

A Comparison of Rhythmic Similarity Measures

Godfried Toussaint*
School of Computer Science
McGill University
Montréal, Québec, Canada

August 18, 2004

Technical Report SOCS-TR-2004.6, August 2004.

Abstract

Measuring the similarity between rhythms is a fundamental problem in computational music theory, with many applications such as music information retrieval and copyright infringement resolution. A common way to represent a rhythm is as a binary sequence where a zero denotes a rest (silence) and a one represents a beat or note onset. This paper compares various measures of rhythm similarity including the Hamming distance, the Euclidean interval-vector distance, the interval-difference distance measure of Coyle and Shmulevich, the swap distance, and the chronotonic distance measures of Gustafson and Hofmann-Engl. Traditionally, rhythmic similarity measures are compared according to how well rhythms may be recognized with them, how efficiently they can be retrieved from a data base, or how well they model human perception and cognition of rhythms. In contrast, here similarity measures are compared on the basis of how much insight they provide about the structural inter-relationships that exist within families of rhythms, when phylogenetic trees and graphs are computed from the distance matrices determined by these similarity measures. For two collections of rhythms, namely the 4/4 time and 12/8 time clave-bell time lines used in traditional African and Afro-American music, the chronotonic and swap distances appear to be superior to the other measures, and each has its own attractive features. The similarity measures are also compared according to their computational complexity.

1 Introduction

Imagine that you are at a concert, after a stunning performance, clapping at a fast uniform pace, much like a heart beat while jogging. Even better, try it out right now, but stop after you reach sixteen claps. Then do it again but this time execute the first, fourth, seventh,

*This research was partially supported by NSERC and FCAR. e-mail: godfried@cs.mcgill.ca





Figure 1: Eight common ways of representing the *clave Son* rhythm.

eleventh and thirteenth claps loudly, and the remaining eleven claps softly. Your clapping pattern could then be represented like this: [1 2 3 4 5 6 **7** 8 9 10 **11** 12 **13** 14 15 16], where the claps shown in bold face are the ones you clap loudly. It may also be helpful to count aloud the sixteen claps in groups as follows: [1 2 3 1 2 3 1 2 3 4 1 2 1 2 3 4], and to clap loudly only on the one's. If you repeat this pattern over and over, making the soft claps completely inaudible, you will be clapping one of the most popular rhythms on the planet, known mainly as the *Clave Son* from Cuba [41].

The *Clave Son* rhythm is usually notated for musicians using standard music notation which affords many ways of expressing a rhythm. Four such examples are given in the top four lines of Figure 1. The fourth line displays the rhythm with music notation using the smallest convenient durations of notes and rests. Western music notation is not ideally suited to represent African music [18], [4]. The fifth and sixth lines show two popular ways of representing rhythms that avoid Western notation. The representation on line five is called the *Box Notation Method* developed by Philip Harland at the University of California in Los Angeles in 1962 and is also known as TUBS (Time Unit Box System). The TUBS representation is popular among ethnomusicologists [18] and invaluable to percussionists not familiar with Western notation. It is also convenient for experiments in the psychology of rhythm perception, where a common variant of this method is simply to use one symbol for the beat and another for the pause [17], as illustrated in line six. In physiology, where the study of cardiac rhythms [6] is important, as well as in computer science the *clave Son* would be written as the 16-bit binary sequence shown on line seven. Finally, line eight depicts the interval length representation of the *clave Son*, where the numbers denote the lengths (in shortest convenient units) of the intervals between consecutive beats (onsets). The compactness and ease of use in text, of this numerical interval-length representation, are two obvious advantages, but its iconic value is minimal.



Clave Son

Figure 2: The clave *Son* represented in 13th Century Arabic cyclic notation.

2 Geometric Representation of Rhythms

Because traditional Western music notation is not well suited for the visual representation of rhythm, scholars of rhythm have sought to develop graphical geometrical representations that afford not only easier manipulation and visualization of rhythms, but that are applicable to domains other than music. However, geometric representations of rhythms have existed for quite some time even before Western music notation was well developed. For example, an Arabic book about rhythm written by Safi-al-Din in 1252 depicts rhythms using a circle divided into “pie slices” [45]. The *Son* represented in this notation is depicted in Figure 2, where the pie slices that correspond to beats are shown shaded, and an arrow indicates the starting note and the direction of the flow of time.

In 1987 Kjell Gustafson, at the Phonetics Laboratory of the University of Oxford, published an original method to represent a rhythm as a two-dimensional graph [21], [22]. His idea (first proposed in 1983) is best explained with an example. Consider first the clave *Son* pattern shown on line five of Figure 1. The interval lengths (3, 3, 4, 2, 4) between the onsets of the beats are ordered along the x -axis as a one-dimensional pattern. Although this representation is realistic in terms of the time at which beats occur, the relative durations of the intervals are not easily observed. In a *histogram* approach to rhythm visualization the intervals are plotted along the y -axis [22], resulting in the *adjacent-interval-spectrum* of the rhythm. Figure 3 depicts the clave *Son* with such a representation. Here the relative lengths of the intervals are clearly visible but the temporal information along the x -axis is lost. To obtain a graphical representation that possesses the advantages of both of these methods, Gustafson simply combines them. The result of this union is illustrated with the clave *Son* in Figure 4. Each temporal element (interval) is now a box and both the x and y axes represent the length of time of the interval. Gustafson refers to such a display as TEDAS (*Temporal Elements Displayed As Squares*). A TEDAS representation of rhythm has other advantages besides those mentioned in the preceding. Some of these will be discussed in the section on measuring similarity between rhythms. For additional advantages the reader is referred to [22].

Unaware of the work of Gustafson [22], Hofmann-Engl [24] in 2002 rediscovered the idea of representing rhythms as a two-dimensional display where both axes measure time. Whereas

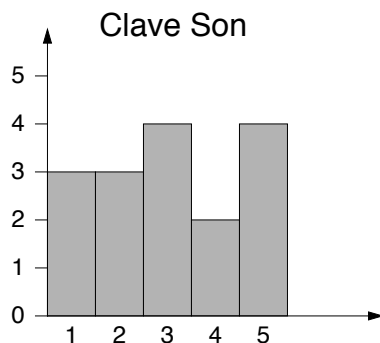


Figure 3: The *adjacent-interval-spectrum* of the clave *Son*.




Figure 4: The TEDAS representation of the clave *Son*.

Gustafson was motivated by visualizability and technical considerations, Hofmann-Engl, on the other hand, was concerned with predicting human perception and cognition of rhythmic similarity. Although the language used by Hofmann-Engl is quite different from that employed by Gustafson, their representations are equivalent. As in the preceding, the idea is illustrated here with the clave *Son* as an example.

First the rhythm is broken down into its smallest time units (16th notes and rests in this case) as in line four of Figure 1. Hofmann-Engl calls these *atomic beats*. Then each atomic beat is assigned a y -coordinate (height) equal to the length of the beat-to-beat interval in which it falls, resulting in a pattern of sixteen points in the plane. Finally these points are connected together in increasing x -coordinate to obtain an x -monotonic polygonal chain. Hofmann-Engl calls such polygonal curves *chronotonic chains*. Figure 5 shows the clave *Son* represented as a chronotonic chain.

A third geometric method of representing a rhythm is illustrated in Figure 6 where the clave *Son* is depicted as a convex polygon. Like the chronotonic representations of Gustafson and Hofmann-Engl, the polygon representation clearly shows the relative lengths of the onset or beat intervals, and is realistic on the circular time scale. In fact this representation is even more realistic than its chronotonic counterparts because rhythm timelines are cyclic, and all the other representations discussed above, apart from the Arabic “pie-slice” representation, have the beginning and end of the pattern located physically distant from each other, when they actually coincide cyclically in time. This representation also yields a powerful visual-



Figure 5: The clave *Son* as a *chronotonic chain*.




Figure 6: The clave *Son* represented as a convex polygon.

ization tool, since a variety of geometric properties of rhythms are immediately perceived from the resulting polygons. For example, the dashed line connecting locations 3 and 11 in Figure 6 indicates an axis of mirror symmetry for the clave *Son*. The polygon can also be used to compute a variety of n -dimensional vectors of geometric features $X = (x_1, x_2, \dots, x_n)$. For example, focusing on the vertices of the *Son* polygon $V = (v_1, v_2, v_3, v_4, v_5)$ one could take v_i to be the internal angle of the i th vertex. On the other hand viewing the polygon as a sequence of edges $E = (e_1, e_2, e_3, e_4, e_5)$ one could take e_i to be the length of the i th edge of the polygon. In addition, one could use a variety of *global* shape features of the polygon such as the ratio of perimeter to area, or various moments of inertia [37]. Recently, geometric features such as these have been used to obtain new classifications of rhythm timelines in African and Afro-American traditional music [41], [43]. Such classifications fall under the more recent *structural* approach proposed by Jay Rahn [35], [34], and contrast sharply with earlier *statistical* methods (for example, in 1966 Abraham Moles [29], classified rhythms based on the entropy function of the histogram of relative frequencies of occurrences of their inter-onset intervals).

3 Measuring the Similarity of Rhythms

At the heart of any algorithm for comparing, recognizing or classifying rhythms, lies a measure of the similarity between a pair of rhythms. The type of similarity measure chosen is in part predetermined by the manner in which the rhythm is represented. Furthermore, the design of a measure of similarity is guided by at least two fundamental ideas: *what* should be measured, and *how* should it be measured.

There exists a wide variety of methods for measuring the similarity of two rhythms represented by strings of symbols [42]. Indeed the resulting approximate pattern matching problem is a classical problem in pattern recognition and computer science in general [16]. Traditionally similarity between two patterns is measured by a simple template matching operation. More recently similarity has been measured with more powerful and complex functions such as the *earth mover's distance* [44], [11].

In this section we describe the five, easy to compute, popular and representative measures of rhythm similarity (or dissimilarity) that will be compared in the following.

3.1 The Hamming distance

For binary sequences a natural measure of distance or dissimilarity is the Hamming distance [23] widely used in coding theory. The Hamming distance is simply the number of places in the strings where elements do not match. When a measure of similarity is preferred some researchers use the number of matches instead [28]. Hence it is one of the simplest kinds of template matching [16]. Note that in this approach each rhythm is represented by a vector $X = (x_1, x_2, \dots, x_n)$ where the x_i represent binary valued features of the rhythm. If a note is played at time unit i then $x_i = 1$ and otherwise $x_i = 0$. Thus rhythms are represented as points in a n -dimensional binary-valued vector space (usually called a *hypercube*). The Hamming distance between two points $X = (x_1, x_2, \dots, x_n)$ and $Y = (y_1, y_2, \dots, y_n)$ in this space is given by:

$$d_H(X, Y) = \sum_{i=1}^n |x_i - y_i| \quad (1)$$

where $|x|$ denotes the absolute value of x .

The Hamming distance is easily computed in $O(n)$ time. However, it is not very appropriate for our problem of rhythm similarity, when used with this particular type of representation, because although it measures the existence of a mismatch, it does not measure how far the mismatch occurs. Furthermore, if a note or beat is moved a large distance the resulting modified rhythm will sound more different than if the beat is moved a small distance.

3.2 The Euclidean interval vector distance

Some rhythm detection algorithms [31] and systems for machine recognition of music patterns [12] use *inter-onset intervals* as the basic features for measuring similarity. These are the intervals of time between consecutive note onsets in a rhythm. This approach and its variants are more appropriate than the Hamming distance for measuring the similarity of rhythms. Each rhythm may be represented by a vector of numbers that characterize these intervals. More specifically a rhythm may be represented by $X = (x_1, x_2, \dots, x_n)$, where x_i is the number

of vertices skipped by the i th polygon edge starting at vertex labelled 0. This is essentially the same as the sequence of inter-onset time intervals since the time interval is the number of vertices skipped plus one. The dissimilarity between two rhythms $X = (x_1, x_2, \dots, x_n)$ and $Y = (y_1, y_2, \dots, y_n)$ may then be measured by the Euclidean distance (or some other metric) between the two interval vectors X and Y , in this n -dimensional interval vector space, given by:

$$d_E(X, Y) = \sqrt{\sum_{i=1}^n (x_i - y_i)^2} \quad (2)$$

The interval vectors may be computed from the given binary sequences in $O(n)$ time, after which $d_E(X, Y)$ can be computed in $O(k)$ time, where k is the number of onsets, thus leading to $O(n)$ time overall.

3.3 The interval-difference vector distance

Coyle and Shmulevich [12] and Shmulevich et al. [36] represent a music pattern by what they call a *difference-of-rhythm vector*. If $T = (t_1, t_2, \dots, t_n)$ is a vector of inter-onset time intervals for the notes of a rhythm then they define the difference-of-rhythm vector as $X = (x_1, x_2, \dots, x_{n-1})$, where $x_i = t_{i+1}/t_i$. In the context of cyclic rhythms considered here $x_n = t_1/t_n$, and the dimensionality of the difference-of-rhythm vector is thus n rather than $n-1$. We shall use the shorter and more descriptive term *interval-difference* vector for the difference-of-rhythm vector. Let $X = (x_1, x_2, \dots, x_n)$ and $Y = (y_1, y_2, \dots, y_n)$ denote two such *interval-difference* vectors. Shmulevich et al. [36] define the dissimilarity between X and Y by:

$$d_{ID}(X, Y) = \left(\sum_{i=1}^n \frac{\max(x_i, y_i)}{\min(x_i, y_i)} \right) - n \quad (3)$$

Clearly this distance measure is also easily computed in $O(n)$ time.

3.4 The swap distance

To combat the inherent weakness of the Hamming distance, many variants and generalizations of the Hamming distance have been proposed over the years. One early generalization is the *edit* distance which allows for insertions and deletions of notes. Discussions of the application of the *edit-distance* to the measurement of similarity in music can be found in Orpen and Huron [33] and Mongeau and Sankoff [30]. A noteworthy more recent generalization is the *fuzzy* Hamming distance [9], [10] which allows *shifting* of notes as well as insertions and deletions. Using *dynamic programming* these distances may be computed in $O(n^2)$ time.

This completely different approach to measuring the dissimilarity between two strings computes the amount of “work” required to transform one string into the other. Such an approach is common in bioinformatics where the two strings to be compared are chain polymers and the work is measured by the minimum number of basic operations required to transform one molecule into the other. The type of basic operation used varies and usually models some kind of mutation relevant to evolution (see [1], [2] and the references therein). Here

we will use a simplified version of the *fuzzy* Hamming distance that is much more efficient to compute.

The problem of comparing two binary strings of the same length with the same number of one's suggests an extremely simple edit operation called a *swap*. A swap is an interchange of a one and a zero that are adjacent to each other in the binary string. Interchanging the position of elements in strings of numbers is a fundamental operation in many sorting algorithms [14]. However, in the sorting literature a swap may interchange non-adjacent elements. When the elements are required to be adjacent, the swap is called a *mini-swap* or *primitive-swap* [7], [8]. Here we use the shorter term *swap* for the interchange of two adjacent elements. The swap distance between two rhythms is the *minimum* number of swaps required to convert one rhythm to the other. For example the rhythm $[x \cdot x \cdot x x \cdot x \cdot x \cdot x]$ can be converted to the rhythm $[x \cdot x x \cdot x x \cdot x \cdot x \cdot]$ by a minimum of four swaps, namely interchanging the third, fifth, sixth, and seventh beats with the corresponding rests preceding them. The *swap* distance may be viewed as a simplified version of the *fuzzy* Hamming distance [9], [10], where only the shift operation is used, and the cost of the shift is equal to its length. Such a measure of dissimilarity appears to be more appropriate than the Hamming distance between the binary vectors or the Euclidean distance between the interval vectors, in the context of rhythm similarity [41], [43]. The swap distance may also be viewed as a special case of the more general *earth mover's distance* (also called *transportation distance*) used by Typke et al. [44] to measure melodic similarity. Given two sets of points called supply points and demand points, each assigned a weight of material, the earth movers distance measures the minimum amount of work (weight times distance) required to transport material from the supply points to the demand points. No supply point can supply more weight than it has and no demand point receives more weight than it needs. Typke et al. [44] solve this problem using linear programming, a relatively costly computational method. The swap distance is a one dimensional version of the earth mover's distance with all weights equal to one. Furthermore, in the case where both binary sequences have the same number of "one's" (or onsets), there is a one-to-one correspondence between the indices of the ordered onsets of the sequences. For example, consider again the two sequences $X = [x \cdot x \cdot x x \cdot x \cdot x \cdot x]$ and $Y = [x \cdot x x \cdot x x \cdot x \cdot x \cdot]$, each with seven onsets. The i -th onset of X must travel the necessary distance to reach the position of the i -th onset of Y . For $i = 1$ this distance is zero. For $i = 3$ the distance is one.

The swap distance may of course be computed by actually performing the swaps, but this is inefficient. If X has one's in the first $n/2$ positions and zero's elsewhere, and if Y has one's in the last $n/2$ positions and zero's elsewhere, then at least a quadratic number of swaps would be required. On the other hand, if we compare the distances instead, a much more efficient algorithm results. First scan the binary sequence and store a vector of the x -coordinates at which the k onsets occur. For example, X and Y above yield the vectors $U = (u_1, u_2, \dots, u_7) = (1, 3, 5, 6, 8, 10, 12)$ and $V = (v_1, v_2, \dots, v_7) = (1, 3, 4, 6, 7, 9, 11)$, respectively. The difference between u_i and v_i is the number of swaps that must be performed to bring the two onsets into alignment. Therefore, in general, the swap distance between two onset-coordinate vectors U and V with k onsets is given by:

$$d_{SWAP}(U, V) = \sum_{i=1}^k |u_i - v_i| \quad (4)$$

Computing U and V from X and Y is done trivially in $O(n)$ time with a simple scan. Therefore $O(n)$ time suffices to compute $d_{SWAP}(U, V)$, resulting in a large gain over using linear or dynamic programming.

3.5 The chronotonic distance

Hofmann-Engl considers the chronotonic chain as a vector of *atomic* units rather than as a piece-wise linear function. For example, the clave *Son* which has the interval set (3,3,4,2,4) and corresponding chronotonic chain shown in Figure 5, yields the chronotonic 16-dimensional vector (3,3,3,3,3,3,4,4,4,4,2,2,4,4,4,4). To measure the similarity between two rhythms Hofmann-Engl computes a weighted Euclidean distance between their corresponding chronotonic vectors. Hofmann-Engl reports that experiments with human subjects supports the hypothesis that this measure agrees with human perceptual and cognitive models of rhythmic similarity [24]. However, viewing the chronotonic representation as a function, such as the piece-wise linear function in Figure 4, opens the door to a large family of possible distance functions with a long history in the fields of statistics and pattern recognition [39]. Given two probability density functions $f_1(x)$ and $f_2(x)$, there are many measures of the distance (or separation) between them. One of the most well known is the *discrimination information*. The discrimination information between $f_1(x)$ and $f_2(x)$, also called the *Kullback-Liebler divergence* between the distributions [5], is given by:

$$KL = \int f_1(x) \log \frac{f_1(x)}{f_2(x)} dx. \quad (5)$$

This measure is closely related to the *Kolmogorov variational* distance [40] given by:

$$K = \int |f_1(x) - f_2(x)| dx. \quad (6)$$

For surveys of these and other information measures, their properties, and their inter-relationships the reader is referred to [38] and [27]. One can of course also compute the Kolmogorov variational distance K with functions that are not necessarily probability density functions. In fact, Ó Maidín [32] proposed exactly such a measure of dissimilarity between two melodies modelled as monotonic pitch-duration functions of time, as depicted in Fig. 7, where the x -axis measures time duration and the y -axis measures the pitch of the notes. In the context of information retrieval of melodies one may also be interested in computing the minimum area between the two curves under translations of one curve in the x and y directions while the second curve remains fixed. Aloupis et al. [3] show that for two melodies (polygonal chains) consisting of n vertices this function can be computed in $O(n^2 \log n)$ time.

Here we will use the measure K to compare rhythms using the chronotonic representation proposed by Gustafson [22]. As an example, consider the rhythm [x . . x x . x . . x x . x . . .] shown in Figure 8 in chronotonic notation. This is the rhythm timeline played on the *tom-tom* drum of *Not Fade Away* by Buddy Holly and the Crickets [13]. The record was released in 1957 and became an instant hit [20]. It is a variant of the famous rhythm dubbed the *Bo Diddley beat* popularised by Bo Diddley in his 1955 hit “Bo Diddley,” given by [x . x x x . x . . . x . x . . .], and also characterized by the mnemonics “*shave and a haircut, two bits*” as well as “*bump bubbly bump bump, bump bump*”. If we superimpose the *Not*




Figure 7: Two melodies as rectilinear pitch-duration functions of time.




Figure 8: The rhythm in *Not Fade Away* in chronotonic notation.

Fade Away and *Shiko* timeline rhythms, as in Figure 9, then the measure K , or area-difference (shown shaded) between the two curves, is equal to 12. A similar calculation superimposing the *Not Fade Away* and clave *Son* timelines yields a value of K equal to 10, thus showing that the *Not Fade Away* rhythm is closer to the *Son* than the *Shiko*. In this discrete setting it is clear that K may be computed easily in $O(n)$ time.

4 Comparing Rhythmic Similarity Measures

Traditionally, rhythmic similarity measures have been evaluated or compared with respect to how well rhythms may be recognized [12], how efficiently they can be retrieved from a data-base [19], or how well they model human perception and cognition of rhythms [24], [5]. In contrast, here we compare rhythmic similarity measures with respect to how much insight they provide concerning the structural interrelationships that exist within families of rhythms, when phylogenetic trees and graphs are computed from the distance matrices determined by these similarity measures.

The similarity measures discussed in the preceding section are compared here using two families of rhythms, namely, the fundamental binary and ternary timelines used in African and Afro-American traditional music. In particular, the first family consists of the six five-beat clave-bell patterns in 4/4 time [41], and the second family is comprised of the ten seven-beat bell patterns in 12/8 time used in West African traditional music [43]. These two families are depicted in the chronotonic notation of Gustafson in Figures 10 and 11, respectively.




Figure 9: The area difference K between *Not Fade Away* and *Shiko*.

4.1 Phylogenetic trees

Several techniques exist for generating phylogenetic trees from distance matrices [25]. Some of these methods have the desirable property that they produce graphs that are not trees, when the underlying proximity structure is inherently not tree-like. One notable example is *SplitsTree* [15], [25]. Like the more traditional phylogenetic trees, *SplitsTree* computes a plane graph embedding with the property that the distance in the drawing between any two nodes reflects as closely as possible the true distance between the corresponding two rhythms in the distance matrix. However, if the tree structure does not match the data perfectly then new nodes in the graph may be introduced, as for example in Figure 12, with the goal of obtaining a better fit. Such nodes may suggest implied “ancestral” rhythms from which their “offspring” may be derived. In addition, edges may be split to form parallelograms, (or more general zonotopes) such as in Figure 16. The relative sizes of these parallelograms are proportional to an *isolation index* that indicates how significant the clustering relationships inherent in the distance matrix are. *SplitsTree* also computes the *splitability index*, a measure of the goodness-of-fit of the entire splits graph. This *fit* is obtained by dividing the sum of all the approximate distances in the splits graph by the sum of all the original distances in the distance matrix [15], [25].

5 Results

In this section we describe the structure of each *SplitsTree* obtained with each of the five distance measures described in the preceeding, when applied to the two families of rhythms. We begin with the binary (4/4 time) rhythms.

5.1 Binary rhythms

5.1.1 The Hamming distance

For the Hamming distance the splits graph shown in Figure 12 is a tree with a perfect fit of 100%. The *Son* pattern lies at the center of the tree. Indeed, the Hamming distance between the *Son* and each of the other rhythms is equal to *two*. The tree also contains one *ancestral* node with *Son*, *Soukous*, *Bossa-Nova*, and *Gahu* as its offspring with each a distance of *one* from the ancestral rhythm. This ancestral rhythm is the four-beat pattern




Figure 10: The chronotonic representation of the 4/4 time *clave* patterns.

[x . . x . . x . . . x]. Note that all four offspring contain this pattern, and they differ only in the location of the last (fifth) beat that falls in positions 12, 13, 14, and 15, for the *Soukous*, *Son*, *Bossa-Nova*, and *Gahu* rhythms, respectively. Finally, it is worth observing that there is no unique pair of candidates that represents the two most different rhythms.

5.1.2 The Euclidean interval vector distance

For the Euclidean interval vector distance the splits graph shown in Figure 13 is tree-like with a good fit of 82.3%. The center of the graph is the *Son*, but it does not have direct paths to all other rhythms. The path from *Son* to *Gahu* is via the *Bossa-Nova*. Here there are two rhythms, *Gahu* and *Soukous*, that are clearly the most dissimilar pair. Furthermore, *Gahu* is the most distinct from all the others.

5.1.3 The interval-difference vector distance

For the interval-difference vector distance the splits graph shown in Figure 14 is almost a chain with a rather poor fit of 58.3%. There is a tight cluster of rhythms comprised of the *Son*, *Rumba*, and *Shiko* patterns. These are the three rhythms that contain identical second bars (the last two beats) given by [. . x . x . .]. Here also *Gahu* and *Soukous* are the undisputed furthest pair of rhythms.

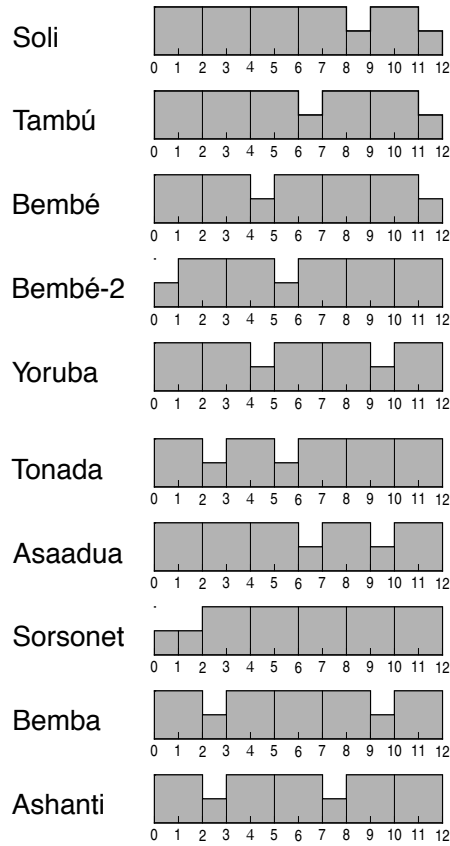


Figure 11: The chronotonic representation of the 12/8 time *bell* patterns.

5.1.4 The swap distance

For the swap distance the splits graph shown in Figure 15 is a tree with a perfect fit of 100%, and it contains no ancestral nodes. The *Son* is the center of the tree and minimizes both the maximum as well as the sum of the swap distances to all the other rhythms. *Gahu* and *Soukous* are again the unique most dissimilar pair. This splits graph is the simplest of all the graphs.

5.1.5 The chronotonic distance

For the chronotonic distance the splits graph shown in Figure 16 is a tree-like graph with a perfect fit of 100%. This distance measure is the only one that exhibits a strong clustering in this family of rhythms. In particular, the long parallelogram clearly separates *Bossa-Nova* and *Gahu* from the rest, and invites one to search for a recognizable structure that explains how these two rhythms differ from the others. An examination of their chronotonic functions of time (see Figure 10) yields a clearly distinguishable discriminating feature. The *Bossa-Nova* and *Gahu* functions have only *one* local maximum, whereas each of the four other rhythms has *two* local maxima.




Figure 12: The SplitsTree with the Hamming distance.

5.2 Ternary rhythms

5.2.1 The Hamming distance

For the Hamming distance the splits graph illustrated in Figure 17 is a complex network with a weak fit of 70.3%. The main observation here is that there is one tightly knit group of five rhythms (located in the upper right) consisting of *Soli*, *Tambú*, *Bembé*, *Yoruba* and *Asaadua*, that are separated from the remaining five which lie scattered about. A discussion of the properties of these five rhythms that sets them apart from the others is deferred to a following section.

5.2.2 The Euclidean interval vector distance

For the Euclidean interval vector distance the splits graph illustrated in Figure 18 is tree-like with a weak fit of 67.9%. Indeed, it is almost a star without any strong clustering evidence. The two parallelograms weakly suggest that *Soli*, *Tambú*, and *Bembé* form one cluster. Analysis of the chronotonic functions in Figure 11 reveals that these three rhythms have a common feature absent in the other seven, namely, all three end with a short onset interval.

5.2.3 The interval-difference vector distance

For the The interval-difference vector distance the splits graph illustrated in Figure 19 has the lowest fit of all, a meager 28.3%. Yet its two parallelograms have the same structure as those in Figure 18 with the leaves pruned off. It also suggests the same weak clustering of the *Soli*, *Tambú*, and *Bembé* rhythms.




Figure 13: The SplitsTree with the Euclidean interval vector distance.




Figure 14: The SplitsTree with the interval-difference vector distance.

5.2.4 The swap distance

For the swap distance the splits graph illustrated in Figure 20 is almost a chain with a perfect fit of 100%. There are no ancestral nodes, and the two most dissimilar rhythms, *Soli* and *Sorsonet*, are strikingly visible. The parallelogram (a rhombus) indicates two weak clusters, one comprised of *Soli*, *Tambú*, and *Bembé*, the same as in Figures 18 and 19, and the other made up of *Soli*, *Tambú*, and *Asaadua*. Examination of the chronotonic functions reveals that these three rhythms are the only ones that have both short inter-onset intervals present in the second halves of their time spans. Note that the subgraph consisting of the rhombus rhythms (*Tambú*, *Bembé*, *Yoruba* and *Asaadua*) together with *Soli* are the same cluster apparent in the upper right-hand portion of Figure 17.

5.2.5 The chronotonic distance

For the chronotonic distance the splits graph shown in Figure 21 exhibits a rich structure with a perfect fit of 100%. There are two highly clustered groups. The first group (already seen in Figures 20 and 17) consists of *Soli*, *Asaadua*, *Tambú*, *Yoruba*, and *Bembé*. The second group is comprised of *Tonada*, *Ashanti*, *Bembé-2*, and *Sorsonet*. Furthermore, *Bemba* lies half-way between the two groups.

An analysis of the chronotonic functions in Figure 11, with the above clustering in mind, yields the following property that explains the clustering. Imagine that each function is a

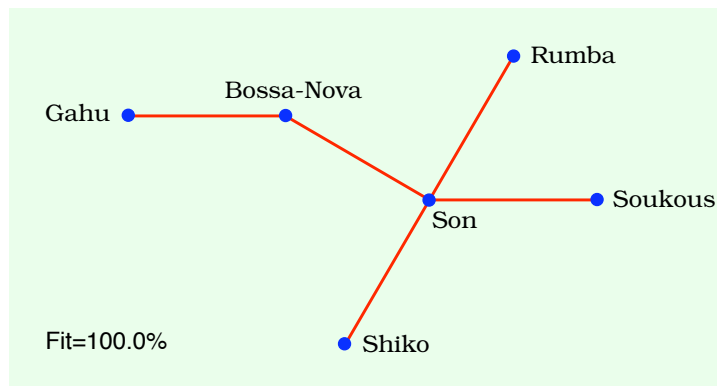


Figure 15: The SplitsTree with the swap distance.

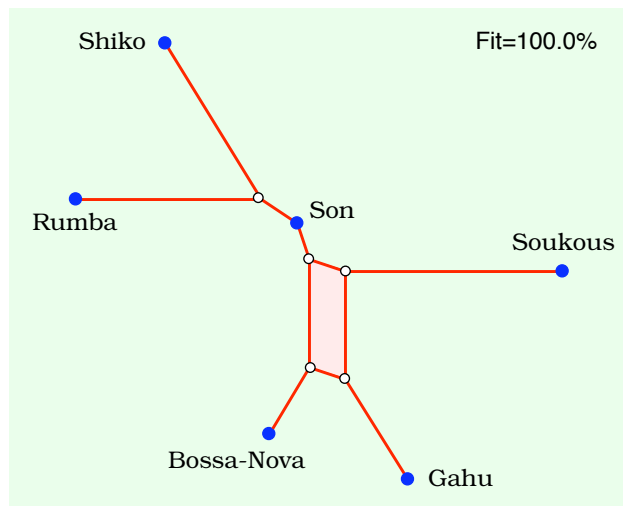


Figure 16: The SplitsTree with the chronotonic distance.

solid object with homogeneous mass equal to the area under the curve, balanced on a pivot located at the midpoint of the time span. For example, Figure 22 illustrates the *Soli* object balanced on the midpoint pivot. Note that this object has more mass on the left side of the pivot and, by “gravity”, would thus rotate in a counter-clockwise direction. So would every other member of group one. On the other hand, all the rhythms of group two have more mass on the right side of the pivot and would rotate in a clockwise direction. Furthermore, the *Bemba* rhythm (which falls half way between the two groups) is perfectly balanced on the pivot.

The splits graph also has nine potential ancestral nodes. One of these is quite distinguished from the rest, not only because of its high degree, but because it is located in the center of the entire collection, and is at a distance of *two* from *all* ten rhythms. For this reason it is labelled the *2-center* of the set. A reconstruction of this ancestral rhythm yields the pattern $[x . x . x . x . x . x .]$, i.e., a steady isochronous “heart-beat” consisting of six beats per measure.

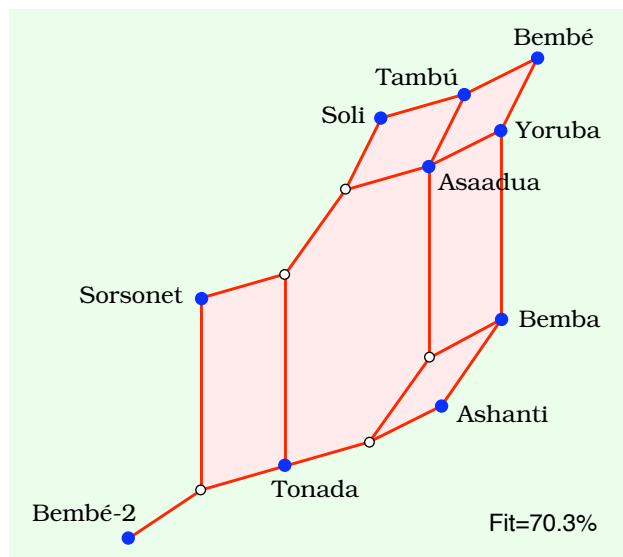


Figure 17: The SplitsTree with the Hamming distance.

6 Concluding Remarks

Several desirable criteria spring to mind for comparing phylogenetic trees of families of rhythms. Such criteria include (1) simplicity, (2) goodness of fit, (3) indication of clustering, (4) generation of interesting “ancestral” rhythms, and (5) support of musicological analyses.

Simplicity: The simplest splits graphs are obtained with the swap distance. For the binary rhythms the splits graph is a tree, and for the ternary rhythms there is only one cycle. Furthermore, neither graph has “ancestral” nodes.

Goodness of fit: The only distance measures that yield a 100% fit for both binary and ternary rhythms are the swap and chronotonic distances. The Hamming distance also produced a 100% fit for the binary rhythms.

Clustering: The most impressive clustering is obtained with the chronotonic distance. For the binary rhythms, *Gahu* and *Bossa-Nova* are clearly differentiated from the rest. For the ternary rhythms the splits graph yields two well separated clusters.

Ancestral rhythm generation: The most noteworthy “ancestral” rhythms were produced with the chronotonic distance for ternary rhythms and with the Hamming distance for the binary rhythms.

Musicological analysis: According to Locke [26] *Gahu* is one of the most unusual binary bell patterns, and he devotes a lengthy musicological analysis to make this point. His conclusions are well supported by the results obtained here with the Euclidean, swap, interval-difference, and chronotonic distance functions.

In conclusion, a comparison of the significance of the role played by the five distance measures according to each of the five criteria outlined in the preceding, suggests that the best overall rhythmic dissimilarity measure is the *chronotonic* distance, followed by the *swap* distance in close second place.

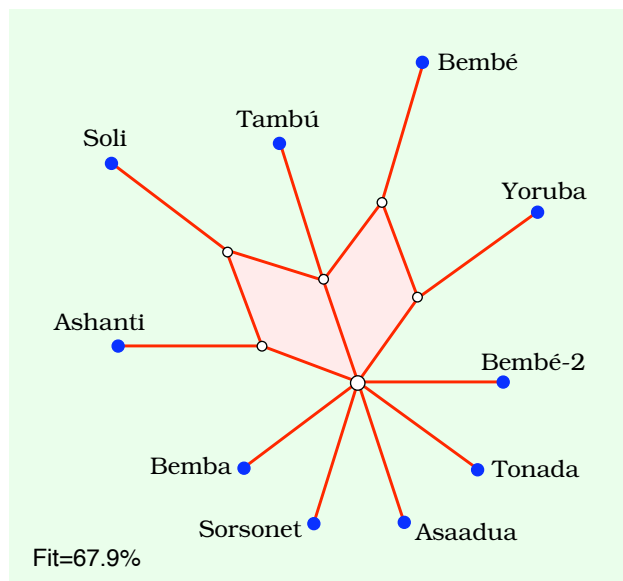


Figure 18: The SplitsTree with the Euclidean interval-vector distance.

With respect to computational complexity, all five measures may be computed efficiently in $O(n)$ time, where n is the length of the binary representation in which the rhythm is given.

7 Acknowledgements

I thank David Bryant for discussions on the *SplitsTree* program, Pat Morin and Tom Shermer for illuminating discussions on the swap distance, and Kjell Gustafson for references on the history of his chronotonic (TEDAS) representation of rhythm.

References

- [1] Greg Aloupis, Prosenjit Bose, Erik D. Demaine, Stefan Langerman, Henk Meijer, Mark Overmars, and Godfried T. Toussaint. Computing signed permutations of polygons. In *Proc. 14th Canadian Conf. Computational Geometry*, pages 68–71, Univ. of Lethbridge, Alberta, 2002.
- [2] Greg Aloupis, Erik D. Demaine, Henk Meijer, Joseph O’Rourke, Ileana Streinu, and Godfried T. Toussaint. On flat-state connectivity of chains with fixed acute angles. In *Proc. 14th Canadian Conf. Computational Geometry*, pages 27–30, Univ. of Lethbridge, Alberta, 2002.
- [3] Greg Aloupis, Thomas Fevens, Stefan Langerman, Tomomi Matsui, Antonio Mesa, Yurai Nuñez, David Rappaport, and Godfried Toussaint. Computing a geometric measure of the similarity between two melodies. In *Proc. 15th Canadian Conf. Computational Geometry*, pages 81–84, Dalhousie University, Halifax, Nova Scotia, Canada, August 11-13 2003.

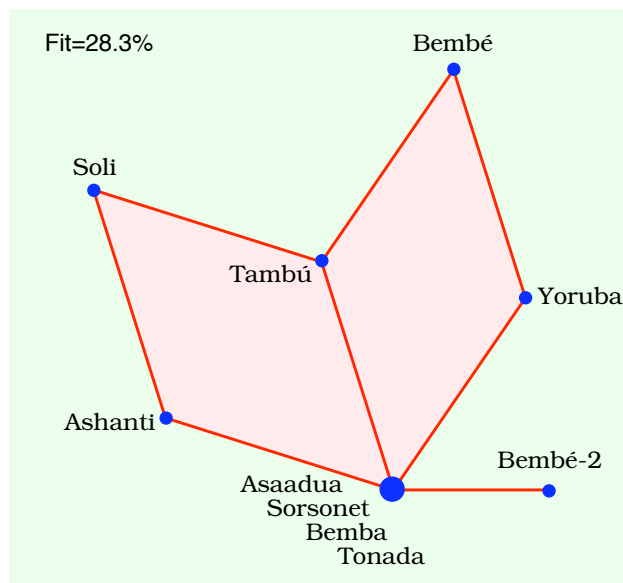


Figure 19: The SplitsTree with the interval-difference vector distance.

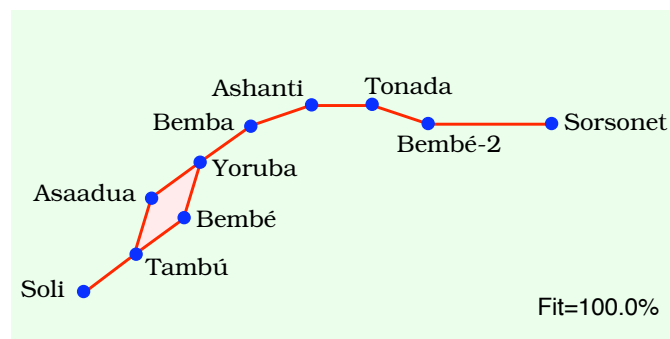


Figure 20: The SplitsTree with the swap distance.

- [4] Simha Arom. *African Polyphony and Polyrhythm*. Cambridge University Press, Cambridge, England, 1991.
- [5] Adam Berenzweig, Daniel P. W. Ellis, and Steve Lawrence. Anchor space for classification and similarity measurement of music. In *Proceedings ICME-03*, Baltimore, U.S.A., July 2003.
- [6] H. Bettermann, D. Amponsah, D. Cysarz, and P. Van Leeuwen. Musical rhythms in heart period dynamics: a cross-cultural and interdisciplinary approach to cardiac rhythms. *Proceedings of the American Physiological Society*, pages H1762–H1770, 1999.
- [7] Therese Biedl, Timothy Chan, Erik D. Demaine, Rudolf Fleischer, Mordecai Golin, James A. King, and Ian Munro. Fun-sort – or the chaos of unordered binary search. *Discrete Applied Mathematics*. to appear.

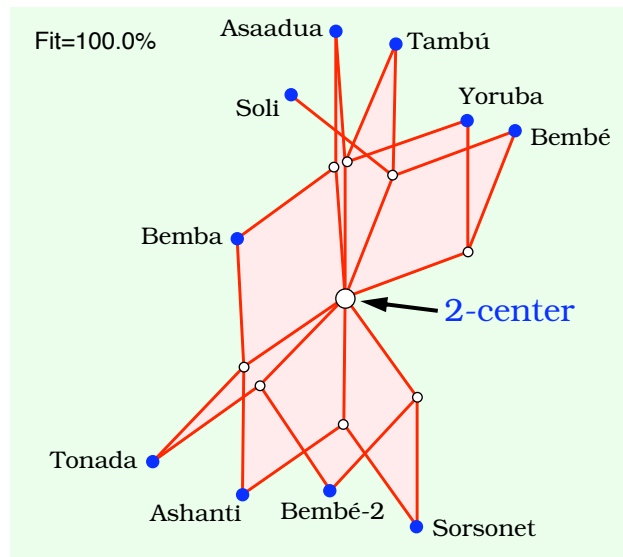


Figure 21: The SplitsTree with the chronotonic distance.




Figure 22: In the *Soli* rhythm “gravity” rotates the object in a counter-clockwise direction.

- [8] Therese Biedl, Timothy Chan, Erik D. Demaine, Rudolf Fleischer, Mordecai Golin, and J. Ian Munro. Fun-sort. In Elena Lodi, Linda Pagli, and Nicola Santoro, editors, *Proceedings of 2nd International Conference on Fun with Algorithms (FUN 2001)*, pages 15–26, Isola d’Elba, Italy, May 29-31 2001.
- [9] Abraham Bookstein, Shmuel Tomi Klein, and Timo Raita. Fuzzy Hamming distance: a new dissimilarity measure. In *Proceedings 12th Annual Symposium on Combinatorial Pattern Matching*, volume LNCS 2089, pages 86–97, Berlin, 2001. Springer-Verlag.
- [10] Abraham Bookstein, Vladimir A. Kulyukin, and Timo Raita. Generalized Hamming distance. *Information Retrieval*, 5(4):353–375, 2002.
- [11] Sung-Hyuk Cha and Sargur N. Srihari. On measuring the distance between histograms. *Pattern Recognition*, 35:1355–1370, 2002.
- [12] E. J. Coyle and I. Shmulevich. A system for machine recognition of music patterns. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, Seattle, Washington, 1998.
- [13] Buddy Holly and The Crickets. Not Fade Away. In *The Best of Buddy Holly*, The Millenium Collection, 1999. Music Corporation of America Records, Inc.

- [14] N. G. de Bruijn. Sorting by means of swapping. *Discrete Mathematics*, 9:333–339, 1974.
- [15] A. Dress, D. Huson, and V. Moulton. Analysing and visualizing sequence and distance data using SPLITSTREE. *Discrete Applied Mathematics*, 71:95–109, 1996.
- [16] Richard O. Duda, Peter E. Hart, and David G. Stork. *Pattern Classification*. John Wiley and Sons, Inc., New York, 2001.
- [17] Douglas Eck. A positive-evidence model for classifying rhythmical patterns. Technical Report IDSIA-09-00, Instituto Dalle Molle di studi sull'intelligenza artificiale, Manno, Switzerland, 2000.
- [18] Laz E. N. Ekwueme. Concepts in African musical theory. *Journal of Black Studies*, 5(1):35–64, September 1974.
- [19] Jonathan Foote, Matthew Cooper, and Unjung Nam. Audio retrieval by rhythmic similarity. In *Third International Conference on Music Information Retrieval*, Centre Pompidou, Paris, France, October 13-17 2002.
- [20] John Goldrosen. *The Buddy Holly Story*. The Bowling Green University Popular Press, New York, 1975.
- [21] Kjell Gustafson. A new method for displaying speech rhythm, with illustrations from some Nordic languages. In K. Gregersen and H. Basbøll, editors, *Nordic Prosody IV*, pages 105–114. Odense University Press, 1987.
- [22] Kjell Gustafson. The graphical representation of rhythm. In *(PROPH) Progress Reports from Oxford Phonetics*, volume 3, pages 6–26, University of Oxford, 1988.
- [23] R. W. Hamming. *Coding and Information Theory*. Prentice-Hall, Englewood Cliffs, 1986.
- [24] Ludger Hofmann-Engl. Rhythmic similarity: A theoretical and empirical approach. In C. Stevens, D. Burnham, G. McPherson, E. Schubert, and J. Renwick, editors, *Proceedings of the Seventh International Conference on Music Perception and Cognition*, pages 564–567, Sidney, Australia, 2002.
- [25] Daniel H. Huson. SplitsTree: Analyzing and visualizing evolutionary data. *Bioinformatics*, 14:68–73, 1998.
- [26] David Locke. *Drum Gahu: An Introduction to African Rhythm*. White Cliffs Media, Gilsum, New Hampshire, 1998.
- [27] A. Mathai and P. Rathie. *Basic Concepts in Information Theory and Statistics*. Wiley Eastern Ltd., New Delhi, 1975.
- [28] Benoit Meudic. Musical similarity in a polyphonic context: a model outside time. In *Proceedings of the XIV Colloquium on Musical Informatics*, Firenze, Italy, May 8-10 2003.

- [29] Abraham Moles. *Information Theory and Esthetic Perception*. The University of Illinois Press, Urbana and London, 1966.
- [30] M. Mongeau and D. Sankoff. Comparison of musical sequences. *Computers and the Humanities*, 24:161–175, 1990.
- [31] Bernard Mont-Reynaud and Malcolm Goldstein. On finding rhythmic patterns in musical lines. In *Proceedings of the International Computer Music Conference*, pages 391–397, San Francisco, California, 1985.
- [32] D. ÓMaidín. A geometrical algorithm for melodic difference. *Computing in Musicology*, 11:65–72, 1998.
- [33] Keith S. Orpen and David Huron. Measurement of similarity in music: A quantitative approach for non-parametric representations. In *Computers in Music Research*, volume 4, pages 1–44. 1992.
- [34] Jay Rahn. Asymmetrical ostinatos in sub-saharan music: time, pitch, and cycles reconsidered. *In Theory Only*, 9(7):23–37, 1987.
- [35] Jay Rahn. Turning the analysis around: African-derived rhythms and Europe-derived music theory. *Black Music Research Journal*, 16(1):71–89, 1996.
- [36] Ilya Shmulevich, Olli Yli-Harja, Edward Coyle, Dirk-Jan Povel, and Kjell Lemström. Perceptual issues in music pattern recognition. *Computers and the Humanities*, 35:23–35, 2001.
- [37] Godfried T. Toussaint. Generalizations of π : some applications. *The Mathematical Gazette*, pages 291–293, December 1974.
- [38] Godfried T. Toussaint. On some measures of information and their application to pattern recognition. In *Proceedings of the Conference on Measures of Information and their Applications*, pages 21–28, Bombay, India, August 16-18 1974.
- [39] Godfried T. Toussaint. On the divergence between two distributions and the probability of misclassification of several decision rules. In *Proceedings of the Second International Joint Conference on Pattern Recognition*, pages 27–35, Copenhagen, Denmark, August 13-15 1974.
- [40] Godfried T. Toussaint. Sharper lower bounds for discrimination information in terms of variation. *IEEE Transactions on Information Theory*, pages 99–100, January 1975.
- [41] Godfried T. Toussaint. A mathematical analysis of African, Brazilian, and Cuban *clave* rhythms. In *Proceedings of BRIDGES: Mathematical Connections in Art, Music and Science*, pages 157–168, Towson University, Towson, MD, July 27-29 2002.
- [42] Godfried T. Toussaint. Algorithmic, geometric, and combinatorial problems in computational music theory. In *Proceedings of X Encuentros de Geometria Computacional*, pages 101–107, University of Sevilla, Sevilla, Spain, June 16-17 2003.

- [43] Godfried T. Toussaint. Classification and phylogenetic analysis of African ternary rhythm timelines. In *Proceedings of BRIDGES: Mathematical Connections in Art, Music and Science*, pages 25–36, Granada, Spain, July 23-27 2003.
- [44] Rainer Typke, Panos Giannopoulos, Remco C. Veltkamp, Frans Wiering, and René van Oostrum. Using transportation distances for measuring melodic similarity. Technical report uu-cs-2003-024, Computer Science Department, University of Utrecht, The Netherlands, 2003.
- [45] O. Wright. *The Modal System of Arab and Persian Music AD 1250-1300*. Oxford University Press, Oxford, England, 1978.