

XENAKIS' KEREN: A COMPUTATIONAL SEMIOTIC ANALYSIS

Christina Anagnostopoulou
Faculty of Music, School of Philosophy, University of Athens, Greece
chrisa@music.uoa.gr

Chris Share
School of Music and Sonic Arts Research Centre, Queen's University Belfast
cshare01@qub.ac.uk

Darrell Conklin
Department of Computing, City University London
conklin@city.ac.uk

*In Makis Solomos, Anastasia Georgaki, Giorgos Zervos (eds.), *Definitive Proceedings of the International Symposium Iannis Xenakis (Athens, May 2005)*.*

Paper first published in A. Georgaki, M. Solomos (eds.), *International Symposium Iannis Xenakis. Conference Proceedings*, Athens, May 2005, p. 288-298. This paper was selected for the *Definitive Proceedings* by the scientific committee of the symposium: Anne-Sylvie Barthel-Calvet (France), Agostino Di Scipio (Italy), Anastasia Georgaki (Greece), Benoît Gibson (Portugal), James Harley (Canada), Peter Hoffmann (Germany), Miha Iliescu (France), Sharon Kanach (France), Makis Solomos (France), Ronald Squibbs (USA), Giorgos Zervos (Greece)

ABSTRACT

This paper describes a computational semiotic analysis of *Keren*, a piece by Xenakis for trombone solo. A knowledge representation method for pattern representation and discovery, previously applied to a large musical corpus, is here applied to the analysis of a single work. Patterns in *Keren* are found within sequences of musical properties, and a statistical model is used to test the significance of patterns within these sequences. In a second stage, following the process of syntagmatic analysis, a search is made for additional patterns, in which each element of the pattern is a musical segment. This reveals how the various segments are placed together in time to create larger scale semiotic structures, and facilitates the identification of the macrostructure of the piece. It is proposed that this method of analysis, shown previously to be effective for a large corpus, can be equally appropriate for the analysis of a single musical piece. The method is suitable for the analysis of challenging post-tonal pieces, where patterns within the various musical dimensions and their intricate repetitions play a fundamental role in the overall music structure.

1. INTRODUCTION

Keren is a piece for trombone solo, composed in 1986 for performer Benny Sluchin. Like other pieces composed by Xenakis during this period, *Keren* makes full use of the instrument's capabilities and stretches them to their limits. This paper presents an approach to multi-levelled computational analysis of music which explores the properties of structures within a piece, how these properties can be represented in a general and abstract way, and how to find musical patterns that reoccur within the piece in a statistically significant way.

The aim of the present work is not to reveal or examine the compositional processes that Xenakis might have used, implicitly or explicitly, for the piece; rather, the focus is on the *neutral level*, on the music object, the score, which is analysed independently of compositional (poietic level) or perceptual (aesthetic) processes (Molino, 1975). The score, as produced by the composer, has its own existence, and an analysis on the neutral level attempts to reveal internal relationships that exist between structures.

While working on the neutral level, an analyst however has an *a priori* perception and interpretation of the piece and of the analytical method to be followed, and makes choices that cannot truly be considered 'neutral'. Analysts thus work on their own poietic level. Although the intension might be scientific objectivity, to the degree that this is possible given the nature of this work, one also makes related analytical choices; these are indicated and discussed in the paper.

Formal music analysis at the neutral level is concerned with understanding pieces of music by identifying their constituent structures and how these are transformed in time. These decisions are primarily based on the various music properties of a piece, and thus it is crucial to be able to distinguish between these properties in a transparent way.

For this analysis, the *viewpoint* formalism (Conklin, 2006) is used to represent the music knowledge of the piece, its segments and its structure. This formalism allows for the expression of features of music objects, and provides constructors that can be used to build new features from existing ones. The aim is to look for melodic patterns within the musical properties of notes, and repetitions of patterns of segments. More specifically, the analysis can be divided into the following stages:

- a. The melodic representation and discovery level;
- b. The segmental representation and discovery level;
- c. The macrostructure level.

This formalism and methodologies are tested for the first time, both in a single piece as opposed to a homogenous corpus of pieces, and in a contemporary atonal style, as opposed to more traditional tonal works.

From a musicological point of view, our approach is paralleled to semiotic analysis, as developed by Nattiez (1975), where the music score is segmented, and segments are grouped into categories according to their similarity. This similarity, although not explicitly defined, is based on the various musical properties, such as melodic contour, intervals, duration, rhythmic patterns and so on. At a second stage, the piece is viewed as a sequence of paradigmatic class labels.

Music analysis is often concerned with the explication of the organisational principles within an individual piece, and this goal can conflict with the goal of making generalizations beyond the particular analysis piece (Brown and Dempster, 1989). The field of *comparative or systematic* musicology (Cook, 1987), being concerned with a corpus of works, can be productively explored with statistical approaches (Huron, 2001).

Data mining is the field of study in computer science concerned with the discovery of interesting patterns in large databases. A general task in data mining of music, shared with that of systematic musicology, is the discovery of patterns or features that are in some way outstanding in an analysis corpus as related to a comparison set of pieces. Data mining is mainly concerned with inferring *predictive models*: for example, computational models that can be used to classify or group objects. In the application of data mining techniques to an individual piece such as *Keren*, however, we are content with *descriptive* aspects; uncovering and evaluating the repeated patterns within the piece. This slight disconnect (descriptive vs. predictive) with the objectives of data mining has led to some interesting methodological issues which are discussed in detail in Methods.

The rest of the paper is organised as follows: first all the methodological framework is presented, concerning the encoding of the piece to MIDI format and the challenges faced, then moving on to music knowledge representation, discussing the viewpoint formalism and the technique for constructing new viewpoints. The Results section presents some preliminary findings and discusses their potential musical validity.

2. METHODS

2.1. Score Structuring

To represent *Keren*, and to structure (segment) the score in various ways, a data type that permits a hierarchical structuring of a melody is used. A *music object* is a **Note** (with a pitch and a duration), or (recursively) a sequence **Seq**(*X*) of music objects all of the same type *X*. For example, a sequence of notes has type **Seq**(**Note**), and a segmented melody has type **Seq**(**Seq**(**Note**)).

It was convenient to first encode the score in a MIDI format using a professional sequencer (Cubase SX v. 2.2, Steinberg Media Technologies GmbH), which served as an intermediate representation between the printed score and the music object data type just described. MIDI includes information on each note's pitch and delta onset times. The challenges faced when encoding *Keren* were mainly on the pitch encoding; how to represent non-discrete pitches such as glissandi, and how to represent quarter tones, since MIDI does not have an equivalent number for these notes. Information on phrases, breaths and fermatas was also added as text annotations, since they form an integral part of the composition.

2.2 The knowledge representation

A central task of a knowledge representation scheme for music is the computational inference of abstract properties of music objects, thereby grouping objects that may be different at the music surface into the same paradigmatic class. For

example, a knowledge representation scheme might infer that two segments have the same contour shape or melodic density, and can therefore be placed in an equivalence class. The *viewpoints* knowledge representation scheme provides such a framework for computing abstract properties for objects within sequences. These properties are called *viewpoint elements*. The *domain* of a viewpoint is of the form **Seq**(X), where X is a variable that can refer to any type of music object. For example, a viewpoint with domain **Seq**(**Note**) is called a *melodic viewpoint*, and a viewpoint with domain **Seq**(**Seq**(**Note**)) is called a *segmental viewpoint*. A *pattern* is a sequence of viewpoint elements. A pattern *occurs* in a piece if it is contained within the sequence of viewpoint elements of the piece. The *count* of a pattern is its number of occurrences within a piece.

Viewpoint constructors

The viewpoints formalism provides some *primitive* features, such as **pitch**, **duration**, and some basic mathematical operators such as modulo arithmetic. From these primitive viewpoints, *composite* viewpoints are created using functions called *constructors*; these are functions that take viewpoints as arguments, returning new viewpoints. Table 1 provides a list of primitive viewpoints and constructors that have been used in this study of *Keren*.

pitch	primitive pitch viewpoint
duration	primitive duration viewpoint
mod12	primitive modulo 12 function
shape	Huron's (1996) melodic shape function
[interval, V]	subtraction of two V viewpoint elements
[contour, V]	the direction of movement in V viewpoint elements
[lift, V]	constructs a segmental viewpoint from a melodic viewpoint V
[set, V]	the set of V viewpoint elements for a segment
[compose, V, f]	applies the function f to V viewpoint elements
[linked, Va, Vb]	pairs the elements of viewpoints Va and Vb
[new, V]	Boolean measure of viewpoint V change
[selected, V, f]	selects V viewpoint elements if f is true

Table 1: Viewpoints and constructors used for the analysis of *Keren*. Top: primitive viewpoints; Bottom: viewpoint constructors used to create composite viewpoints

2.3 Pattern significance

In systematic musicology studies, it is necessary to identify those features that are in some sense salient and unique to the particular analysis piece or corpus (Huron, 2001). Patterns should reoccur within the corpus, but should not be so general or trivial as to occur with equal probability within unrelated pieces. The attainment of this objective can be assessed using a statistical hypothesis testing framework. Significant patterns are sought within the analysis corpus; those that are over- (or under-) represented in terms of their expected count (here, only over-represented patterns are explored). The expected count of a pattern can be based on its relative frequency within a comparison repertory. For example, in Huron (2001), the analysis corpus is the first movement of Brahms' Opus 51, No. 1 quartet, and the comparison set contains the first movement of Brahms' Opus 51, No. 2, and Opus 67 quartets. Significant deviation

from the expected count can be measured using the chi-square statistic. A requirement of the chi-square significance test is that the comparison repertory is of sufficient size to include the pattern a sufficient number of times, and this implicitly limits the method to short patterns.

Without a large comparison repertory, it is necessary to use alternative procedures to estimate the expected counts of patterns. One way is to design an analytic method to directly model the expected pattern counts (Conklin and Anagnostopoulou, 2001; Apostolico and Crochemore, 2002). For example, Conklin and Anagnostopoulou (2006) noted that the background distribution of counts for a pattern may be approximated by a normal distribution with a variance that depends on the *period* of the pattern (its amount of self-overlap). From this distribution, significance levels as p-values are computed for all patterns, and all patterns meeting a specified significance level are reported to the analyst.

Central to this analytic framework is the computation of the background probability of a pattern, because this probability dictates how many times one expects to see the pattern in an analysis corpus. Here, a type of high-order Markov model is used, constructed from all sub-pattern counts in *Keren*. The construction of the background model from the analysis corpus itself raises the possibility of overestimating pattern probabilities, and thereby not reporting potentially significant patterns. However, the position that this is far less of a problem than reporting too many spurious patterns is adopted here, and in the studies below the significance level has been slightly reduced, in places, to reveal a larger set of patterns.

2.4 Segmentation

There are several simple segmentation points in this piece, which guide the process of manual segmentation of the score into **Seq(Note)** objects. Xenakis has incorporated breath and fermata markings, which naturally break the score into smaller units of diverse lengths. Where breaths and fermatas are scarce, the melodic phrasing indications and the dynamics, both of which are very meticulously notated by the composer, contribute to the segmentation.

There are a small number of breaths and fermatas used for articulation purposes. These tend to be at every single note, for a few consecutive notes (for example, events 212–217, 726–729). These are not interpreted as indicative of segmentation points in our analysis.

The process resulted in 43 segments of significantly different lengths. One should stress here that there can be a number of ways to segment the piece, and each one would yield different results. For this analysis one musically sensible way was chosen, but there are certainly other meaningful segmentations.

2.5 Pattern discovery

To find all significant patterns in *Keren*, for a specified score structuring and viewpoint of the appropriate type, the viewpoint sequence is computed, and all patterns occurring more than once are found using an efficient algorithm. Patterns not meeting the significance level are discarded. This process can still produce a large set of patterns, and

therefore a further filtering is applied to the set of all significant patterns in several applications with real music, we have found that the extrema of chains of sub-patterns are of the most interest. Therefore, only the shortest significant (no other significant pattern is a subpattern) and the longest significant (sub-patterns to no other significant pattern) are reported.

3. ANALYSIS AND RESULTS

Using these methods an analysis layout is created which has three stages: the melodic, the segmental, and the macrostructural. In the following, these stages are discussed in detail, presenting a selection of the results. It should be noted that there are many representation possibilities of the musical surface using constructed viewpoints, and each viewpoint will produce different patterns. This paper reports only on a small number of viewpoints and patterns which appear to have the most apparent musical interest (independent of compositional processes), as well as some results that demonstrate the viewpoint representation and the abstraction which is made possible by using the formalism.

3.1. Melodic level

Paradigmatic analysis, the first part of a semiotic analysis, would segment the score according to repetition, and classify the resulting segments according to similarity. While the concept and rationale here is similar, there are two main differences. First, the score is not segmented, but with the pattern discovery algorithm shortest and longest significant patterns (or segments) are revealed. Second, similarity between segments is explicitly defined, as class equivalence segments share the same viewpoint patterns.

The melodic viewpoints considered ranged from basic ones, such as **pitch** and **duration**, to constructed and gradually more abstract ones, such as pitch classes, melodic intervals, intervals of pitch classes **mod12** and viewpoints using the constructors **linked**, **new** and **selected**. To illustrate some melodic viewpoints, Table 2 shows an example of the opening four-note figure of *Keren*, and how this is translated to a number of melodic viewpoint sequences.

The following results will illustrate some shortest and longest significant patterns when using the pattern discovery algorithm with a p-value threshold of 0.01.


				
Pitch	60	56	55	56
[compose, pitch, mod12]	0	8	7	8
[interval, pitch]	undef	-4	-1	1
[compose, [interval, pitch], mod12]	undef	8	11	1
[contour, pitch]	undef	-	-	+
[contour, duration]	undef	=	=	+

Table 2: The first four notes of *Keren*, comprising the opening motif, translated into a number of melodic viewpoint sequences. Interval and contour viewpoints are undefined for the first element in the sequence. Score fragment reproduction from Editions Salabert, Paris, 1989.

3.1.a pitch

This basic viewpoint returns the pitch of an event. We observed that all reported significant patterns made use of only very specific pitches:

- a large number of shortest and longest significant patterns, made exclusive use of the pitches Fs4 (MIDI number 66), G4(67), B4(71);
- a number of patterns, in addition to the above, also used Gs3(56), A3(57), B3(59), Cs4(61), D4(62);
- a few more patterns additionally made use of D3(50), G3(55), As3(58), C4(60).

The large number of the patterns that consist of various combinations of Fs4, G4, and B4 are all found in the section of event numbers 430-645. This is a section comprising fast, isochronous events of these pitches, and will henceforth be referred to as section B.

One of the patterns reported was the opening motif of the piece: C4, Gs3, Gn3, which is encountered five times in the piece. Table 2 shows the opening motif of the piece, containing these three notes. In the following, where more abstract viewpoints are investigated, more occurrences are found that are equivalent to this motif.

What these results convey is not that these pitches are important on their own (though they might be), but that their co-occurrences into patterns are significant. No other pitch patterns were reported.

3.1.b [compose, pitch, mod12]

This viewpoint is more abstract than **pitch** in that it checks for patterns of pitch classes, disregarding octave information (**mod12**). This viewpoint therefore had slightly more general results than the **pitch** viewpoint. A number of patterns

were reported, the most interesting being a large group of shortest and longest significant patterns which made exclusive use of the pitch classes 6,7 and 11 (Fs, G, B) in various combinations. This is in accordance with the results of the **pitch** viewpoint above. These patterns were also found in the same distinct section B of the piece.

3.1.c [interval, pitch]

This viewpoint looks at patterns of pitch intervals in numbers of semitones. The patterns reported for both shortest and longest significant patterns made use of the intervals ± 1 , ± 4 , ± 5 , and to a much lesser extent ± 6 and ± 2 . It is very interesting that no other intervals were reported to be part of the significant results: The intervals ± 7 , ± 8 , ± 9 and above did not appear at all in the set of significant patterns.

The majority of the interval patterns, when realised to actual instances of pitches, in fact use the same three pitches. One realisation could be the pitches Fs, G and B, although this is not always the case in the score because this viewpoint allows for transposition. However, it is interesting to note that Xenakis realised most of the interval patterns to the same pattern of notes in each case. Examples of this in shortest significant patterns, with their count in *Keren*, are:

[4,-5] (34),
 [1,4] (31),
 [5,-4,4] (18),
 [-5,5,4] (14),
 [-4,-1,5] (10),
 [4,-4,-1,1] (5),

and a few examples of longest significant patterns are:

[1,4,-5,5,-4,4,-5,5] (2),
 [-5,1,4,-4,4,-5,5,-4,4,-5] (2),
 [1,4,-4,-1,5,-5,1,4,-4,-1,1,4,-5,1,4,-4,-1,5,-5,1,4,-4] (2),

and several others. Most patterns instances are distributed evenly throughout the piece. The longest patterns, however, tend to appear in the same distinct fast section B that was mentioned above in relation to the previous viewpoints.

The results also contained the following interesting **[interval, pitch]** patterns: a. [1,1,1,1] (count 3) which is within the glissando passage of the piece, and b. patterns with a downward trend: [-1,-4,-1,-2] (count 4), [-4,-1,-2,-2,-1] (count 3), and others.

The initial motif [C,Af,G] which has been discussed in relation to the previous viewpoints (Table 2), creates the interval pattern [-4,-1], and it is also found here as part of longer patterns (for example, the pattern [-4,-1,-2,-2,-1] just mentioned).

3.1.d [compose, [interval, pitch], mod12]

This viewpoint looks at interval patterns between pitch classes in number of semitones. The intervals within the resulting patterns are mainly: 1,4,5,7,8,11. These are symmetric to the unison, and they show two symmetric structures: 1,4,5 and 7,8,11. The shorter versions of these patterns are found everywhere in the piece, whereas the longer versions, as expected, in the fast passage section B (events 430-645).

Another interesting observation here is that the intervals 3 and 9 do not appear anywhere in the significant pattern set. If we take a possible realisation of these two intervals, starting from pitch class C, we have:

3: C-to-Ds

9: C-to-A.

The interval Ds-to-A is the tritone. This is interesting because Xenakis, as reported in the interval results above, tends to avoid the intervals 2 and 6 (tritone) in this piece. Furthermore, the three intervals, 3, 6 and 9 are omitted completely from the isochronous fast section B.

Another interesting result is that patterns beginning with the two intervals **0,0** appear only in the first and third sections of *Keren*.

3.1.e [contour, pitch]

This viewpoint describes melodic contour. Oscillating textures — patterns with interchanges of a melodic contour of up and down motion at every viewpoint element — were very prominent in the results and were found throughout the piece.

Patterns with exclusive downward trend also appeared: [-,-,-,-,-] (count 32) and [-,-,-,-,-,-] (count 20) were part of shortest and longest reported patterns respectively. This is in accordance with the results on intervals above: because contour is more abstract representation than intervals, there are more occurrences in this case. Again here we notice a distinction with section B: these patterns are used throughout the piece apart from that specific fast section.

There were no patterns with two or more consecutive upward motions (contiguous + elements), and only very few with same motion (contiguous = elements).

3.1.f [contour, duration]

Analogous to pitch contour, duration contour involves the comparison between each rhythmic value with its predecessor.

This viewpoint pointed out the two fast isochronous sections in the piece, events 223-388 and 430-645 (which we call B). These were, as expected, instances of the long pattern [=,=,=,...]. However, a very long pattern [=,=,...], with more

than 70 elements, occurs 17 times in an overlapping way only in the second fast section. The shortest significant pattern [=,=] appears 553 times throughout the piece.

Some other tested viewpoints were:

[linked, pitch, duration]

[linked [interval pitch] duration]]

[linked, [compose [interval, pitch], mod12], duration]

These viewpoints link constructed viewpoints based on **pitch** with **duration**. In the first one, all patterns reported are from the fast section B mentioned above. In the second one, there is only one occurrence, in one of the patterns, which is outside that B section, while in the third viewpoint, there are 4 extra occurrences pre-B, and 2 post-B (the rest being found at section B).

3.2 Segmental level

As mentioned above, the first step of a semiotic analysis would be the score segmentation and the categorisation of the segments. The second step, which is known as *syntagmatic analysis*, would involve the order in which the segments appear.

For the segmental level, the score was manually segmented as described in Section 2.4, using score indications for segment boundaries. The algorithm then looked for segmental patterns; patterns whose elements are not single notes, but segments of notes.

Several segmental viewpoints were constructed and considered. The viewpoints reported here are: the set of pitch classes, the melodic shape as this was introduced by Huron (1996), contour of density, contour of duration, and lifted selected contour at new contour points. Table 3 shows an example of the first two segments of *Keren*, and how this is translated to a number of segmental viewpoint sequences.


		
[set, [compose, pitch, mod12]]	[0,7,8]	[0,1,7,8,10]
shape	concave	convex
[contour, density]	undef	+
[lift,[selected,[contour,pitch], [new,[contour,pitch]]]]	[-,+]	[+,-]
[contour, duration]	undef	+

Table 3: The first two segments of *Keren*, translated into segmental viewpoint sequences. Density contour and duration contour viewpoints are undefined for the first segment in the score. Score fragment reproduction from *Keren*, Editions Salabert, Paris, 1989.

The p-value threshold for all segmental pattern experiments was ignored; any pattern that appeared twice or more was reported. This was because the piece, when segmented, is short and there were not enough segments to justify a significance threshold. Below the most interesting patterns taken out of the segmental stage are reported:

3.2.a [set, [compose, pitch, mod12]]

This viewpoint shows the set of pitch classes encountered in a segment. The only repeated patterns found here were:

pattern **[[0,1,5]]** occurred at events 6 and 42;

pattern **[[10,11]]** occurred at events 8 and 16;

pattern **[[7,8]]** occurred at events 13 and 35.

This viewpoint shows interesting musical relations, where pitches in a segment are shared, but segments can differ in any other property: the order and the repetition of the pitches in the segment, the register, the length of the segment, rhythm, and so forth. All results obtained with this viewpoint were of length 1: no longer pattern appeared more than once in *Keren*.



Figure 1: Two instances of the segmental pattern **[[0,1,5]]** in *Keren*. Top: segment 6; Bottom: segment 42. Score fragment reproduction from *Keren*, Editions Salabert, Paris, 1989.

Figure 1 illustrates the two instances of the pattern **[[0,1,5]]**: segments 6 and 42. In Table 4, a representation of the whole piece in terms of segment numbers and their pitch class sets is displayed.

3.2.b shape

This viewpoint describes the overall melodic shape of a segment, as defined by Huron (1996). A set of rules relating the first and last pitches in a segment to the mean of the internal pitches leads to a segment being in one of nine shape classes. The most common pattern that is met is **[descending]**, with 16 occurrences. This result is in line with the melodic contour viewpoint examined above, where there was a number of descending patterns found. The one-element patterns **[convex]** and **[concave]** were also important, with 10 and 8 occurrences respectively. Table 4 shows a classification of all segments according to melodic shape. The longest pattern that was reported was **[descending, concave, convex, descending, convex]**, occurring at segments 4 and 26.

It is important to note here that these results are based on a particular chosen segmentation, and a different segmentation might have resulted in different results altogether.

<i>Segment</i> <i>t</i>	<i>Pitch Class Set Viewpoint</i>	<i>Shape Viewpoint</i>
1	[0,7,8]	concave
2	[0,1,7,8,10]	convex
3	[0,1,7,8]	descending
4	[1,2,4,6,7,8,9,10,11]	descending
5	[10]	concave
6	[0,1,5]	convex
7	[1,2,4,6,7,8,9]	descending
8	[10,11]	convex
9	[0,10,11]	descending
10	[4,5,9,10]	ascending
11	[0,5,6,11]	descending
12	[1,2,5,7,8,9,10]	concave
13	[7,8]	convex
14	[0,1,5,6,7,8,9,11]	convex
15	[0,11]	convex
16	[10,11]	ascending
17	[0,2,4,5,6]	descending
18	[10,1,2,5,6,7,8,9,10,11]	concave
19	[0,1,2,5,6,7,8,9]	concave
20	[0,6,7,8,11]	descending
21	[4,5,6,7,10]	descending
22	[1,2,8,9,11]	concave
23	[0,1,5,6,7,8,10,11]	convex
24	[0,1,2,4,6,7,8,9]	descending
25	[4,6,10]	ascending
26	[0,1,6,11]	descending
27	[0,1,8,9]	concave
28	[0,6,7,11]	convex
29	[0,6,7,9,11]	descending
30	[3,5,6,7,11]	convex
31	[0,1,6,7,8,9]	concave
32	[0,1,2,4,5,6,7,8,10]	descending
33	[4,5,6,7,8,9,10,11]	ascending
34	[0,1,2,4,5,7,8,9,10,11]	ascending
35	[7,8]	ascending
36	[0,7,9]	ascending
37	[0,3,4,5,9,10,11]	ascending
38	[2,5,7]	descending
39	[0,4,5,6,7,9,11]	concave
40	[5,6,7,11]	horizdes
41	[0,1,2,4,5,6,7,9,11]	descending
42	[0,1,5]	convex
43	[4,5,6,7,8,9,10]	descending

Table 4: A syntagmatic analysis of all segments in *Keren*, in terms of two segmental viewpoints: set of all pitch classes and shape.

3.2.c [contour, density]

This viewpoint compares the density (number of notes divided by length of segment) between successive segments. One of the patterns discovered here is [+ ,+ ,+], which appears at locations 26 and 27. This means that the longer pattern [+ ,+ ,+ ,+] occurs only once in the piece, in segments 25–29. This is interesting because this building of density appears at the place leading towards and beginning of the fast section B, which featured in a number of other results in the melodic points above.

3.2.d [lift, [selected, [contour, pitch], [new, [contour, pitch]]]]

This viewpoint shows the segment sequence (lifted) of the contour value when there is a change of contour. There were 15 repeating segments were found; One of them was [[+ ,+ ,+ ,+], [-]], which has 2 occurrences, at segments 3-4 and 17-18.

Figure 2 demonstrates this pattern and its instances. The two instances have an interesting similarity between them, the first pattern in each being the characteristic melodic pattern which makes use of the intervals ± 1 , ± 4 , ± 5 , and the second being this downward characteristic movement, also seen in the melodic viewpoint results above.




Figure 2: Two instances of the segmental pattern $[[+, -, +, -], [-]]$ in *Keren*. Top: segments 3 and 4; Bottom: segments 17 and 18. Score fragment reproduction from *Keren*, Editions Salabert, Paris, 1989.

3.3 Macrostructure

The third level of the analysis, the macrostructure, is indicated by the results in the other two levels, the melodic and the segmental. These results seem to suggest that there is one section in the piece where the music material is transformed: certain features or patterns disappear, while others become more prominent and intensify. This section, which has been named B above, has the following features:

Pitch patterns make use of pitches Fs4, G4 and B4, as opposed to the rest of the piece. The same goes for pitch class patterns. In terms of intervals, the longest significant patterns that make use of the intervals ± 1 , ± 4 , ± 5 appear again in this section. Rhythmic values are equal and very fast. Finally, contour duration patterns single out this section, as do the linked pitch with duration longest significant patterns. There are no breath marks or fermatas to break the intensity, so the whole section is represented as one segment in the segmental level.

The distinction of this section from the rest of the piece results in a loose ABC, where:

- section A**, between events 1–429, where all the compositional material is first presented;
- section B**, between events 430–645, where there is a climax, with very specific persistent patterns;
- section C**, between events 646–788, where the piece winds down.

At the same time, while most of the results differentiate between section B and the rest of the piece, some of the more abstract patterns are found throughout the piece, and thus contribute towards the cohesion of the piece. These include some of the shortest and longest significant pitch patterns, some of the shortest pitch class patterns, some of the shortest significant pitch and pitch class interval patterns, and all of the melodic contour patterns.

It is worth noting that in section A we can also observe a subsection between events 223–388, which has at a first glance a lot in common with section B: fast isochronous rhythmic values, with some repetition of pitches, and no breath marks. However, most of the viewpoints presented in this paper showed that this section was not as interesting as section B in terms of patterns and their significance.

4. DISCUSSION

This paper has presented a computational analysis of *Keren* based on representation of music properties and discovery of interesting patterns. As a first step, melodic patterns are found within the various musical dimensions, starting from the basic ones, pitch and duration, and gradually progressing to more abstract ones such as pitch classes, pitch intervals, melodic contour, contour of duration, linked intervals with duration. As a second step, the piece is segmented following the composer's indications for breaks, and various segmental patterns are discussed.

A characteristic of the music of Xenakis is that the unit of composition is not the note, or even the simple melody, but rather much more complex structural units, which can differ from piece to piece (Papaioannou, 1994). The viewpoint formalism and the pattern discovery method we use supports the discovery of some of these structures and complexities, since the level of representation can be tuned at any level of abstraction, taking as basic information pitch and duration.

The manual segmentation in the segmental stage above, although following the composer's indications, had the following problems: first, segments were too unequal in length to be able to demonstrate repetition in the segmental level. For example, section B all formed part of a single segment. As a result, the findings from the segmental level were not as interesting as the ones from the melodic level. At the same time, we observed that some of the melodic patterns we discovered spanned across segments. This means that the composer himself did not follow the break points of the score to enclose his compositional ideas.

For Xenakis, the concept of symmetry has been a very important one in his work. With the theory of *sieves* he shows one way of formalising symmetry: sieves are symmetries in 'any set of characteristics of sound or of well-ordered structures, and especially to any group which entails an additive operation and whose elements are multiples of a unity' (Xenakis, 1992, p. 268).

In *Keren*, Xenakis has been very selective with his compositional material; not so much of specific pitches or classes of pitches, although these are salient too, but of intervals and interval structures. The interval structure of 1,4,5, and its symmetric structure of 7,8,11 are very prominent in this work. In the section that we have called B, this structure is all we hear, which in fact is specified to certain pitches. All permutations and combinations, using different articulation accents, are heard in this section. In the rest of the piece those interval structures are also very noticeable, but are diluted with other compositional material.

The very characteristic opening motif of the piece is a statement of this interval structure, downwards and upwards. Another characteristic pattern used in this piece, the downward movement, also makes use of this interval structure (for

example, events 17–25, 91–98, etc.). The sequence of intervals 1,4,5,7,8,11 could be thought of as a sieve that comprises the union of two residual classes, 3.1 (intervals 1,4,7) and 3.2 (intervals 5, 8,11).

5. CONCLUSIONS AND FURTHER WORK

The results presented in this work are still very preliminary: we only considered a small number of viewpoints, and out of those we selected a small number of results. We have demonstrated that there are interesting music relations within the piece, especially related to pitch classes and intervals, which point out a loose structure of ABC. At the same time, while compositional intentions were not considered, this paper has demonstrated that this generalised method of representing music and discovering patterns might be a suitable method for analysing single atonal pieces.

A feature of this work, discussed in Methods, is that the piece *Keren* itself was used to construct an analytic background model for evaluation of the statistical significance of patterns. In the future we would like to use a wider set of Xenakis pieces to create background models, and also to do inter-opus pattern discovery within each piece.

Another problematic issue, as mentioned above, has been segmentation. Although we segmented the piece at places indicated by the composer as breaks, it was obvious that this segmentation did not yield the best results in the segmental patterns. This suggests that the compositional material had been distributed across segment boundaries. A more reliable way for the segmentation, given the compositional style of Xenakis, might have been to use the longest significant patterns of certain viewpoints to denote segment boundaries. Segments in *Keren* seem to be *textural*, and require a special approach to auto-segmentation: we are currently exploring some melodic viewpoints that might reveal changes in texture.

Further work also includes the following aspects: a) the use of dynamics in the experiments, since these are so carefully notated by the composer, and it is obvious that their role in the overall structure of the piece should not be neglected; b) the use of sieves in Xenakis' music might not only stop at pitches and intervals, and aim to investigate relations between values and patterns in various viewpoints, especially related to durations, dynamics and musical textures; c) the macro-structure revealed with our results here could be further explored by studying the musical material in sections A and C in more depth.

6. REFERENCES

- Apostolico, A. and Crochemore, M. (2002). String pattern matching for a deluge survival kit. In Abello, J., Pardalos, P., and Mauricio, G., editors, *Handbook of Massive Data Sets*, pages 151-194. Kluwer Academic Publishers, Boston, MA.
- Ariza, C. (2004). An Object-Oriented model for the Xenakis Sieve for algorithmic pitch, rhythm and parameter generation. In *Proceedings of the International Computer Music Conference*, pages 63-70, Miami.
- Conklin, D. (2006). Melodic analysis with segment classes. *Machine Learning*, to appear.

- Conklin, D. and Anagnostopoulou, C. (2001). Representation and discovery of multiple viewpoint patterns. In *Proceedings of the International Computer Music Conference*, pages 479-485, Havana, Cuba.
- Conklin, D. and Anagnostopoulou, C. (2006). Segmental pattern discovery in music. *INFORMS Journal on Computing*, 18(3):285-293.
- Cook, N. (1987). *A Guide to Musical Analysis*. Oxford University Press.
- Huron, D. (1996). The melodic arch in Western folksongs. *Computing in Musicology*, 10:3-23.
- Huron, D. (2001). What is a musical feature? Forte's analysis of Brahms's Opus 51, No. 1, revisited. *Music Theory Online*, 7(4).
- Molino, J. (1975). Fait musical et sémiologie de la musique. *Musique en Jeu*, 17:37-61.
- Nattiez, J. (1975). *Fondements d'une sémiologie de la musique*. Union générale d'éditions, Paris.
- Xenakis, I. (1986). *Keren pour trombone solo*. Editions Salabert, Paris, 1989.
- Xenakis, I. (1990). Sieves. English translation by J. Rahn, *Perspectives of New Music*, 28(1):58-78.
- Xenakis, I. (1992). *Formalized Music*, Revised Edition Pendragon Press, NY.
- Papaioannou, I. (1994). Iannis Xenakis, in National Technical University of Athens, ed., *Iannis Xenakis, Ena afieroma tou Ethnikou Metsoviou Politehniou pros enan apofito tou*. Sigxroni Epoxi, Athens.

Acknowledgments: Special thanks to Makis Solomos and Jason Dixon.