The section begins with each of the two subprocesses A and BC lasting from between 11 to 42 measures long, and ends with each lasting only half a measure. The perceptual effect is that of two alternating processes that gradually fuse and become one due to the increase in frequency of alternation.

[18] The coincidences between strokes in layers A and BC at measures 643, 700, 701, 702 and 703 contradict this hypothesis of a strict alternation between two processes A and BC. One could certainly view the whole section as a single high-dimensional process composed of all three layers A, B and C, but this also presents the practical problem of the *curse of dimensionality*, which we will encounter in Section 7. From a purely pragmatical point of view, these exceptions might be considered as posterior additions to a sequence resulting from the strict alternation of independent processes.

[19] A linked list is a data structure that consists of a sequence of data records, each of which references the next record in the sequence.

[20] This model is interesting in that it is a single "flat" potentially rhizomeatic graph but with fuzzy hierarchies implied by the state (node) transition probabilities.

2.3.2 Subprocesses A and B

The short isolated stroke fragments found in $\mathbf{A}[519 : 740]$ have exactly the same properties that we observed in $\mathbf{A}[380 : 518]$, with two exceptions: the doubly repeated strokes (three consecutive strokes on the same sublayer) in the ranges $[702 : 703]$ and $[724.5 : 756.5]$.[21] Likewise, the sequences found in $\mathbf{B}[519 : 740]$ recall those found in $\mathbf{B}[0 : 518]$. Thus we continue to see $\mathfrak{A}_1$-like sequences in layer A and $\mathfrak{B}$-like sequences in layer B. The $\mathfrak{B}$-like sequences we find here have two speeds: the original rate found in $\mathbf{B}[0 : 39]$ of 1 s/m and a slower one of 1/2 s/m.[22] The first three sequence fragments that appear in layers BC have a rate of 1/2 s/m. After this (tick 631), all the strokes in BC imply a constant rate of 1 s/m. In the range $[631 : 670]$, the strokes in layer B become more erratic and the sequence almost completely monophonic. At tick 672, the layer recovers its $\mathfrak{B}$-like characteristics.

2.4 Section 3 $[740 : 989]$

The whole section is a poly-tempi canon. The opening $\mathbf{B}[0 : 39]$ sequence reappears verbatim three times (with the exception of a few temporal displacements, to be discussed below) in layers A, B and C. Each layer carries the sequence at a different rate. After opening with five foreign strokes ($[740 : 759]$), layer C presents the first instance of $\mathbf{B}[0 : 39]$ at a rate of 5.5 m/s. Layer A follows, coming in at tick 772.5 and reproducing $\mathbf{B}[0 : 39]$ at a rate of 2.5 m/s. $\mathbf{B}[0 : 39]$ is restated one last time on layer B with 3.5 m/s, starting at tick 790.

There'd be little to nothing more interesting to say about this section if it weren't for the displacements of some strokes from their expected time-points. In this context there are two kinds of displacements: (1) *time* based displacement, or displacement relative to the absolute time grid (time jitter) , and (2) *delta-time* displacement, or displacement relative to the expected ISI (delta-time jitter). All the displacements found in the sequences on layers A and C are time based. This can be concluded because all other points in the sequence coincide with a constant periodic delta function (see Appendix D). In other words, the lengthening of the ISI preceding a displaced point is compensated by the shortening of the following ISI, and vice versa. For example, strokes on time-point 796.5, layer C, are displaced half a measure from 797. The ISI between this and the previous point is 5 measures, but the ISI between this point and the next is 6 measures: $5 + 6 = 11$, the expected time interval between the previous and the next points. Displaced points in layer A are at ticks 824.5 and 862. Displaced points in layer C are at ticks 796.5, 830.5 and 858.5.

[21] Only four of these doubly repeated strokes (at time-points 702, 705.5, 717 and 724.5) appear in this section. Repeated strokes, however, become recurrent units later in the piece.

[22] $\mathbf{B}[537 : 549]$ is identical to $\mathbf{B}[182 : 187]$, but twice as slow.

Except for point $\mathbf{B_1}[807]$, all displaced points in layer B are *delta-time* displacements. These altered and non-compensated ISI occur at time-point pairs $(825, 829.5)$, $(829.5, 833.5)$, $(844, 848)$, $(851.5, 854.5)$, $(854.5, 859)$, $(876.5, 880.5)$, $(884, 888)$, $(898.5, 902.5)$, $(913, 917)$ and $(927.5, 931.5)$. What motivates these displacements? The purpose of the displacements of points on layers A and C might have been the avoidance of concurrent strokes (simultaneities) between layers. Stroke $\mathbf{A}[824.5]$ is strange, however, because while it avoids a collision with a stroke on layer B, it causes a collision with a stroke on layer C. On layer B, only some displaced points seem to be motivated by collision avoidance. These are at time-points 807, 848, 854.4 and 862.5. All other displacements seem to happen "spontaneously". Note that by maintaining a strict 3.5 m/s in $\mathbf{B}[790 : 935]$, only five strokes would collide with strokes in the other layers. i.e. only five displacements on layer B would be needed to avoid collisions, half of the displacements that are actually found in the score. So, in addition to the displacements for the purpose of collision avoidance, Xenakis may have also wanted jitter in the B layer sequence. This added jitter is something that appears at various places in the score.

2.5 Sublayer C_3

Sublayer C_3 plays an interesting role in the piece. It appears first at tick 519 as a point of arrival from what might be described as the winding down of a failing machine that opened the piece in layer B. At time-point 747.5 C_3 becomes part of multi-sequence $\mathbf{C}[747.5 : 984]$, which is in turn part of the canonic structure described above. From 1000 to 2160, sublayer C_3 carries out a completely autonomous *fuzzy pulsation* process. In the range $[1000 : 1905]$, this *fuzzy pulsation* carries a relatively constant period of 20 m/s. The actual mean period is 21.047 m/s, with a variance of 4.126 m/s. This is the longest continuous process in the piece and it spans four sections (sections 4, 5, 6 and 7). At tick 1905, the pulsation gradually accelerates until it reaches a constant 1 m/s in tick 2160. Starting at tick 1727 (beginning of section 7), however, there is an interesting asymmetric doubling of the period. i.e., the pulsation period is split in two asymmetric halves. This asymmetric period is continued until the end of the acceleration at tick 2160 (see Figure 14).

2.6 Section 4 $[990 : 1410]$

Section 4 begins quite suddenly with what might be called a "resonance interlude". The range $[990 : 1202]$ is the sparsest of the whole piece. It is comprised of the opening strokes of the *fuzzy pulsation* process in sublayer C_3 and a regular pulsation of 30 m/s in sublayer B_2. This sparsity pulls our attention towards the decay of sound and the resonance of the space.

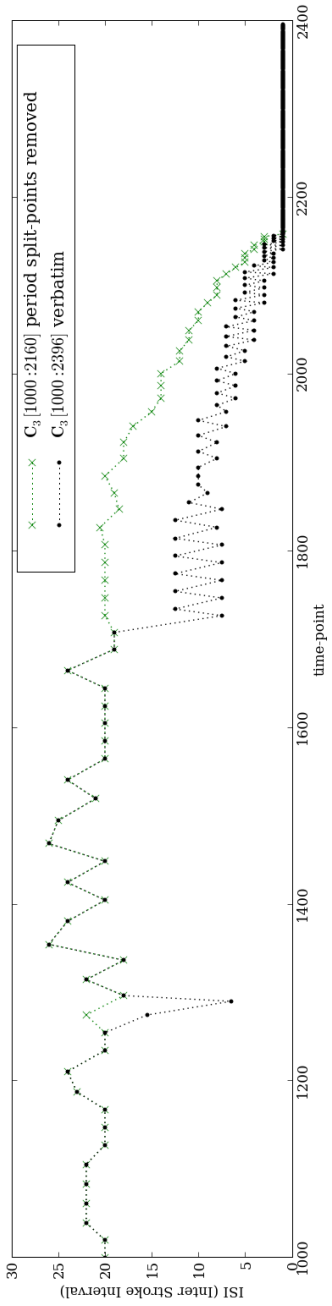

Figure 14: *Fuzzy pulsation* sequence $\mathbf{C_3}[1000 : 2396]$. The plot shows two graphs, the original sequence $\mathbf{C_3}[1000 : 2396]$ and the same sequence with period-splitting strokes removed.

This whole section seems to be composed of two distinct processes: one is the long *fuzzy pulsation* process in sublayer C_3 that starts in this section (see Section 2.5). The other process is like a multi-layer expansion of a quasi-$\mathfrak{A}_1$-like processes, similar to the two-layer expansion of a $\mathfrak{B}$-like process found in section 2, layers B and C. Here, the layers involved are A, B, D and E. Layers D and E, two of the three metal layers, first appear in this section. Layer D is introduced at 1238.5 and E at 1321. While very distinct in timbre, the two layers appear to belong together with layers A, B and D in the same process.

This multi-layer process has it peculiarities. Consecutive strokes do sometimes fall on the same sublayer. This breaks property 2 of $\mathfrak{A}_1$-like sequences. The other two properties are fairly well maintained. Exceptions to property 1 (strict mono-stoke configurations) appear at ticks 1323.5, 1342, 1346.5, 1352, 1392, 1402, 1405.5 and 1407. Those sequences satisfying properties 1 and 3 of $\mathfrak{A}_1$-like sequences I shall refer to as $\mathfrak{A}_2$-like.

Strokes on sublayer A_1 almost always appear in groups of repeated strokes with a rate of 2 s/m. These repeated strokes, which we just got a glimpse of at the end of Section 2 [702 : 726], now become quite present. This "stuttering" element is more prominent in sublayer A_1, but it also "infects" sublayers 1 and 3 of layers B and D, as well as layer E. Their appearance is modest at the beginning of the section, but they gradually become more present as the section progresses.

The repeated strokes on sublayer A_1 seem to be coupled with the *fuzzy pulsation* process; repeated strokes in A_1 fall immediately before, immediately after, or around the strokes on sublayer C_3, but never simultaneously: sublayer A_1 "knows" about the *fuzzy pulsation* in C_3. At the same time, these strokes on sublayer A_1 alternate with the strokes in the other layers, suggesting that they belong together. Thus, while the *fuzzy pulsation* process is independent, this other multi-layer process seems to be somewhat coupled with it.

This multi-layer process also displays two mildly gradual transitions. One is the progressive elimination of silence breaking up the continuity of the sound. The other is the progressive move from clean $\mathfrak{A}_2$-like sequences (range [1203 : 1237]) to a sequence of relatively long stretches of repeated strokes on the same sublayer.

2.7 Section 5 [1411 : 1589]

Section 5 has three independent simultaneous sequences. One is the *fuzzy pulsation* sequence that comes from section 4 and continues into section 6. The other two are $\mathfrak{A}_2$-like monophonic sequences, each spanning two layers (AB and DE) and each running at a different rate: one at 1 s/m, the other at 2 s/m. In the range [1411 : 1538], sequence **AB** runs at 1 s/m and **DE** at 2 s/m. At tick 1539, the speeds are exchanged: sequence **AB** runs at 2

s/m and **DE** at 1 s/m.

Glitches The $\mathfrak{A}_2$-like sequences on layer groups AB and DE have some interesting "anomalies" that call our attention. All the strokes in **AB**[1411 : 1538] fall on a tick, except one: the stroke expected at tick 1478 seems to have been delayed by 1/2 measure, falling at time-point 1478.5. In both **AB**[1411 : 1538] and **DE**[1411 : 1538], almost every time-point has at most one stroke (i.e., no more than one sublayer is sounding at a given time). In **DE**[1411 : 1538] there are three instances where this is not the case. At time-points 1423.5 and 1463.5 we find two simultaneous strokes in two sublayers. In both cases there is a gap half a measure later: an empty time-point where we expect a stroke for the continuity of 1/2 m/s to be maintained. It seems, again, as if strokes had been displaced from their expected places, creating a glitch in an otherwise constant pulse. In sequence **AB**[1539 : 1589] we find two more instances of this at time-points 1557 and 1583.5. Again we find that the strokes seem to have been anticipated by half a measure, creating simultaneities at 1557 and 1583.5 and gaps at 1557.5 and 1584. In **DE**[1488] there is another pair of simultaneous strokes, although this time not followed by a gap. Here, the extra stroke seems to have been simply added. The same thing occurs at **AB**[1418], **AB**[1426] and **DE**[1586]. Where do these "glitches" come from? Why are they there? They are relatively few and not very perceptually salient; yet, they are too frequent to be accidental...

2.7.1 **DE**[1411 : 1538]

There are 8 distinct stroke configurations in sequence **DE**[1411 : 1538]. The simultaneities at time-points 1423.5, 1463 and 1488 are unique and statistically negligible so we remove them from the alphabet,[23] leaving us with the following 5 letters:

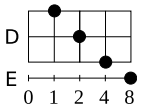

While we have decided to remove these letters from the alphabet, it is still not clear how we should handle them in the sequence. The fact that they seem like "anomalous" temporal displacements suggests that instead of eliminating them we should simply move one of the coinciding points so that both the simultaneities and the time-gaps are "fixed" in one move. Having done this we are left with only two time-points with no strokes in the whole sequence, making them also quite negligible. Removing these silent time-points from the alphabet leaves us with only four letters:

[23] Remember that the larger the alphabet size, the longer the sequence must be to extract from it meaningful statistics.

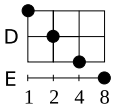

From this new "clean" sequence we compute the generating ϵ-machine. **DE**[1411 : 1538] is 255 tokens long, so the trust boundary is very close to $L_{tb} = 4$ given our new alphabet or cardinality 4 ($4^4 = 256$).[24] As can be seen in Figure 15, the predictability gain for **DE**[1411 : 1538] starts to grow at $L = 3$ and peaks at $L = 5$. Because $L_{tb} = 4$ we cannot know if

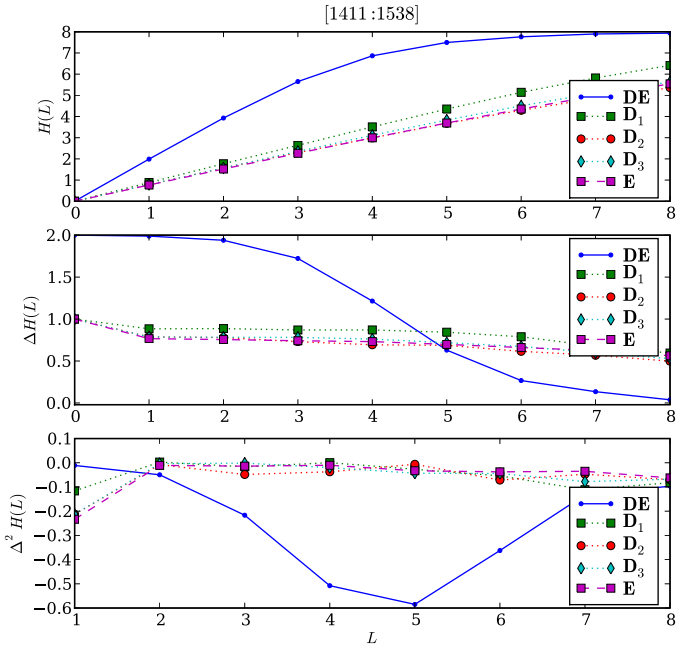

Figure 15: "Cleaned" **DE**[1411 : 1538] entropy measurements.

the gain at $L = 5$ is due to insufficient data or if it's a true attribute of the process. Using a significance level of 0.1 and $L_{max} = 4$, we obtain the ϵ-machine shown in Figure 16.

2.7.2 **AB**[1411 : 1538]

If we take every time-point as a valid configuration, then sequence **AB**[1411 : 1538] has 9 letters in its alphabet. For an alphabet this big, **AB**[1411 : 1538] is too short (128 points) to extract from it any meaningful statistics. Thus, the two only instances of simultaneities in **AB**[1411 : 1538] at 1418 and

[24] The soft statistical trust boundary is actually 4.0137.

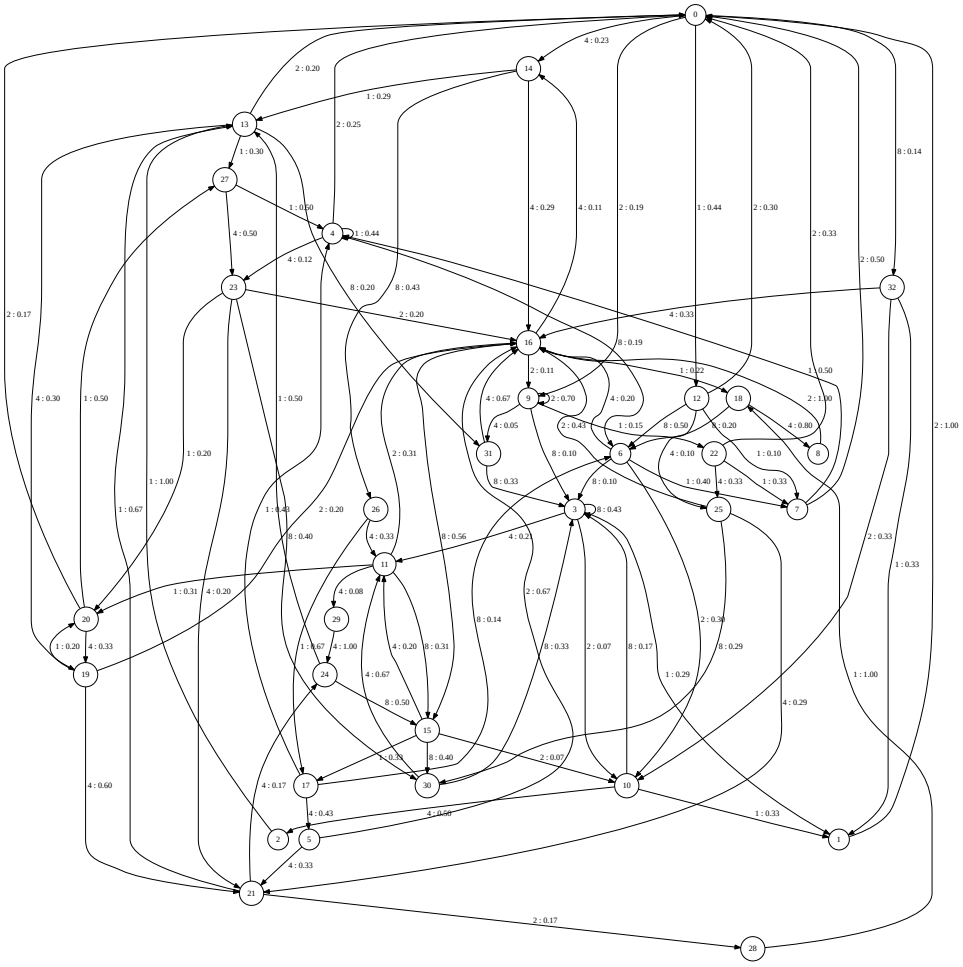

Figure 16: ϵ-machine reconstructed from **DE**[1411 : 1538] with parameters $L_{max} = 4$, $\alpha = 0.1$.

1426 we take as negligible and eliminate them from the sequence and the alphabet, leaving us with the following 7 letters:

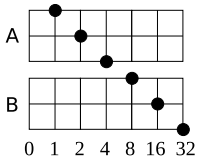

The computed ϵ-machine is shown in Figure 18.

2.7.3 **AB**[1539 : 1589]

As we did with sequence **DE**[1411 : 1538], we modify the statistically negligible time-point configurations in the sequence in order to obtain a reduced

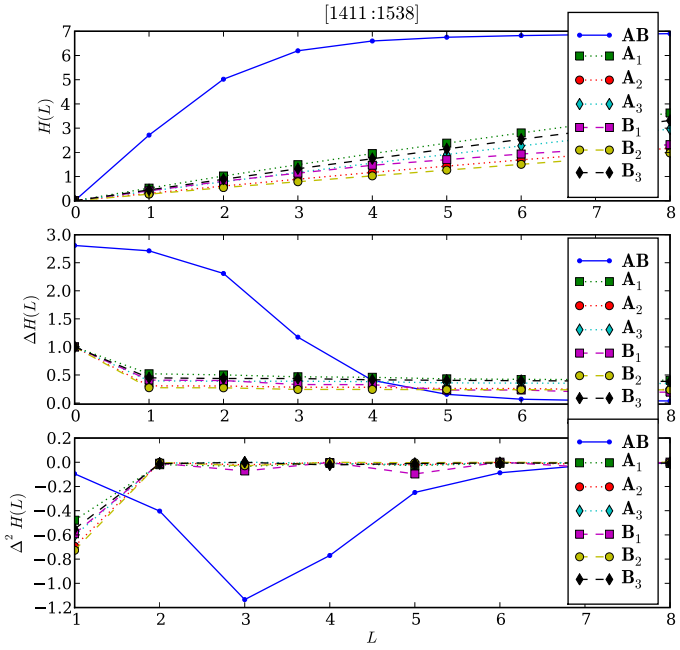

Figure 17: "Cleaned" **AB**[1411 : 1538] entropy measurements.

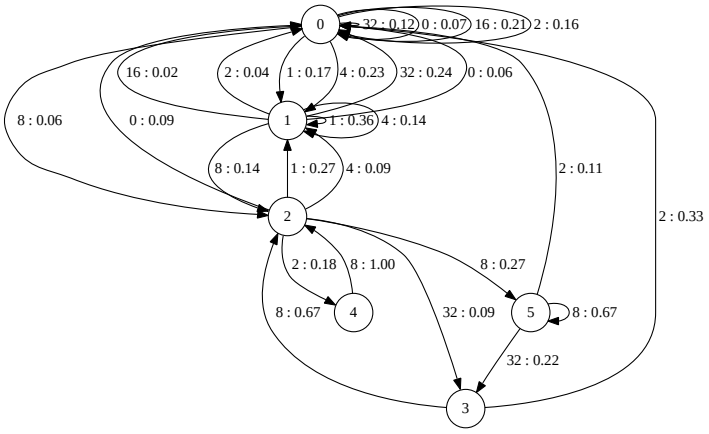

Figure 18: ϵ-machine reconstruction of **AB**[1411 : 1538]. Parameters: $L_{max} = 3$, $\alpha = 0.1$.

alphabet. The three points we modify are the simultaneities followed by a rest at [1557 : 1557.5] and [1583.5 : 1584], and the rests at 1562. For the first two, we displace one of the simultaneous points by half a measure. This simultaneously covers the empty spaces and makes the sequence

strictly monophonic. The emtpy time-point at 1562 we simply remove. The resulting alphabet is:

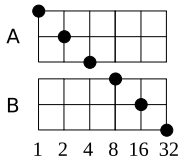

As can be seen in Figure 19, the maximum predictability gain is at $L = 2$. This is fortunate, since the trust boundary for a sequence of length 95

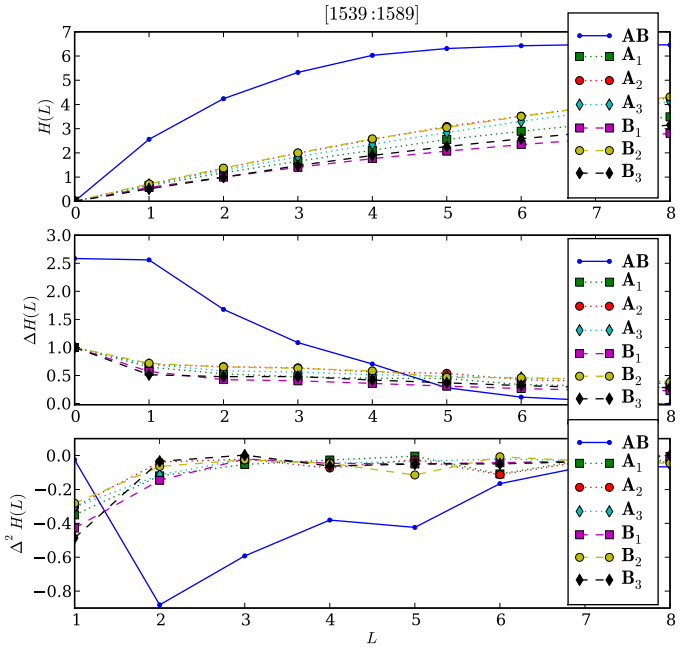

Figure 19: "Cleaned" **AB**[1539 : 1589] entropy measurements.

and alphabet cardinality 6 is just $L_{tb} = 2$. The ϵ-machine recovered from **AB**[1539 : 1589] is shown in Figure 20.

2.7.4 **DE**[1539 : 1589]

Except for the last strokes at tick 187, this short sequences is strictly monophonic. For the ϵ-machine reconstruction, we remove the last token from the sequence. This leaves us with an alphabet of cardinality 4. Figure 21 plots the entropy measures of the sequence and its sublayer sequences; Figure 22 shows the recovered ϵ-machine.

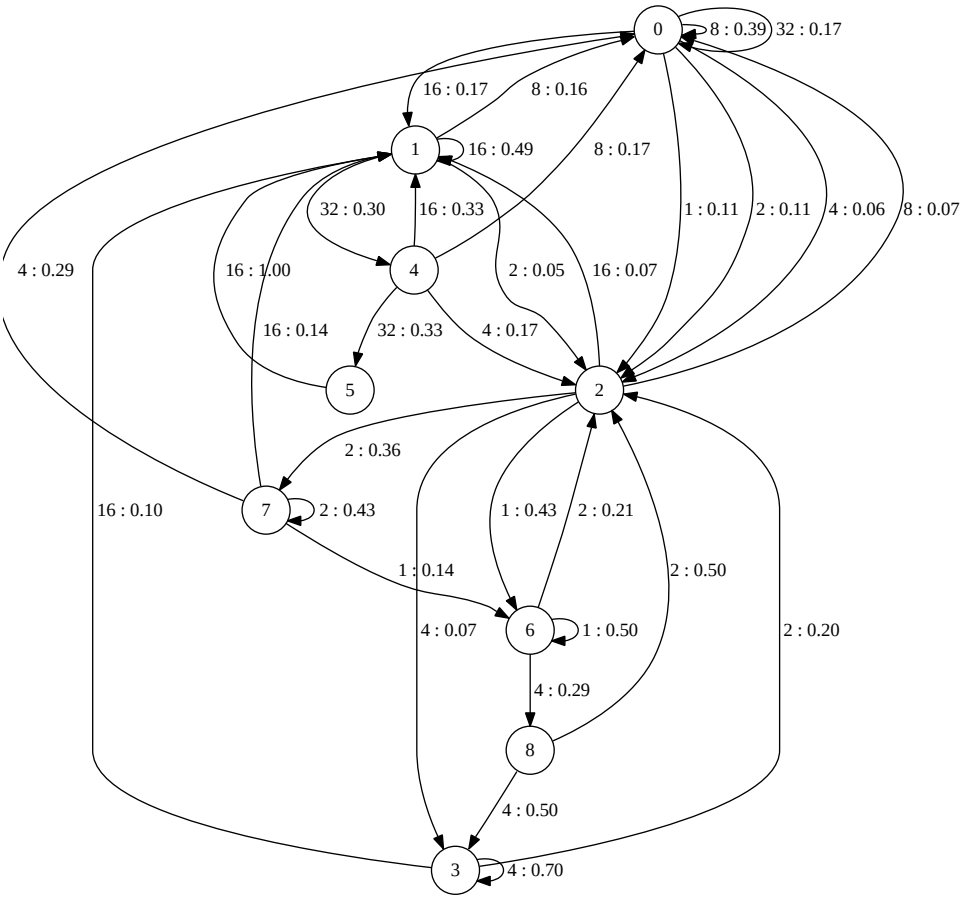

Figure 20: ϵ-machine reconstruction of **AB**[1539 : 1589]. Parameters: $L_{max} = 3$, $\alpha = 0.1$.

2.8 Section 6 [1590 : 1720]

In this section, Xenakis introduces a new kind of sequence characterized by pairs of sequential strokes. The two strokes in these pair units are always one measure apart, and one of the two strokes is always accented while the other is not. This *twin-strokes* sequence spans 384 measures, from tick 1614 to tick 1998. Of the 24 twin-strokes total, 20 fall on sublayer B_3, three fall closely together on sublayer D_3 ([1715 : 1724]) and one, at tick 1897, on sublayer D_1. Notice that the thee stroke pairs in $\mathbf{D_3}$[1715 : 1724] are very close to each other (3 measures apart) compared to the stroke pairs in sublayer B_3. This fact and the collision at tick 1719 suggests that the two sequences belong to separate, albeit identical, processes.

In addition to this new *twin-strokes* sequence and the *fuzzy pulsation* continuing its way towards the end of the piece, the section also contains what appears to be a multi-layer $\mathfrak{A}_2$-like process very much like the one

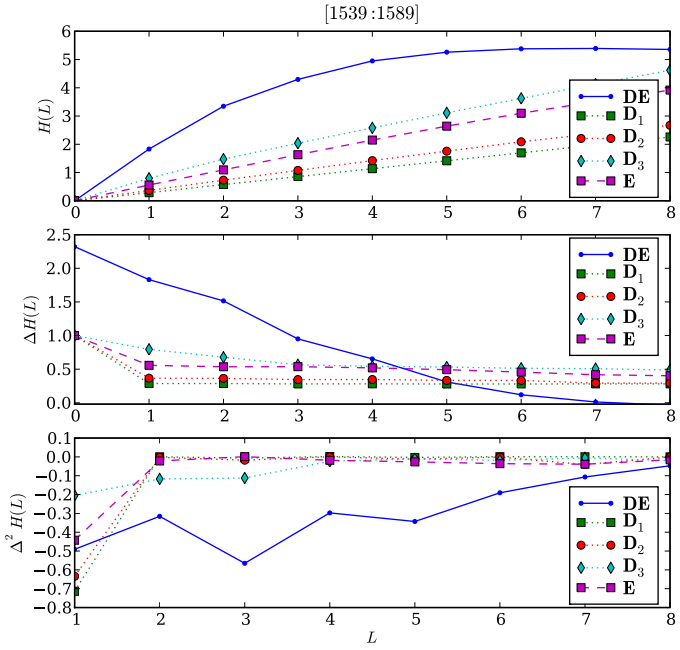

Figure 21: "Cleaned" **DE**[1539 : 1589] entropy measurements.

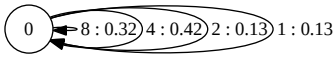

Figure 22: ϵ-machine reconstruction of **DE**[1539 : 1589]. Parameters: $L_{max} = 3$, $\alpha = 0.1$.

found in section 4. Again, we find the repeated strokes on sublayer A_1 falling exactly around the strokes of the *fuzzy pulsation* in sublayer C_3. The layers involved in this multi-layered process are A, B_2, D and E. Short sequences of strokes again alternate with silence. Notice that the strokes of the three processes (the *fuzzy pulsation* with its halo of A_1 repeated strokes, the *twin-strokes* and the multi-layer $\mathfrak{A}_2$-like process) never coincide, but are arranged in a strictly sequential stream, a single stroke at a time.

2.9 Section 7 [1720 : 2175]

This section sees the continuation of two independent processes: the *fuzzy pulsation* and the *twin-stroke* process that began in the previous section. This *twin-stroke* sequence thins out towards the end of the section, appearing for the last time at tick 1998. The *fuzzy pulsation*, however, moves into the perceptual foreground. It is in this section that it doubles its density (tick

1734.5) by asymmetrically inserting one more stroke per period, and where it accelerates until reaching a rate of 1 s/m (see Section 2.5).

Arborescent pulsations A third, new kind of "texture" dominates this section. The little "tutter virus" that was introduced at the end of Section 2 has here taken over. Layers A, B, D and E all serve as the canvas for what could be described as a gradually evolving arborescent contour[25] filled in with regular machine-gun-like repeated strokes. The texture is slightly reminiscent of the arborescences found in *Evryali* (1973), written two years before *Psappha*. These "machine-guns" fire at constant rate of 2 s/m until tick 2023, where the rate is doubled (or tripled) via tremolo indications (the score reads: "2 or 3 strokes per point"[26]). Notice that strokes in the different sublayers overlap, sometimes for more than 3 measures. The distribution of these "machine-gun" sequences is broad and irregular at the beginning of the section ([1748 : 1820]) but it gradually narrows and branches, ending in two oscillating and accelerating streams on layers A and D ([2130 : 2173]).

Deriving a stochastic model from this multi-sequence is a challenge with the methods we've been employing. The two problems are the large alphabet of the sequence relative to its length, and the very clear and gradual trend that the sequence has. Such a trend implies a large amount of memory but, given the length of the sequence relative to its alphabet size, there is no way of reaching that length before hitting the trust boundary. The actual number of stroke configurations found in the sequence is 50,[27] but the sequence is only 850 time-points long. $50^2 = 2500$, so the trust boundary L_{tb} for this arborescent sequence is only 1. Thus, another approach is needed.

2.10 Section 8 [2176 : 2397]

The final section comprises two streams. One is the continuation of the constant 1 m/s pulse in sublayer C_3. The other is a retrograde presentation of a segment of the long $\mathbf{A}[380 : 518]$ sequence. After a short 6 measure return of the cyclic sequence $\mathbf{A}[47.5 : 60]$ at tick 2266, Xenakis finishes the piece with an almost exact retrograde of $\mathbf{A}[400 : 518]$. The only difference is a three stroke gap in the range [2285.5 : 2286.5]. This time, however, these sequences appear in layer F, making this the first and only appearance of the layer in the whole piece.

Xenakis breaks the monotony of the constant 1 m/s sequence in sublayer C_3 by accenting some of the strokes in the sequence. All the IAIs (Inter Attack Intervals) belong to the Fibonacci series. From tick 2176 to tick 2265

[25] In [12], Harvey describes arborescences as "... proliferations of melodic lines created from a generative contour".

[26] "*2 ou 3 coups par point*"

[27] Although the 10 sublayers composing this sequence/texture imply a potential $2^{10} = 1024$ alphabet size.

we find that only the first four numbers of the series are used: $\{1, 2, 3, 5\}$. From 2255 to the end of the piece, the accents follow the series exactly, starting from 2 and ending on 55: 2, 3, 5, 8, 13, 21, 34, 55.

3 Closing Remarks

As noted in the preamble, Xenakis' interest in composing *Psappha* was in the "architecture" of pure, abstract, rhythm. He was interested in the *structural relationships* between the sound tokens, not in the sound tokens themselves.[28] In keeping with this line of thought, I have here presented a purely structural analysis, focusing on the relationships between the strokes and the way these outline the different rhythmic patterns.

I have dissected *Psappha* while simultaneously attempting to partially resuscitate it by deriving localized stochastic models, in the form of ϵ-machines, from the fragments obtained. These models pertain mostly to local sequences. With the exception of the meta-process discussed in section 2, I made no attempt to derive a model of process interactions or a global model for the whole piece. Further, I did not attempt to model all the local stroke sequences. In particular, the arborescent sequence of section 7 poses a special challenge. In general, the main difficulty in deriving these models is the estimation of probability measures from relatively little data. The insatiable hunger for data that these models have makes them hard to estimate from sequences that are short relative to their alphabet size. There is an alternative to the standard ϵ-machine models we have used here; rather than taking the arborescent contours as a single multi-dimensional sequence from which to derive a global model, we could consider them as spatio-temporal evolutions, like cellular automata. From this perspective, instead of deriving a single global model from a high dimensional sequence, multiple local models can be computed from purely local interactions. Shalizi[20] proposed a particularly interesting approach to this problem using the notion of "light cones" to define the space of interactions, and he applies this technique precisely to cellular automata. At any rate, it seems clear that when confronted with non-periodic or semi-periodic sequences, global ϵ-machines can provide a more compact representation with more relevant parametric handles than sieves do.

[28] To what extent did Le Corbusier influence Xenakis' aesthetic?: "[...] architecture which is everything—but is not the "decorative arts." Tail pieces and garlands, exquisite ovals where triangular doves preen themselves or one another, boudoirs embellished with "pours" in gold and black velvet, are now no more than the intolerable witnesses to a dead spirit. These sanctuaries stifling with elegancies, or on the other hand with the follies of "Peasant Art," are an offence.[p.91] Art is an austere thing which has its sacred moments. We profane them. A frivolous art leers upon a world which has need rather of organization, of implements and of methods and which is pushing forward in travail towards the establishment of a new order.[p.100][5]

In the process of analyzing and modeling *Psappha*, a couple of fundamental structural patterns and procedures have become manifest. Some are obvious but nevertheless important and worth mentioning:

1. *Parallel* vs. *sequential* arrangement. Two elements (sequences) can be arranged in sequence or in parallel (simultaneously). Elements arranged in parallel can be composed as independent sequences or as a single multidimensional sequence (i.e., the vertical configurations between two or more simultaneous sequences may be accounted for or not.). The former we call *layering*, the later, *braiding*; $\mathbf{B}[0:39]$ is an example of braiding. The relationships between parallel sequences can, however, have multiple degrees of interdependence (coupling). Thus, there is a continuous spectrum between total independence and total coupling between parallel sequences.
2. *Modulation.* A process modulates another process (or a set of processes) if the first governs a parameter of the second.
3. *Continuous* vs. *discontinuous.* Generally speaking, change can occur gradually (continuously) or abruptly (discontinuously). Examples of continuous change: the *fuzzy pulsation* process, the gradual perceptual merging of alternating processes by increasing the frequency of alternation between processes, and the gradual slowing down in $\mathbf{B}[200:519]$. Example of discontinuous changes: the alternating entrances of stroke sequences in $\mathbf{A}[0:519]$.
4. *Repetition.* Repetition can occur at various scales and levels of abstraction, i.e., at the level of sequences (verbatim repetition) and at the level of processes. e.g., Section 8 is a kind of coda that ends the piece with a reversed restatement of material appearing early in the piece.
5. (Time) *scaling*: e.g., the poli-tempi canon in section 3.

From a purely formal/structural perspective, time scale seems like a secondary issue. But the impact this has on perception cannot be ignored (or minimized) given the enormous qualitative variations one can experience from the exact same sequence rendered at different speeds. The continuum rhythm-pitch observed by several composers ([9, 22]) is a most obvious acknowledgment of this. In the score, Xenakis uses a wide range of time scales. Pertaining ISIs specifically, we find a proportion of approximately $97:1$ between the longest (the opening of section 4, with ISIs of 20 measures at a tempo of $\geqslant$ 110MM approximately) and the shortest (the strokes with tremoli in section 8 (assuming 2 strokes per point) with a tempo of $\geqslant$ 134MM) ISIs.

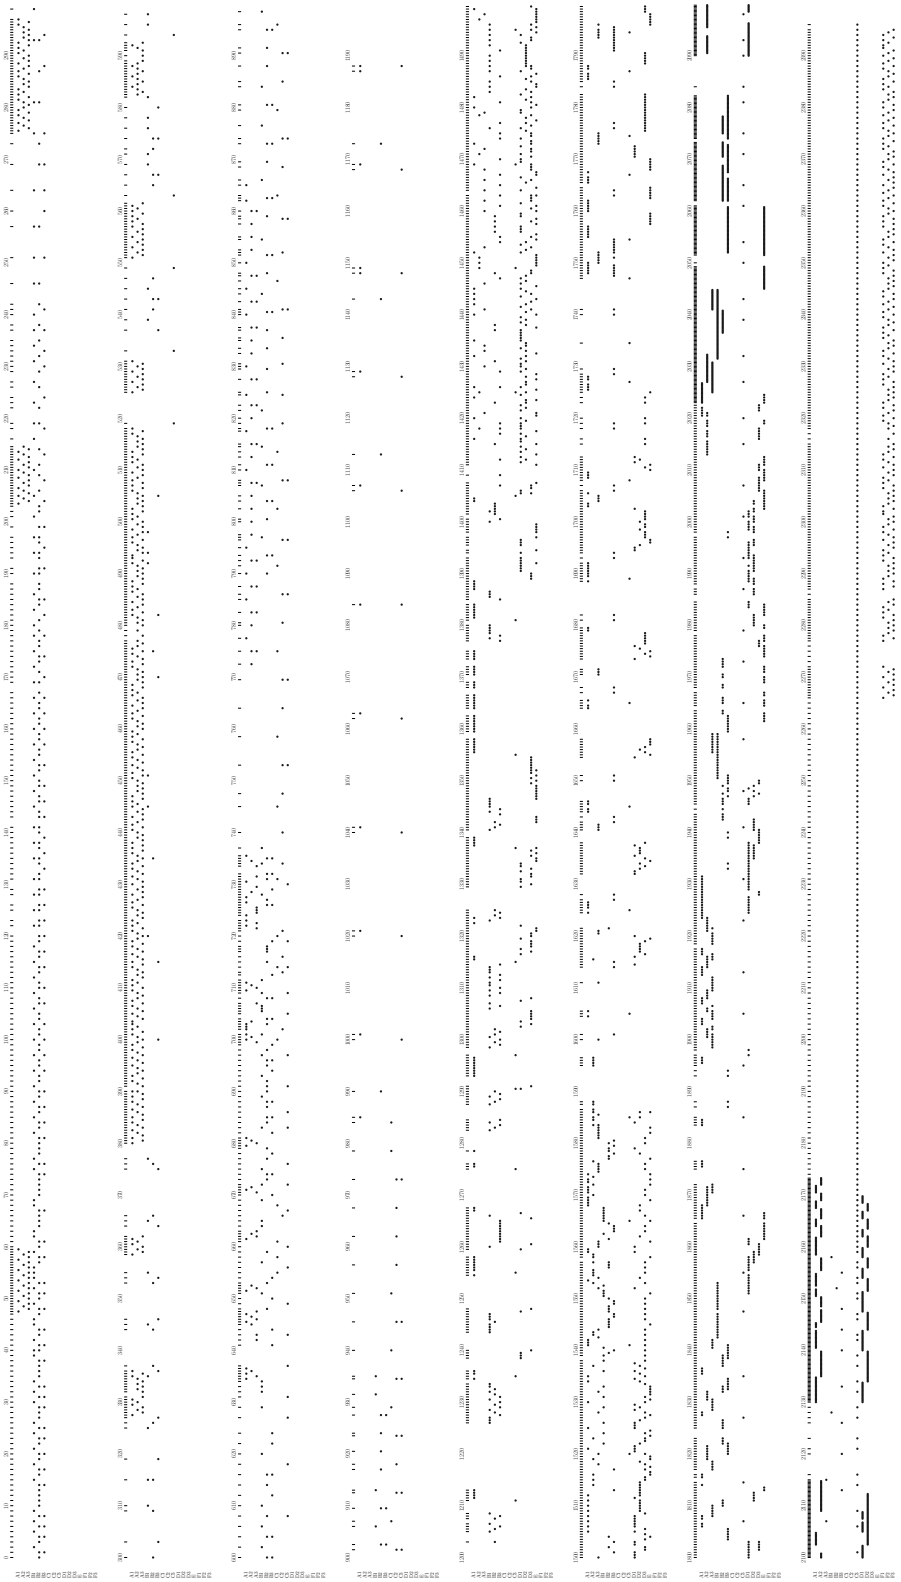

Figure 23: Bare bone score of *Psappha.* The score shows all the strokes found in the original. Here, however, no accents, dynamic markings, or tempo indications are shown. The time-grid has also been removed for clarity. Instead, tick marks above each system are drawn at the time-points where there is at least one stroke.

A Glossary

configuration a particular combination of strokes and non-strokes on different sublayers, at a single time-point.

IAI Inter Accent Interval.

ISI Inter Stroke Interval.

layer a stratum in *Psappha* indicated with upper case letters. e.g. **A**, **B**, etc.

m/s measures per stroke.

measure the time interval spanned by two consecutive *ticks*.

multi-sequence in *Psappha* specifically, a succession of stroke configurations ranging more than one sublayer.

sequence in *Psappha* specifically, a succession of strokes on a single sublayer.

s/m strokes per measure.

stream segregation the perceptual grouping and separation of stimuli to form coherent mental representations of the sensorial stream.[2]

sub-layer an instrumental sub-stratum in *Psappha* indicated by a number 1, 2, or 3, forming part of a layer.

tatum the greatest common divisior of all durations in the score. In *Psappha* this is one quarter or a third (depending on the subdivision taken at measure 2023) of a tick (with the notable exception of the quintuplet found at measure 70).

tick a vertical line demarkating the boundaries of a measure.

time-point a time "slot", multiple of the *tatum*, that can be occupied by a stroke.

B Notation

$\cup, \cap$	set theoretic union and intersection.
$\vee, \wedge, \oplus$	logical OR, AND and XOR (exclusive OR).
$rc(m, r)$	a residue class with modulo m and residue r.
m_r	in the context of sieves, a residue class $rc(m, r)$.
$\mathrm{P}(a)$	the probability of a.
$\mathrm{P}(a\|b)$	the probability of a given that we know b.
$[a : b]$	a range of numbers between a and b, inclusive on both ends.
$\mathbf{X}[a : b]$	a multi-sequence (layer sequence) segment in the range $[a : b]$.
$\mathbf{X}_n[a : b]$	a sublayer sequence segment in the range $[a : b]$.
$\|\mathcal{A}\|$	cardinality of alphabet $\mathcal{A}$

C Residue Classes

Two integers c and r are said to be congruent modulo m if there exists an integer q such that $c = qm + r$. This is written $c \equiv r \pmod{m}$. The set of all integers c congruent r modulo m is called a *residue class* modulo m, here denoted $rc(m, r)$.

Definition 1 (Residue class)**.**

$$rc(m, r) = \{x \in \mathbb{Z} : x = qm + r\}$$

Example. The residue classes $rc(12, i)$ for $i \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\}$ correspond to the pitch-class sets we all know and love. Pitch class 0 (C): $rc(12, 0) = \{\ldots, -12, 0, 12, 24, 36, \ldots\}$, pitch class 1 (C♯): $rc(12, 1) = \{\ldots, -11, 1, 13, 25, 37, \ldots\}$, etc.

D Xenakis Sieves

> ... every well-ordered set can be represented as points on a line, if it is given a reference point for the origin and a length u for the unit distance, and this is a sieve.[25]

There are a couple of ways in which one can formalize sieves. Xenakis' discussion on sieves is set theoretic. In [24], he defines an elementary sieve as a *residue class*, an equivalence relation of congruence modulo m (See Appendix C). In his various texts, Xenakis uses different kinds of notation to represent a residue class. In [24], he notates m_r, the modulo subscripted with the residue, while in [25] he uses the tuple notation (m, r). Here we will also use $rc(m, r)$ to make the residue class more explicit. A residue class (or elementary sieve) is thus an infinite sequence with its difference always equal to a constant m, the modulo of the residue class. Thus all basic sieves are periodic with period m.

A *complex sieve* is defined by an expression that combines multiple elementary sieves, or residue classes, via set operators like union ($\cup$) , intersection ($\cap$) and complementation (superscript c). Also important is the operation of exclusive disjunction ($\oplus$); Xenakis, however, never makes use of it.

Example (union). The union of $rc(2, 0)$ and $rc(3, 0)$:

$$\begin{aligned} rc(2, 0) \cup rc(3, 0) &= [\ldots, -2, 0, 2, 4, 6, 8, \ldots] \cup [\ldots - 3, 0, 3, 6, 9, \ldots] \\ &= [\ldots, -3, -2, 0, 2, 3, 4, 6, 8, 9, \ldots] \end{aligned}$$

Example (intersection).

$$rc(2, 0) \cap rc(3, 0) = [\ldots, -6, 0, 6, 12, 18, \ldots]$$

Example (exclusive disjunction)*.*

$$rc(2,0) \oplus rc(3,0) = [\ldots, -2, 2, 3, 4, 8, 9, 10, \ldots]$$

Xenakis' sieves can also be thought of as a boolean analogue of the Fourier transform in the sense that a complex sequence of values can be decomposed into a collection of simple, periodic functions, each with a different period and/or phase. Inversely, a complex sequence can be constructed from a combination of simple basis.

Let's define a function $\delta : \mathbb{N} \to \{0, 1\}$ to be an elementary sieve basis function:

Definition 2 (Sieve basis function)**.**

$$\delta_m^r[n] \triangleq \begin{cases} 1, & \text{if } n \equiv r \pmod{m} \\ 0, & \text{otherwise.} \end{cases}$$

To make the distinction between sets of residue classes and these boolean sequences, we shall use the standard boolean operators when dealing with δ functions. The above examples would then be expressed as as:

$$\delta_2^0[n] \vee \delta_3^0[n] = [\ldots, 1, 0, 1, 1, 1, 0, \ldots]$$

$$\delta_2^0[n] \wedge \delta_3^0[n] = [\ldots, 1, 0, 0, 0, 0, 0, \ldots]$$

$$\delta_2^0[n] \oplus \delta_3^0[n] = [\ldots, 0, 0, 1, 1, 1, 0, \ldots]$$

In his 1990 *Sieves* article[25], Xenakis discusses both the construction of complex sequences by the combination of elementary sieves and the inverse transformation: the decomposition of a complex sequence into simple sieves. The algorithm he presents might be call a XOR decomposition because the basis δ functions obtained from the decomposition are mutually exclusive (the intersection of their corresponding residue classes yields the empty set) and thus can be XORed to obtain the original complex sequence. For example, the sequence $[1, 0, 1, 1, 1, 0, \ldots]$ would be decomposed into the two basis $\delta_2^0[n]$ and $\delta_6^3[n]$: $[1, 0, 1, 0, 1, 0, \ldots] \oplus [0, 0, 0, 1, 0, 0, \ldots]$. Jones[14] considers this XOR decomposition to be "flawed" in part because the basis resulting from the decomposition are orthogonal.[29] This is not a bad thing *per se* of course, but one might indeed want the decomposition of $[1, 0, 1, 1, 1, 0, \ldots]$ to yield $rc(2, 0)$ and $rc(3, 0)$. This we call an OR decomposition because the δ basis obtained can be ORed together to recover the original complex sieve.

[29] In [14], Jones proposes an alternative algorithm for sieve decomposition that offers a compromise between exactness and compact representation. It is a lossy OR decomposition that reconstructs an approximation to a complex sieve by combining the extracted basis via the OR operator.

Definition 3 (Normalized δ correlation)**.** Given a boolean sequence $\mathbf{s}$, for a period p and a phase (or index) i, with p a divisor of len($\mathbf{s}$) and $0 \leq i < p$

$$\langle \mathbf{s}[n], \delta_p^i[n] \rangle \triangleq \frac{1}{\sum_{n=0}^{len(s)-1} \delta_p^i[n]} \sum_{n=0}^{len(s)-1} \mathbf{s}[n]\delta_p^i[n] \tag{1}$$

Definition 4 (δ spectrum)**.**

$$S[p, i] \triangleq \langle \mathbf{s}[n], \delta_p^i[n] \rangle \tag{2}$$

Example. Compute the δ spectrum of $\mathbf{s} = \delta_2^0[n] \vee \delta_3^0[n] = [1, 0, 1, 1, 1, 0]$.

(3)

p							
1	0.6	0	0	0	0	0	
2	1	0.3	0	0	0	0	
3	1	0.5	0.5	0	0	0	
6	1	0	1	1	1	0	
	0	*1*	*2*	*3*	*4*	*5*	*i*

Notice in the table that periods 4 and 5 are skipped because 4 and 5 are not divisors of 6.

E Information Theory

The most fundamental notions in information theory are those of information and *entropy*. Let X be a random variable taking the possible values x from the finite alphabet $\mathcal{X}$. We write $\mathrm{P}(x)$ to denote the probability of X taking value x.

Definition 5. The *Shannon entropy* of X:

$$H[X] \triangleq -\sum_{x \in \mathcal{X}} \mathrm{P}(x) \log_2 \mathrm{P}(x) \tag{4}$$

Definition 6. The *total Shannon entropy*[7][30] of length-L subsequence blocks, for $L > 0$:

$$H(L) \triangleq -\sum_{s^L \in \mathcal{A}^L} \mathrm{P}(s^L) \log_2 \mathrm{P}(s^L), \tag{5}$$

Definition 7. The *mutual information* between two random variables X and Y:

$$\begin{aligned} I[X;Y] &\triangleq H[X] - H[X|Y] && (6)\\ &= H[Y] - H[Y|X] && (7)\\ &= H[X] + H[Y] - H[X,Y] && (8)\\ &= H[X,Y] - H[X|Y] - H[Y|X] && (9) \end{aligned}$$

[30] In [7], Crutchfield and Feldman present a very clear and thorough exposition of various information theoretic measures; highly recommended.

It is non-negative and zero when X and Y are independent.

Definition 8. The *mutual information distance* between X and Y:

$$d(X,Y) \triangleq H[X,Y] - I[X;Y] \tag{10}$$

Definition 9. The *relative mutual information distance*[15]:

$$D(X,Y) \triangleq \frac{d(X,Y)}{H[X,Y]} \tag{11}$$

$$= 1 - \frac{I[X;Y]}{H[X,Y]} \tag{12}$$

The normalized distance $D(X,Y)$ is useful in comparing the mutual information between probability distributions derived from varying length-L sequence blocks by keeping all distances between 0 and 1 always.

Computing Probability Distributions

Estimating accurate probability distributions from limited amounts of data is a non-trivial problem. Because probabilities form the basis of many measurements in information theory (most notably *entropy*), a good amount of literature is devoted to the problem of estimating probabilities from finite data and new approaches are still being proposed (e.g., [4][13]).

The simplest way of estimating probabilities is by counting the relative frequencies of all possible outcomes. Given a sequence $\overline{s}$ of length N with letters from an alphabet $\mathcal{A}$, we compute the marginal probability of each letter in the alphabet by counting the number of times each letter appears in $\overline{s}$ and dividing by the total number of letters:

$$\mathrm{P}(a) = \frac{\mathrm{count}(a)}{N} \tag{13}$$

All probability estimates in this analysis are computed in this way.

Statistical Trust Boundary

The entropy of a process depends on two things: the probability distribution, as is clear from the definition, and the cardinality $|\mathcal{A}|$ of the alphabet $\mathcal{A}$ (or $\mathcal{A}^L$ for length-L subsequences with letters taken from $\mathcal{A}$). The maximum entropy value that a process with an alphabet of size $|\mathcal{A}|$ can have is $\log_2(|\mathcal{A}|)$, and it is reached when the probability distribution over the alphabet is uniform. i.e., when all symbols in the alphabet are equally probable. This implies that out of two processes with identical probability distributions, the one with the larger alphabet will have the greater entropy. This in turn means that, for any given process, $H(L)$ must be a monotonically

increasing function. i.e., $H(L) \leq H(L+1)$ always. But when deriving probabilities statistically from finite data we find that, past a certain length value L_{tb}, the probability estimates computed with the simple counting method described above begin to "narrow", so to speak, and become less and less reliable due to the reduction in the number of data points to count. As can bee seen in Figure 24, after a certain length L the rate of growth of the entropy $H(L)$ function begins to drop, to the point where the entropy gain turns negative. This contradicts the fact that $H(L) \leq H(L+1)$ always. Here, $H(L)$ is not monotonically increasing because the number of words found in the sequence begins to decrease as L increases; at $L = N$, there is just one word (the whole sequence) and the entropy is 0: $H(N) = 0$. The boundary length L_{tb} at which the entropy gain drops due to the finite length of the sequence and not due to the inherent properties of the sequence, we call the *statistical trust boundary.*

Definition 10 (Hard statistical trust boundary)**.** Let $\overline{S}$ be a sequence of random variables of length N with elements taken from alphabet $\mathcal{A}$:

$$L_{tb} = \lfloor \log_2(N) / \log_2(|\mathcal{A}|) \rfloor$$

L_{tb} tells us what the longest word length (subsequence length) in a sequence $\overline{S}$ is for which it is still possible to find *at least* one of each of the $|\mathcal{A}^{L_{tb}}|$ words in alphabet $\mathcal{A}^{L_{tb}}$.

Example. Let $\overline{s}$ be a binary sequence of length $N = 140$, with alphabet $\mathcal{A} = \{0, 1\}$. There are 2 possible words of length $L = 1$, 4 possible words of length $L = 2$ (2^2), 8 possible words of length $L = 3$ (2^3), etc. In this case, the maximum word length L_{tb} for which there could possibly be at least one of each of the words of length L_{tb} of smaller in the sequence, is 7: $2^7 = 128$. Eight is too long: $2^8 = 256 > N = 140$. This is what the *hard statistical trust boundary* tels us:

$$\begin{aligned} L_{tb} &= \lfloor \log_2(140) / \log_2(2) \rfloor \\ &= \lfloor 7.129/1.0 \rfloor \\ &= 7 \end{aligned}$$

The *statistical trust boundary* defined above is a *hard* boundary. This boundary can be relaxed if one considers that some of the letters in the alphabet may actually be quite rare. $\log_2(|\mathcal{A}|)$ can be replaced with the entropy of the process:

Definition 11 (Soft statistical trust boundary)**.** Let $\overline{S}$ be a sequence of random variables with length N.

$$L_{stb} = \lfloor \log_2(N) / H(\overline{S}) \rfloor$$

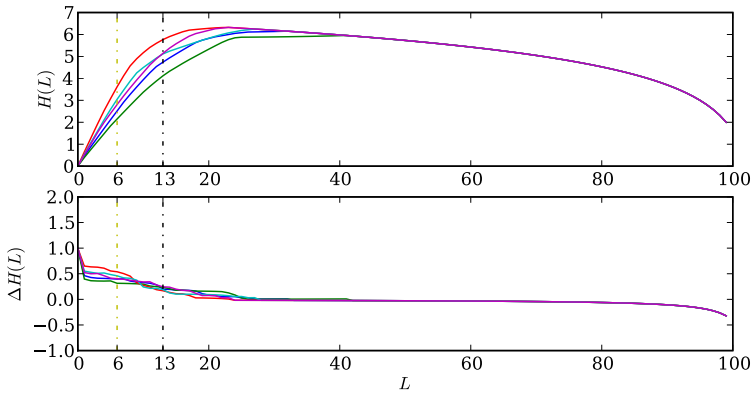

Figure 24: Entropy estimates for five sequences of length $N = 100$, generated from the same random process, with alphabet $\mathcal{A} = \{0, 1\}$ and probability distribution $\mathrm{P}(0) = 0.9$, $\mathrm{P}(1) = 0.1$. The *hard statistical trust boundary* is at 6, the *soft statistical trust boundary* at 13.

Example. Suppose we have a sequence of length 100 generated from a process with alphabet $\mathcal{A} = \{0, 1, 2\}$, and a probability distribution $\mathrm{P}(0) = 0.9$, $\mathrm{P}(1) = 0.05$, $\mathrm{P}(2) = 0.05$. With the given distribution, letters 1 and 2 will together appear only 10% of the time, while letter 0 will appear 90% of the time. It seems reasonable to believe that, in this case, the trust boundary should be larger than 4 ($3^4 = 81$) since the subsequences will be composed mostly of just the letter 0. The *soft statistical trust boundary* of such a sequence is thus 10 (see Figure 25).

Figures 24 and 25 shows typical $H(L)$ curves for a random sequences of finite length. Each figure depicts five curves corresponding to a different trial of the same random process. Both the hard and soft statistical trust boundaries are shown.

F Computational Mechanics

Here we review the basics of computational mechanics needed to understand the stochastic models (in the form of ϵ-machines) computed in this analysis. While information theory is concerned mainly with measuring and encoding information, computational mechanics is concerned with how information is processed and how structure can be quantified.

Some notation: $\overline{S}$ is a sequence of random variables $\overline{S} = \ldots, S_{i-1}, S_i, S_{i+1}, \ldots$. S_i is the i^{th} random variable of the sequence. s_i is the actual value returned by S_i. $\overrightarrow{S_i}^L$ is the sequence $S_i, S_{i+1}, \ldots, S_{i+L-1}$ of L random variables. $\overleftarrow{S_i}^L$ is the sequence $S_{i-L}, \ldots, S_{i-2}, S_{i-1}$ of L random variables immediately preceding S_i.

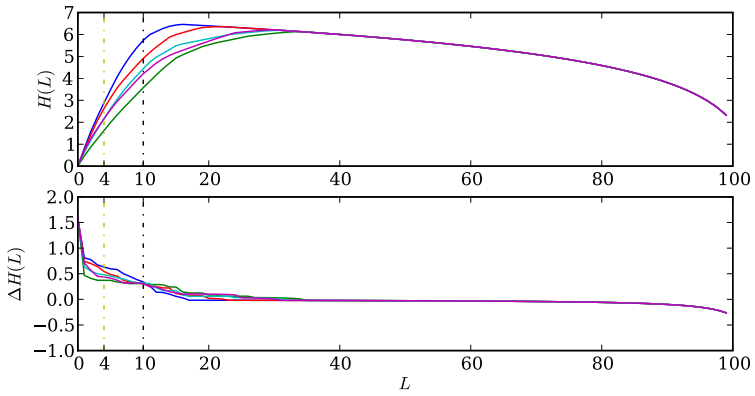

Figure 25: Entropy estimates for five sequences of length $N = 100$, generated from the same random process, with alphabet $\mathcal{A} = \{0, 1, 2\}$ and probability distribution $\mathrm{P}(0) = 0.9$, $\mathrm{P}(1) = 0.05$, $\mathrm{P}(2) = 0.05$. The *hard statistical trust boundary* is at 4, the *soft statistical trust boundary* at 10.

Given a sequence of values $\overline{s} = s_1, s_2, s_3, \ldots$, is it possible to find a pattern that can explain and compactly represent $\overline{s}$? What causes s_2 to follow s_1? What causes s_3 to follow s_1, s_2? Inversely, having seen s_1, s_2, s_3, what will s_4 most likely be? Or, better yet, what is the distribution of futures given s_1, s_2, s_3? The probably obvious but important implication of the statement "having seen..." is the requirement of memory. If we had no memory about the immediate past we would have difficulties predicting the future. Of course, the more we remember the better our estimate of the future will be...one would think. But this is not always true. Sometimes retaining more information about the past does not give us a better prediction about the future. In a simple memoryless random process, such as the toss of a coin, no amount of memorizing past trials will give us a better estimate about the future than a simple guess. So given a particular sequence, how much of it must we remember in order to make the best possible prediction? More specifically, given a particular past sequence $\overleftarrow{s}$ what is the most likely future $\overrightarrow{s}$? Or better yet, what is the probability distribution of futures $\mathrm{P}(\overrightarrow{s}|\overleftarrow{s})$ for an observed sequence $\overleftarrow{s}$? How can we go about finding the minimum memory required for maximal prediction?

From a statistical point of view, we might do the following: Given a sequence $\overleftarrow{s}$ of length N, first find the alphabet $\mathcal{A}$ of the sequence. Then count how many times each letter appears in the sequence. From the estimate of the relative frequencies of the letters in the alphabet we can make a first informed prediction. Obviously we would predict that the most frequent is the most likely to come next. Can we improve our prediction by keeping track of the past subsequences of length $L = 1$, $\overleftarrow{s}^1$? Count how many times each letter in $\mathcal{A}$ is followed by every other letter. If we can better predict

the future with this new calculation , then we continue to $L = 2$ and so on until our prediction can no longer be improved or until we have reached a predefined maximum length limit L_{max}. Note that what we are computing is the probability distribution of the alphabet conditioned on each possible past subsequence $\overleftarrow{s}^L$ of length L. i.e., $\mathrm{P}(a|\overleftarrow{s}^L)$, $\forall\, a \in \mathcal{A}$. In the case of $L = 1$, we have k conditional probability distributions because there are k letters in the alphabet. For $L = 2$ we could potentially have k^2 distinct probability distributions and so on. Some of these distributions, however, could be the same, or practically the same. Thus, it makes sense to group together all subsequences $\overleftarrow{s}^L$ having the same conditional distribution of futures, and to treat them as equivalent given that they have the same "shape" and predictive power. i.e., there is no need to distinguish between different past sequences that give us the exact same information about the future. Doing so would simply add more memory requirements without giving us any added predictive power.[10] At the end of the analysis, having reached an L_{max}, we end up with a collection Σ of sets σ of sequences $\overleftarrow{s}^L$ grouped together because they share the same probability distribution of futures.

What we have just done is effectively partitioned the space of all past subsequences $\overleftarrow{s}^L$ on the basis of their probability distribution of futures. This set Σ gives us an interesting first characterization of the process that might have generated the sequence observed. The next stage is to see how these σ sets relate to each other; i.e., how they connect and how the process jumps from one set to the next.

ϵ-machines

What is an ϵ-machine? An ϵ-machine is a computational model reconstructed from some data stream.

Definition 12 (ϵ-machine)**.** An ϵ-machine of a process is the pair $\{\Sigma, \mathbf{T}\}$, where Σ is the set of *causal states*, and $\mathbf{T}$ is the set of labeled transition matrices $T_{ij}^{(s)}$. ϵ-machines can be visualized as graphs, where vertices (nodes) are the *causal states*, and the edges are the labeled transition probabilities.

Definition 13 (Causal State)**.** The *causal states* of a process are the equivalence classes σ_i induced by the equivalence relation $\sim_\epsilon$ defined as:

$$\overleftarrow{s_i} \sim_\epsilon \overleftarrow{s_j} \text{ iff } \mathrm{P}(\overrightarrow{s}|\overleftarrow{s_i}) = \mathrm{P}(\overrightarrow{s}|\overleftarrow{s_j}) + \delta \quad \forall \overrightarrow{s},$$

where δ is a tolerance. We write σ_i the i^{th} *causal state* Σ the set of all causal states and $\mathcal{S}$ the corresponding random variable.[31]

Each *causal state* has the following attributes:

[31] A couple of equivalent definitions have been given by different authors. See, for example, [20] and [10].

1. An index i, or "name".
2. A set of histories $\{\overleftarrow{s} \in \sigma_i\}$.
3. A conditional distribution over futures: $\mathrm{P}(\overrightarrow{S}|\sigma_i) = \mathrm{P}(\overrightarrow{S}|\overleftarrow{s})$, $\overleftarrow{s} \in \sigma_i$, called the *morph*[8].

The *morph* is the probability distribution of futures of a given *causal state* σ. Subsequences $\overleftarrow{s_i}$ are grouped together in the same *causal state* precisely because they share the same *morph*.

Definition 14 (Causal Transitions)**.** The labeled transition probability $T_{ij}^{(s)}$ is the probability of making the stransition from state σ_i to state σ_j while emitting the symbol (letter) $s \in \mathcal{A}$:

$$T_{ij}^{(s)} \triangleq \mathrm{P}(\mathcal{S}' = \sigma_j, \overrightarrow{S}^1 = s|\mathcal{S} = \sigma_i),$$

where $\mathcal{S}$ is the current *causal state* and $\mathcal{S}'$ its successor. $\mathbf{T}$ is then the set $\{T_{ij}^{(s)} : s \in \mathcal{A}\}$.[20]

ϵ-machine reconstruction

The "intuitive" algorithm described above is a very rough analogue of the first stage in the ϵ-machine reconstruction algorithm proposed by Crutchfield and Young.[6] The actual computation of the ϵ-machines derived in this paper are performed with an algorithm proposed by Shalizi in [20] and [21] and implemented by Klinkner and Shalizi. The ϵ-machine reconstruction algorithm tries to infer the minimal Markovian model capable of generating a given sequence. It begins by assuming that the sequence to be analyzed comes from a simple memoryless random process having a single *causal state*. i.e., it assumes the simplest possible process as point of departure for the reconstruction. Broadly, the algorithm is as follows:[32]

1. *Homogeneity.* Compute the causal states whose member "words" (subsequences) have the same *morph*.
   (a) For each $\hat{\sigma} \in \Sigma$, compute its *morph*:
      i. For each sequence $\overleftarrow{s}^L \in \hat{\sigma}$, compute its *morph*.
      ii. Average the *morphs* of the sequences in $\hat{\sigma}$ to obtain the *morph* of $\hat{\sigma}$.
   (b) For each $\hat{\sigma} \in \Sigma$, test the null hypothesis.[33]
   (c) Find the *causal state* of suffix sequence $a\overleftarrow{s}^L$ with length $L+1$, for each length L subsequence $\overleftarrow{s}^L \in \hat{\sigma}$ and each $a \in \mathcal{A}$.

[32] Please refer to [21] for a detailed definition, explanation and discussion.

[33] All the ϵ-machine reconstructions computed in this paper use the Kolmogorov-Smirnov test to evaluate the goodness-of-fit of the probability distributions derived from the data.

i. Compute the *morph* of $a\overleftarrow{s}^L$.
ii. See if this *morph* matches the *morph* of a previously obtained state.
iii. If it does, add it to the matching state, otherwise create a new state and add the sequence $a\overleftarrow{s}^L$ to it.

2. *Determinization.* We now compute the state transitions to connect the states, and we make sure every member (subsequence) of each state $\hat{\sigma}_i$ has the same *successor state* for the same symbol $a \in \mathcal{A}$. i.e., we determinize them.

   For each state $\hat{\sigma} \in \Sigma$:

   (a) For each $a \in \mathcal{A}$:
      i. For all $\overleftarrow{s} \in \hat{\sigma}$ find the successor state on a. i.e., Find the state $\hat{\sigma}'$ that has the same morph as $\overleftarrow{s}\,a$, for all $\overleftarrow{s}$.
      ii. If there is only one successor state on a, go to the next a.
      iii. If there are $n \geq 2$ successor states on a, create $n-1$ new states and move histories from $\hat{\sigma}$ to the new states so that each state has the same successor on a.

   (b) If every history $\overleftarrow{s} \in \hat{\sigma}$ has the same successor on a, for every a, go to the next state.

References

[1] Roland Auzet. *Roland Auzet: Percussion(s).* Mode, 2008. CD recording.

[2] Mark A. Bee and George M. Klump. Primitive auditory stream segregation: A neurophysiological study in the songbird forebrain. *Journal of Neurophysiology*, 92:1088–1104, 2004.

[3] Pedro Carneiro. *Iannis Xenakis - Psappha, Rebonds A & B, Okho Pour Trois Djembés.* Zig Zag Territoires, 2004. CD recording.

[4] Anne Chao and Tsung-Jen Shen. Nonparametric estimation of Shannon's index of diversity when there are unseen species in sample. *Environmental and Ecological statistics*, 10:429–443, 2003.

[5] Le Corbusier. *Towards a new architecture.* Dover, 1986. Originally published in 1931.

[6] James P. Crutchfield. The calculi of emergence: computation, dynamics, and induction. *Physica D*, 75:11–54, 1994.

[7] James P. Crutchfield and David P. Feldman. Regularities unseen, randomness observed: Levels of entropy convergence, 2001. http://arxiv.org/abs/cond-mat/0102181v1.

[8] James P. Crutchield and Karl Young. Inferring statistical complexity. *Physical review letters*, 63(2):105–108, 1989.

[9] Julio Estrada. *Théorie de la composition : discontinuum continuum.* PhD thesis, l'Université de Strasbourg II, Sciences Humaines, France, 1994.

[10] David Feldman. A brief introduction to: Information theory, excess entropy and computational mechanics. Technical report, Department of Physics, University of California, Berkeley, 2002.

[11] Ellen Rennie Flint. Metabolae, arborescences and the reconstruction of time in Iannis Xenakis' Psappha. *Contemporary Music Review*, 7(2):221–248, 1993.

[12] James Harley. *Xenakis: his life in music.* Routledge, 2004.

[13] Jean Hausser and Korbinian Strimmer. Entropy inference and the James-Stein estimator, with application to nonlinear gene association networks. *Journal of machine learning research*, 10:1469–1484, 2009.

[14] Evan Jones. Residue-class sets in the music of Iannis Xenakis: An analytical algorithm and a general intervallic expression. *Perspectives of New Music*, 39(2):229–261, 2001.

[15] Alexander Kraskov, Harald Stogbauer, Ralph G. Andrzejak, and Peter Grassberger. Hierarchical clustering based on mutual information, 2003. http://arxiv.org/abs/q-bio/0311039.

[16] Jean-Jaques Nattiez. *Musicologie générale et sémiologie*. Christian Bourgois Éditeur, 1987.

[17] Nicolas Ruwet and Mark Everist. Methods of analysis in musicology. *Music Analysis*, 6(1/2):3–36, 1987.

[18] Steven Schick. *Xenakis, Iannis: Complete Percussion Works*. Mode, 2006. CD recording.

[19] William A. Sethares. *Rhythm and transforms*. Springer-Verlag, 2007.

[20] Cosma Rohilla Shalizi. *Causal Architecture, Complexity and self-organization in time series and cellular automata*. PhD thesis, University of Wisconsin at Madison, 2001.

[21] Cosma Rohilla Shalizi. Blind construction of optimal nonlinear recursive predictors for discrete sequences. In Max Chickering and Joseph Halpern, editors, *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence (UAI 2004)*, pages 504–511. AUAI Press, 2004.

[22] Karlheins Stockhausen. ...how time passes... *Die Reihe*, 3, 1957. Translation by Cornelius Cardew in the English edition of Die Reihe (1959), pages 10–40.

[23] Iannis Xenakis. The crisis of serial music. *Gravesaner Blatter*, 1955.

[24] Iannis Xenakis. Towards a metamusic. *Tempo*, 93:2–19, 1970.

[25] Iannis Xenakis. Sieves. *Perspectives of New Music*, 28(1):58–78, 1990.