

2021-07-01

Multileveled rhythmic structure of ragtime

J. Yust, P. Kirlin. 2021. "Multileveled rhythmic structure of ragtime." HCII International Conference on Human-Computer Interaction: Culture and Computing. Design Thinking and Cultural Computing. International Conference on Human-Computer Interaction. Online (Washington DC), 2021-07-24 - 2021-07-29. https://doi.org/10.1007/978-3-030-77411-0_22
<https://hdl.handle.net/2144/44331>

"Downloaded from OpenBU. Boston University's institutional repository."

The Multileveled Rhythmic Structure of Ragtime

Jason Yust¹ and Phillip B. Kirlin²

¹ Boston University, Boston, MA, 02215

² Rhodes College, Memphis TN, 38112

Abstract. Syncopation in ragtime music has been defined in multiple ways. In this study we propose a method using the Hadamard transform. We extract four-measure phrases from a corpus of ragtime pieces by Scott Joplin, James Scott, and Joseph Lamb, and convert them to 32-element binary onset vectors. The Hadamard transform converts this to another 32-element vector that can be interpreted as representing syncopation at various metrical levels. This method is closely related to a similar application of the discrete Fourier transform. Using the Hadamard representation, we show that syncopation is strongest at the quarter-note level, and that tresillo-like rhythms are especially characteristic of the genre. We identify a number of significant differences based on the position of a phrase in a sixteen-measure strain, the position of the strain in the rag, and the composer. The Hadamard representation also facilitates discovery of relationship between different levels of rhythmic organization.

Keywords: Ragtime, Syncopation, Rhythm, Hadamard transform, Discrete Fourier transform.

1 Introduction and Methodology

1.1 Introduction

When Scott Joplin’s “Maple Leaf Rag” was published in 1899, it and other ragtime pieces created a decades-long sensation in American music publishing. Contemporaneous critics consistently described the new and distinctive element of ragtime as syncopation (Blesh and Janis 1971). But syncopation is not a precisely defined term, as is evidenced by the numerous and divergent ways that it has been defined for empirical purposes (Temperley 2010, Longuet-Higgins and Lee 1984). Nevertheless, a number of recent studies have been performed with the intention to uncover patterns in the way syncopation has been used in ragtime music. Volk and de Haas (2013), for instance, were able to confirm a musicological hypothesis originally put forth by Berlin (1980) that *untied* syncopations (those occurring completely in the first or second half of a measure) dominated in the early era of ragtime before 1900, whereas *tied* syncopations (those occurring across the midpoint of a measure or across a measure boundary) only became prevalent in the later ragtime period starting in roughly 1902. Koops, Volk, and de Haas (2015) were able to precisely identify the most prevalent syncopated rhythmic patterns used in the ragtime era and also showed that the overall amount of syncopation in ragtime music tended to increase over time. Kirlin (2020) showed that certain ragtime composers in general used more syncopations than their contemporaries, and also that

syncopation is typically not equally distributed throughout a composition: highly syncopated measures tend to immediately follow other highly syncopated measures.

In this paper, we propose a new way of parsing a rhythm using the Hadamard transform, and an interpretation of this transform as a measure of different kinds of syncopation. The Hadamard transform requires periodic rhythms in cycles that multiply a basic pulse by powers of two. Ragtime is well-suited for this type of analysis, since it is consistently written in multi-leveled pure-duple meters. A typical ragtime piece is made up of about four sixteen-measure strains, which divide symmetrically into four-measure phrases of $2/4$. With a basic rhythmic unit of a sixteenth note, these can be treated as 32-element rhythmic cycles, with five potential duple metrical levels.

An important theoretical advantage of the Hadamard transform is that it is a reversible transformation, so it translates the rhythm into an organized metrically interpretable form without any loss of information. In this respect it is similar to the discrete Fourier transform (DFT), which has been used to give metrical interpretations of periodic rhythms by Amiot (2016) and Yust (2021a–b). The DFT and Hadamard transformations are, in fact, closely related; we will explore this connection further below.

1.2 Introduction

Musicologists who study ragtime often refer to Scott Joplin (1867/8–1917), James Scott (1885–1938), and Joseph Lamb (1887–1960) as the “big three” ragtime composers, due to their prolific output and the adherence of their compositions to what would become known later as the “classic ragtime” musical form, namely a composition having a sequence of sixteen-measure strains made up of four-measure phrases. The music of these three composers, while not always acknowledged as such during their lifetimes, is now seen as best exemplifying the ragtime genre.

Due the big three’s adherence to the classic ragtime form, specifically the conformance to sixteen-measure strains, we chose to focus our study solely on their music. We collected 71 compositions in total, 31 by Joplin, 28 by Scott, and 12 by Lamb. We restricted ourselves to compositions known to follow the “classic ragtime” form, which in this case meant pieces for solo piano, in duple meters (mostly $2/4$, some in $4/4$), with clearly defined sixteen measure strains. We specifically excluded songs and ragtime waltzes from our corpus. We used the sheet music on the International Music Score Library Project website (www.imslp.org) and located computer-readable scores of each composition (Blesh and Janis 1971, Magee 1998, McNally 2015).

For each composition, we manually identified the starting and ending measures of each sixteen-measure strain, labeling them “a,” “b,” “c,” and so forth, in the order in which they first appear in the music. Because most strains are repeated, sometimes with slight variations, we identified repetitions and used only the first version in our analyses. We then isolated the right-hand part of each strain and converted each one into a representation called a *binary onset pattern* (Sethares 2007). These patterns illustrate the rhythmic content of a musical passage by using ones and zeroes to represent a note onset and the absence of a note onset, respectively. Binary onset patterns are produced by examining a musical phrase at a specific level of rhythmic granularity; here, at the sixteenth-note level for music in $2/4$, and at the eighth-note level for music in $4/4$. For

every each possible “beat” of the music at that metrical level, if there is a note onset on that beat, a “1” is placed into the binary onset pattern, and a “0” otherwise. For instance, the binary onset pattern 10101010 represents the rhythm of a 2/4 measure of four eighth notes or a 4/4 measure of four quarter notes. The different metrical granularities for 2/4 and 4/4 compositions reflect the notational conventions of the genre, where the 4/4 notation simply doubles the notational values of the more common 2/4 notation. This is clearly evident in the conventional oom-pah patterns of the left hand, which proceed in eighth-notes in 2/4 and quarter-notes in 4/4. Discussions below will ignore the 2/4 versus 4/4 notational distinction and describe note values with reference to the 2/4 convention, which are implicitly doubled to apply to the 4/4 convention.

The end result of this conversion was a dataset consisting of, for each composition, a set of sixteen-measure strains with accompanying binary onset patterns, each pattern having 128 bits (8 bits per measure times 16 measures). We then took each strain and corresponding binary onset pattern and divided it into four four-measure phrases. We then ran the DFT and Hadamard transformation on each of the phrase-level binary onset patterns, resulting in a series of vectors which we analyzed.

1.3 Defining Syncopation with the Hadamard Transform

The Hadamard transform converts a vector of length 2^n into another vector of the same length, such that one can recover the original vector by multiplying by the Hadamard matrix. Figure 1 shows Hadamard matrices for $n = 1, 2,$ and $3,$ where “+” stands for 1 and “-” stands for $-1.$ The matrix can be constructed by a recursive rule: make four copies of the 2^{n-1} Hadamard matrix arranged in a 2×2 grid, and multiply the last one, in the lower righthand corner, by $-1.$

+	+
+	-

+	+	+	+
+	-	+	+
+	+	-	-
+	-	-	+

+	+	+	+	+	+	+	+
+	-	+	-	+	-	+	-
+	+	-	-	+	+	-	-
+	-	-	+	+	-	-	+
+	+	+	+	-	-	-	-
+	-	+	+	-	+	-	+
+	+	-	-	-	-	+	+
+	-	-	+	-	+	+	-

Fig. 1. Hadamard matrices for $n = 1, 2,$ and $3.$

In our interpretation, each row of the Hadamard matrix is a rhythm in a metrical cycle of length $2^n,$ with an onset present wherever a “+” appears. The number n corresponds to the number of metrical levels. At $n = 1,$ there is only one level. The zeroth row of the matrix, ++, is a trivial rhythm, and the second, +-, is a basic rhythm distinguishing on-beat from off-beat. The Hadamard transform of a rhythm will

decompose it into these two elements: how many onsets does it have in total (level 0) and what is the difference between on-beat and off-beat (level 1). A positive value for coefficient 1 indicates an *unsyncopated* rhythm, while a negative value indicates a *syncopated* rhythm, one that favors off-beats over beats.

At $n = 2$, we reproduce the trivial rhythm and the level-1 rhythm by repetition of the ++ and +- rhythms in rows 0 and 1. The new rows give two basic rhythms for level 2 by taking these two patterns and negating them in the second half of the rhythm, giving ++-- and +--+ . This process can be understood as contrasting the two halves of the rhythm using the patterns of coefficients 0 and 1. So coefficient 2 (++--) contrasts the two halves according to total number of onsets, and coefficient 3 (+--+) by the difference in syncopation at level 1. These give a complete account of metrical level 2, with positive values again associated with unsyncopated and negative values with syncopated rhythms. In our interpretation of ragtime, level 1 is the eighth-note level and level 2 is the quarter-note level, with the sixteenth note as the basic unit. Together, coefficients 2 and 3 amount to a description of quarter-note level syncopation that contrasts on-beat onsets with those that occur on the second eighth.

We can extend this logic recursively as we increase the number of levels. A $2/4$ measure is described by the 8×8 matrix, where coefficient 0 counts the number of onsets, coefficient 1 measures eighth-note syncopation, coefficients 2 and 3 measure quarter-note-level syncopation, and coefficients 4 – 7 represent four basic types of half-note level syncopation, defined by contrasting the two halves of the measure according to coefficients 0 – 3 of the 4×4 matrix. So coefficient 4 weights the first half of the measure against the second, coefficient 5 weights the eighth-note-level syncopation of the first half against that of the second half, and coefficients 6 and 7 weight the two kinds of quarter-note level syncopation between the two halves of the measure.

According to this logic, the Hadamard matrix losslessly converts the 2^n -cycle rhythm into a cardinality plus $2^n - 1$ measures of syncopation of different kinds. These break up into 1 type of eighth-note-level syncopation, 2 kinds of quarter-note-level syncopation, 4 kinds of half-note-level syncopation, etc. A given coefficient may be directly interpreted by converting it into a binary number. Each place in the binary number is a metrical level, and if a 1 appears in a place, it means we make a contrast based on that level. For instance, for $n = 3$, coefficient 1 is 001, meaning it makes simple eighth-note level contrasts. Coefficient 5 is 101, meaning it contrasts eighth-note level syncopation at the quarter-note level. Coefficient 7 is 111, meaning it contrasts type-11 syncopation at the half-note level, where type-11 syncopation is the contrast of eighth-note level syncopation at the quarter-note level.

Each coefficient of the Hadamard transformation has an associated coefficient of the DFT on the 2^n -place rhythmic vectors. The DFT converts a rhythm of length 2^n into 2^n complex numbers (coefficients) where coefficients 0 and 2^{n-1} are real (1-dimensional) and coefficients of index greater than 2^{n-1} are conjugates of those less than 2^{n-1} . The zeroth coefficients of each transform are equivalent, and the first Hadamard coefficient is equivalent to the 2^{n-1} -th DFT coefficient. Each of the other DFT coefficients, from 1 to $2^{n-1} - 1$, has two oblique axes in complex space such that each corresponds to one Hadamard coefficient. Rhythms with a high value on that Hadamard coefficient will also have a large value on the DFT coefficient. The two Hadamard rhythms that relate

to the same DFT coefficients are rotations of one another, such that they have oblique phase values in the complex space of the DFT coefficient. Since each DFT coefficient k corresponds to a sinusoidal function with a periodicity of c/k for cycle length $c (= 2^n)$, the DFT suggests an alternate interpretation of the Hadamard coefficient.

For example, consider the Hadamard transform for $n = 3$ ($2^n = 8$). The matrix rows for coefficients 5 and 7 are $+ - + -$ $- + - +$ and $+ - - +$ $- + + -$ respectively. The second is a rotation of the first back two places (an eighth note). If we rotate it back another eighth note we get the opposite (complementary) rhythm $- + - +$ $+ - + -$, another such rotation gives the opposite of the second rhythm, $- + + -$ $+ - - +$, and one more such rotation returns us to the first rhythm. These are all prototypes of coefficient 3 of the DFT for an 8-cycle, and the eighth-note rotation is a 90° rotation in the complex plane for coefficient 3. The best approximation to an even rhythm dividing the 8-cycle into 3 is the *tresillo* rhythm, 10010010. Its complement, rotated to align with the downbeat, is known as the *cinquillo* rhythm, 10110110. Both rhythms are common in a variety of popular and traditional musics, especially African diasporic traditions of the Americas. Both are also *maximally even* rhythms, which are prototypes of the corresponding DFT coefficient (Amiot 2007). While maximally even rhythms give a maximum value on the corresponding DFT coefficient for their cardinality (the definition of a prototype), the absolute maximum for any cardinality occurs instead at half the cycle length, for a rhythm exactly halfway between the two complementary maximally even collections, in the sense that it is a subset of one and a superset of the other. For coefficient 3 of an 8-cycle, this is the rhythm of Hadamard matrix rows 5 and 7, 10010110.

Cohn (2016) has pointed out that the *tresillo* and *cinquillo* are typical ragtime rhythms. We verify this assertion empirically below, and in fact can go somewhat farther to assert that these are defining rhythms for the style.

2 Results and Discussion

2.1 ANOVAs

We conducted two sets of ANOVAs to help identify significant trends in the data according to a few available factors in our dataset. For the first set of ANOVAs, we used the raw Hadamard transform results (Table 1), and for the second we used the absolute values of these (Table 2). The absolute values indicate the presence of a certain rhythmic type, regardless of its orientation with respect to the downbeat. The raw values indicate whether rhythms of that type tend to be syncopated or not. For instance, higher absolute values of coefficient 1 in some condition would indicate greater presence of eighth-note-based rhythms, without distinguishing on-beat eighth notes from syncopated eighth notes. A trend in raw values of coefficient 1 would indicate a tendency for eighth-note-based rhythms to be more or less syncopated in one condition or another.

Table 1. ANOVA results on raw Hadamard data. F-values are given for significant results only.
 $*p < .01$, $**p < .001$, $***p < .0001$

	$\frac{1}{8}$	$\frac{1}{4}$	$\frac{1}{2}$	2-m.	4-m.	Betw.
Coefficient	—		57***	42***	8.6*	310***
Position	34***					
Composer		40***	33***			
Strain		5.7***	5.6**	9.2***		4.4*
Coefficient \times Position	—					33***
Coefficient \times Composer	—			7.2**		7.2**
Coefficient \times Strain	—			4.9*		7.0**
Composer \times Strain		6.9***				
Coeff. \times Comp. \times Strain			3.4*	3.8**		3.4*

Table 2. ANOVA results on absolute values of Hadamard data.

	$\frac{1}{8}$	$\frac{1}{4}$	$\frac{1}{2}$	2-m.	4-m.	Betw.
Coefficient	—		83***	115***	209*	2456***
Position		8.5*			451***	119***
Composer			19***	13***	47***	13***
Strain			4.1*	3.9*	18***	
Coefficient \times Position	—					128***
Coefficient \times Composer	—	5.9*			4.9*	34***
Coefficient \times Strain	—	8.5***	10***			13**
Position \times Strain					5.2*	
Position \times Composer				7.0**	8.4**	
Composer \times Strain		4.7**		4.4**	5.4***	
Coeff. \times Pos. \times Comp.						6.8*
Coeff. \times Comp. \times Strain			4.4*	3.4*		3.5*
Pos. \times Comp. \times Strain					3.3*	

In each set of ANOVAs, we treated coefficient number as a factor, in order to compare different coefficients. Since the number of coefficients was large, to narrow down the sources of significant results, we split the analysis into six separate ANOVAs, by dividing the coefficients into the five different levels (eighth, quarter, half, two-measure, four-measure), and then averaging across each level for a sixth between-level ANOVA. Because these ANOVAs have high statistical power, and to avoid false positives due to multiple tests, we set a relatively conservative α of $p < .01$, and focus on highly significant ($p < .001$) results in our analysis and discussion.

The factors included in all the ANOVAs were • coefficient, • composer, • position in strain, • strain position in piece, and all possible interactions. The great majority of the pieces in the corpus consist of four strains of sixteen measures each. There are four possible positions in each strain for the four-measure phrases (1–4). The position of the strain in the piece is indicated by the letters *a*, *b*, *c*, and *d*. Isolated introductions and interludes between the strains were excluded to better target possible effects of the strain position factor.

The results of greatest interest are those involving the coefficient factor. There are highly significant simple effects of coefficient on both dependent variables at all levels except eighth-note (which has only one coefficient and therefore no possible contrasts of coefficient), and quarter-note, which has only two coefficients (numbers 2 and 3). The between-level coefficient factor also interacts with all of the other factors (position, composer, and strain) significantly on both dependent variables. Interactions of coefficient number and strain are significant on both dependent variables within all the high levels (half-note, two-measure, and four-measure). We also found a number of three-way interactions of coefficient, composer, and strain at the higher levels, and between levels.

Clearly all factors were significant in a number of respects. To investigate these further we consider two different descriptive statistics for different combinations of factors below: (1) The mean absolute values of coefficients, and (2) what we will call the *bias*, the mean raw value divided by the mean absolute value of a coefficient. The bias varies from -1 to 1 , where -1 indicates that all instances of a given coefficient are negative, 1 means they are all positive, and 0 means they are perfectly balanced between positive and negative. The bias is a measure of the tendency towards syncopation for each rhythmic type.

In all the analyses that follow, when averaging across the three composers, we weight each composer equally, so as not to privilege Joplin and Scott, who are represented by a larger number of pieces than Lamb.

2.2 Effects of Position and Coefficient

The position factor varies from 1 to 4 and indicates the position of each four-measure phrase within the sixteen-measure strains. We can see very clearly in the data that differences relating to position are largely due to position 4, the concluding phrase. Figure 2 shows the mean absolute values of all coefficients for the different positions. Positions 1 and 3 are nearly identical, and show a clear between-level trend, with high values at the eighth and half-note levels, smaller values at the two-measure level, and still smaller at the four-measure level. Final phrases are different in two salient ways: they have larger values across the board at the four-measure level (coefficients 16 through 31), especially the simple measure-by-measure contrasts of coefficients 16 and 24. The reason for this is obvious: since cadences typically occur in these measures, there is less overall rhythmic activity in the last two measures, leading to high values for all contrasts at the four-measure level. The other difference of position 4 is at the half-note level, with lower values for all coefficients at this level except for number 7. This indicates less overall contrast between the two halves of each measure.




Fig. 2. Absolute value by coefficient number, averaged over the position in the strain

Position number 2 shows similar tendencies to position 4, but to a lesser extent. The reason for this is also evident: since this phrase occurs at the midpoint of the strain, it can sometimes also have cadential rhythms, similar to those in the final phrases. The cadential rhythms in that position are either less frequently occurring, or less extreme, leading to a smaller difference.

We can also use Figure 2 to make some observations about the main effects of coefficient. One coefficient stands out distinctly: number 7. This coefficient identifies tresillo-type rhythms, and its prominence is quite striking. As noted above, Cohn (2016) and others have pointed out that this rhythm is distinctive of ragtime, and our data, as it turns out, bears out this observation quite clearly. Not only is coefficient 7 by far the largest one across the corpus, it is also consistently so, across all phrase positions, and, as we will see below, all strain positions and composers. In particular, position 4, even though it generally gets lower values at the half-note level, still has about the same values on coefficient 7.

We found highly significant main effects of coefficient number at all levels from the half note up, and between levels, for both absolute values and raw values. The primary between-levels effect on absolute values is a gradual reduction in coefficient size at higher levels. At the 2-measure level, coefficients 8 and 9, and to a lesser extent 12, are more heavily weighted, while numbers 10 and 13 less so. Numbers 14 and 15 appear to be more important in first and third phrases (coefficient \times position effects at the half-note and 2-measure level only barely fell short of our significance criterion, at $p = .01$ and $p = .02$ respectively). At the four-measure level, we see consistently higher values for coefficients 16, 17, 23, 24, and 25. The general pattern at higher levels is that multiples of 8 and 4, which make contrasts of density between measures or half-measures, are typically important, as are coefficients one higher than a multiple of eight, which contrast the eighth-note syncopation from one measure to the next. This means that eighth-note syncopation is an important element of the rhythmic style, and tends to be maintained within measures, and contrasted, as opposed to sustained, from one measure to the next.

Other higher-level coefficients of possible interest other than these are 14, 15, and 23. These all share a relationship with the tresillo-like rhythm of coefficient 7. Coefficient 15 contrasts this rhythm from one measure to the next: +---+---+---+---+.

We can also interpret it as *extending* the dotted-sixteenth-generated pattern of the tresillo through two measures (with the basic framework 10010010 01001001), a possibility discussed by Cohn (2016). Coefficient 23 indicates a similar kind of contrast of coefficient-7-type rhythms at the four-measure level (from the first two to the second two). Finally, coefficient 14 may be understood as an augmentation of the coefficient-7 pattern: ++-----++ ---++++--.

Figure 3 compares biases of coefficients between the positions. Again, position 4 behaves very differently than the others. Of particular interest are coefficients 1–3. Coefficient 1, the eighth-note level, gets amongst the highest absolute values (Fig. 2). However, its bias is relatively low for positions 1, 2, and 3 (Fig. 3), meaning that, while rhythms are more often unsyncopated at the eighth-note level (the bias is positive), syncopated rhythms at this level are fairly common (high absolute magnitude despite low bias). The quarter-note level (coefficients 2 and 3) is negative-biased for positions 1, 2, and 3, meaning that this form of syncopation (favoring the weak eighths) is prevalent. Both of these trends are weak or absent for position 4: these phrases have a high positive bias for coefficient 1, and zero bias for coefficients 2 and 3, meaning that eighth-note-level syncopation is largely absent in final phrases, and quarter-note level syncopation is not so prevalent.

The bias pattern for cadential rhythms includes high positives for coefficient 16 and high negatives for coefficient 24, which together indicate a lower density of attacks in the fourth measure. At the same time, we see the opposite trend in coefficients 17–19 (negative bias) and 25–27 (positive bias). These coefficients all involve contrasts of eighth- and quarter-note-level syncopation, and indicate that fourth measures, in addition to being lower density, have less of these kinds of syncopation.

In section 1 we noted the relationship of the Hadamard transform to the DFT. We briefly illustrate the correspondence in Figure 4 by displaying the average magnitudes of DFT coefficients. Along the x-axis are DFT coefficient numbers, which refer to a possible division of the 32-beat cycle. The corresponding Hadamard coefficients are labeled above each point. The lines connect points at three levels, multiples of 4, 2, and 1, which correspond to the first three levels up to the measure, the two-measure, and the four-measure levels respectively. This illustrates the general tendency for lower magnitudes at higher levels, already observed in the Hadamard transformed data. The high value at DFT coefficient 12 matches the high value at Hadamard coefficient 7. At the two- and four-measure levels, we see a tendency for higher values on the lower, density-based coefficients (1 and 2), lower values just above this, and a slight increase approaching 16. We could continue to use DFT phase values to detect syncopations, but since an analysis of DFT transforms would essentially duplicate our analysis using the Hadamard transform, we will take it no further at present and continue the analysis of Hadamard-transformed data below.




Fig. 3. Bias by coefficient number, averaged over the position in the strain




Fig. 4. Average DFT magnitudes across the corpus. The corresponding Hadamard coefficients are listed above each point.

2.3 Differences between the Three Composers

In light of the evident special status of final phrases, we will largely focus our analysis of differences between the three composers on positions 1 and 3. Before putting final phrases aside, however, consider the pattern of results, separated by phrase, for Lamb alone, in Fig. 5. While Joplin and Scott (not shown) tend to look a lot like the overall trends in Fig. 3, with Joplin having slightly greater tendency to use cadential rhythms in position 2, Lamb seems to treat position 2 very differently. The higher values at the four-measure level overall suggest that he uses cadential rhythms frequently in this position, but the pattern of four-measure-level coefficient values is very different from that for final phrases, which means that his mid-strain cadential rhythms are distinct from the usual final-phrase cadential endings. In particular, the mid-strain cadential rhythms for Lamb seem to have more to do with contrasts of syncopation (coefficients 17, 23, 25) than overall density of the measures (coefficients 16 and 24).




Fig. 5. Absolute values by Hadamard coefficient for Lamb, separated by position in the strain

The ANOVAs in Tables 1 and 2 found significant interactions of composer and coefficient between levels and at the quarter-note level. Figure 6 breaks the data up by composer for phrase positions 1 and 3 only. An evident between-levels difference is that the half-note level (coefficients 4–7) is stronger for Scott overall, indicating a tendency to use more measure-to-measure rhythmic repetition, especially with the tresillo-type rhythms (coefficient 7). At the quarter-note level, we find that Lamb has a preference for simple quarter-note syncopations (of the form $--++$, coefficient 2) while Joplin prefers those compounded with eighth-note level contrasts (of the form $-++-$, coefficient 3)




Fig. 6. Average absolute magnitudes of positions 1 and 3 by coefficient number, for each composer

Figure 7 shows the biases for positions 1 and 3 separated by composer. Overall, Joplin has a noticeably less predictable rhythmic palette, resulting in a flatter bias curve. Lamb stands out for his large biases. The ANOVA for raw values found a significant composer \times coefficient interaction at the two-measure level (coefficients 8–15). There is a large difference between Lamb and the others on coefficient 9 where Lamb strongly prefers eighth-note syncopations in the second and fourth measures of the phrase, while

the others prefer it in the first and third measures. There is also a divergence between Lamb and Scott on coefficient 15, with Scott tending to start two-measure units with downbeat oriented tresillo-type rhythms, and Lamb instead putting these in the second of a two-measure unit. In salsa terminology, we might say Scott's rhythms tend to suggest a 3-2 clave, while Lamb's more of a 2-3 (Peñalosa 2012)




Fig. 7. Biases for positions 1 and 3 by coefficient number, for each composer

2.4 Effects of Strain Order

The ANOVAs showed a number of effects having to do with the ordering of strains in the rag. Although there is a tendency to regard the strains of a rag as arbitrarily strung together, in fact the variation of rhythm from strain to strain create an overall shape to the typical rag. Our data shows that there are consistent patterns of rhythmic development over the strains that lead to reliable identifying traits of the strains based on their position in the rag. Furthermore, it shows that there are composer-specific strategies to how this works. Results involving strain in the ANOVAs mostly appeared within the half-note and 2-measure levels, and the strongest results were interactions of coefficient, strain, and composer, an indication of composer-specific strategies.

Figure 8 shows the averages by strain for each composer separately. Since we found that final positions behave substantially differently, these averages include only positions 1 and 3, although averages including all positions would present a similar picture.

An especially notable trend at the quarter-note level appears in Lamb and Scott's data, involving coefficients 5 and 7. Though coefficient 7 is high overall, for these composers it is especially high in *c* strains. On the other hand, coefficient 5 stands out in *d* strains. This coefficient detects rotated versions of the tresillo pattern, with the basic rhythm 10100110, and eighth-note rotation of the coefficient-7 rhythm 10010110. This suggests an importance of the tresillo-type rhythms to larger formal narratives for these composers, where the eighth-note rotation of these rhythms is distinctive of the transition from *c* strains to final *d* strains. Considering biases however, shown in Figure 9, we see an important difference between Lamb and Scott: Lamb has a positive bias for coefficient 5 in *d* strains where Scott's is negative. This means that, for Lamb, the

transition into the final strain is characterized by a rotation of the tresillo that remains unsyncopated, where for Scott is a rotation to a syncopated tresillo, the former typified by the rhythm 10100101 and the latter by 01011010.




Fig. 8. Absolute values by Hadamard coefficient number and strain, for each composer




Fig. 9. Biases by Hadamard coefficient number and strain, for each composer

Coefficients 5 and 7 also appear to be important to distinguishing strains for Joplin. Like Lamb and Scott, Joplin's strain *d* rhythms show a sharp reduction in coefficient 7 and increase in coefficient 5. We do not see the increase in coefficient 7 for *c* strains in Joplin's data, however. Instead, we see similar values across all the earlier strains, and coefficient 5 has a similar prominence in *b* strains as well as *d* strains. A difference between these appears in the bias data in Figure 9, which is about zero for coefficient 5 in *d* strains, and strongly positive in *b* strains. Thus, Joplin tends to use both syncopated and unsyncopated versions of the coefficient-5 rhythm in *d* strains, while strongly favoring the unsyncopated type in *b* strains.

A general between-level trend, evident most strongly in Joplin's data but also consistent with Scott's, is smaller values at the four-measure level for strain *d*, meaning

these tend to have more rhythmic repetition between the two-measure units. For Joplin specifically, *c* strains show higher values across the four-measure level, indicating greater contrasts between the two measure units.

The 2-measure level (coefficients 8–15) seems to be a particularly important locus on different treatments of the strains between composers. For Lamb, coefficient 12, with a strong negative bias (----++++ +++++----), is especially distinctive of *a* strains, while for Scott, coefficient 11 (+----+---+ -++---++-) with a strong positive bias is especially distinctive of *d* strains.

Coefficient 9, which contrasts eighth-note syncopation from measure to measure, is one of the most important at the two-measure level. Relatively high absolute values of this coefficient are consistent across composers and strains, but we see distinct differences in bias. For Scott, coefficient 9 is negative-biased in all strain positions, though weakly so in *c* strains. A negative bias indicates a tendency to put eighth-note syncopation in the first measure of a pair. Lamb is the opposite, with positive biases in all strain positions, especially *b* strains. Joplin, like Scott, usually has a negative-biased coefficient 9, except in *d* strains, where he puts more eighth-note syncopation in second measures, like Lamb.

Each composer appears to have a distinctive method of syncopation for *a* strains. Scott and Lamb stand out with almost perfect negative bias for coefficient 2, an extremely consistent use of quarter-note syncopation. Lamb also has near perfect negative bias in coefficients 12 and 15, indicating syncopation at the two-measure level that favors the second beat of an initial measure and the downbeat of a second measure. The prominence of these coefficients in *a* strains is also visible in the absolute value data. Coefficient 15 may also indicate the presence of tresillo-type rhythms oriented around the downbeats of second or fourth measures of phrases.

3 Concepts of Syncopation and Multi-Levelled Structure

Concepts of syncopation based on the Hadamard transform or DFT are not equivalent to ordinary common-sense usage of the term, but there is a definite relationship between the two. Two special properties of the definition proposed here are its balance and completeness. It is balanced in the sense that there are exactly as many theoretically possible unsyncopated rhythms as syncopated ones on every dimension. It is complete in the sense that every possible difference between two rhythms can be characterized as in the direction of more or less syncopation on some set of individual dimensions. Ordinary concepts of syncopation have neither of these properties: they are unbalanced, since there are more ways to be syncopated than unsyncopated, and incomplete, since certain differences between rhythms may not affect how it is classified. More specifically, ordinary concepts of syncopation primarily focus on the ones with simpler Hadamard representations, especially low-level coefficients that are powers of 2 (1, 2, 4) or sums of two powers of two (3 and 6). A rhythm might be syncopated, in conventional terms, not only by having a negative value on these coefficients, but also by having a zero or small positive value. Rhythms that concentrate their energy in other coefficients, especially 5 and 7, may be regarded as syncopated even if they have high positive values on

those coefficients. For instance, a version of the tresillo rhythm like 10100100 has positive values on coefficients 1, 2, and 4, but still be understood as syncopated because it has a much larger coefficient 5, even though it is a positive value.

Therefore, we propose that the Hadamard representation can be calibrated fairly well to common-sense notions of syncopation by focusing on certain parameters and adjusting thresholds. With that in mind, consider the question: in what sense is ragtime, the “syncopated idiom,” syncopated? Our analysis identifies three relevant features:

1. Eighth-note-based rhythms (large absolute values on coefficient 1) are common, but they have a bias typically close to zero, meaning that they are often highly syncopated (negative).
2. At the quarter-note level, levels of syncopation are very high, such that coefficients 2 and 3 have negative biases in most conditions.
3. The most prominent type of rhythm is represented by coefficient 7. This has a high positive bias, but rhythms with large positive values on coefficient 7 count as “syncopated” in conventional terms, because they concentrate energy in a more complex coefficient. The same can be said of coefficient 5, which, though not nearly as prominent as coefficient 7, has a substantial presence and seems to be important in certain circumstances, such as in concluding (*d*) strains.

All of these points focus on lower-level rhythmic organization, from the measure level down. At higher levels, concepts of syncopation become more diffuse and less salient as such. Our results at the two- and four-measure levels are more readily understood by thinking of them as patterns of distribution of onset density and lower-level forms of syncopation across the four measures of a phrase.

The Hadamard transform helps us relate forms of syncopation at different levels in this way. As an illustration, let us consider a specific case, comparing Joplin and Scott’s practice in *c* and *d* strains, using raw Hadamard values rather than splitting them into absolute values and biases. Figure 10 averages the four conditions across levels, using just first and third phrases. The basic strategy is the same, only more extreme for Scott: the eighth-note level is balanced, quarter-note syncopated, and half-note unsyncopated, more strongly in *c* strains and less so in *d* strains. At the two-measure level, contrastingly, we see balance in *c* strains and positive values in *d* strains. Since the two-measure level contrasts pairs of measures, this means that *c* strains usually have similar rhythms from measure to measure, while in *d* strains they are more varied. In other words, *d* strains trade low-level rhythmic contrast for between-measure contrast, shaping rhythmic ideas over somewhat longer spans.

We can get a closer look at this by comparing the first three levels directly to the two-measure level. Figure 11 overlays them, separately for Joplin and Scott. The solid lines show coefficients 1–7, and the dotted lines show 8–15, such that coefficients differing by 8 are aligned. An outward movement going from solid to dashed (away from zero) indicates that a balance in some low-level parameter becomes differentiated when split into paired measures. An inward movement shows that some prominent low-level rhythm type is simply repeated measure to measure. For both composers, *c* strains have purely inward motion: the low-level rhythms tell the whole story. In *d* strains, there is outward motion at coefficients 1 and 3. The comparison of coefficients 1 and 9 for Scott

show that what appears to be a balance between notes on and off the eighth-note metrical positions, when viewed at the two-measure level, turns out to a tendency to alternate from eighth-note syncopation in odd-numbered measures to straight eighths in even-numbered measures. The similar comparison of coefficient 3 to 11 (and 2 to 10 for Joplin) shows something similar at the quarter-note level. A balance between on-beat and weak eighth is actually an alternation of straight quarter-note and syncopated quarter-note rhythms.




Fig. 10. Average of Hadamard coefficients over levels, for Joplin and Scott in *c* and *d* strains

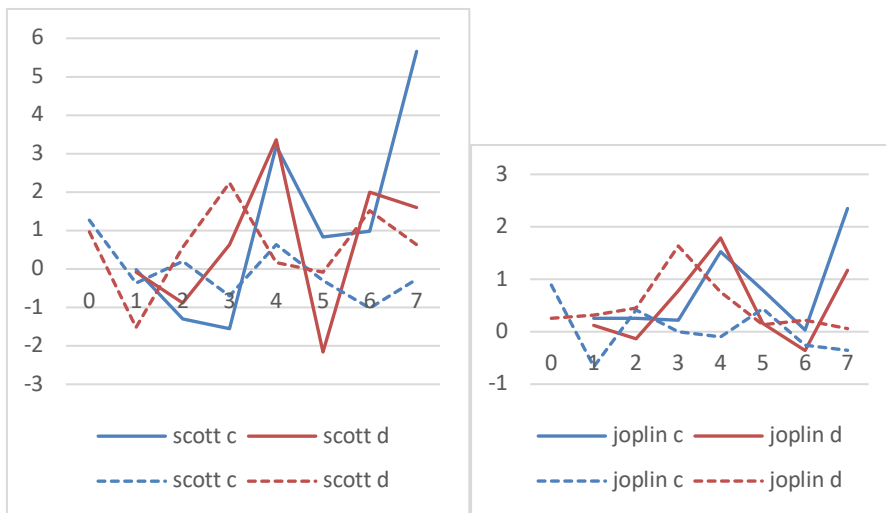


Fig. 11. Average of Hadamard coefficients 1–7 and 8–15, for Joplin and Scott in *c* and *d* strains. These are overlaid such that coefficient numbers 8–15 appear at the same x-positions as those less by eight.

This specific case is illustrative. The structure of the Hadamard transform is well suited to detecting these sorts of interactions between levels. It also provides a way of characterizing syncopation that, unlike other definitions, has the character of a lossless transform, making it particularly flexible and useful for the kind of data analysis research represented here. This allows us to answer a question like, “in what sense are ragtime rhythms syncopated?”, very precisely. It also has the potential to reveal many stylistic features with great detail, of which we have only highlighted a select few here.

References

1. Amiot, E.: David Lewin and Maximally Even Sets. *Journal of Mathematics and Music* 1(3), 157–172 (2007).
2. Amiot, E.: *Music through Fourier Space*. Springer: Heidelberg.
3. Berlin, E.A.: *Ragtime, A Musical and Cultural History*. University of California Press: Berkeley (1980)
4. Blesh, R., Janis, H.: *They All Played Ragtime*. 4th edn. Oak Pub., New York. (1971)
5. Cohn, R.: A Platonic Model for Funky Rhythms. *Music Theory Online* 22(2). <http://mtosmt.org/issues/mto.16.22.2/mto.16.22.2.cohn.html>
6. Huron, D., Ommen, A. An Empirical Study of Syncopation in American Popular Music, 1890–1939. *Music Theory Spectrum* 28(2), 211–230 (2006).
7. Kirilin, P. A Corpus-based Analysis of Syncopated Patterns in Ragtime. *Proceedings of the 21st International Society for Music Information Retrieval Conference*, 647–653 (2020).
8. Kooops, H.V, Volk, A., de Haas, W.B.: Corpus-based Rhythmic Pattern Analysis of Ragtime Syncopation. *Proceedings of the 16th International Society for Music Information Retrieval Conference*, 483–489 (2015).
9. Longuet-Higgins, H.C., Lee, C.S.: The Rhythmic Interpretation of Monophonic Music. *Music Perception* 1(4), 424–441 (1984).
10. Magee, J.: *Ragtime and Early Jazz*. *The Cambridge History of American Music*, David Nicholls, Ed. Cambridge University Press, 388–417 (1998).
11. McNally, W.M.: *Ragtime Then and Now: Composers and Audiences From the Ragtime Era to the Ragtime Revival*. Ph.D. dissertation, The City University of New York (2015).
12. Peñalosa, D.: *The Clave Matrix: Afro-Cuban Rhythm: Its Principles and African Origins*. Bembe Books, Redway (2012).
13. Sethares, W.A.: *Rhythm and Transforms*. Springer, London (2007).
14. Temperley, D. Modelling Common Practice Rhythm. *Music Perception* 27(5), 355–376 (2010).
15. Volk, A., de Haas, W.B.: A Corpus-Based Study of Ragtime Syncopation. *Proceedings of the 14th International Society for Music Information Retrieval Conference*, 163–168 (2013).
16. Yust, J.: Steve Reich’s Signature Rhythm, and an Introduction to Rhythmic Qualities. *Music Theory Spectrum* 43(1) (2021a).
17. Yust, J.: Periodicity-Based Descriptions of Rhythms and Steve Reich’s Rhythmic Style. *Journal of Music Theory* 65(1) (2021b).