MICROTIMING ANALYSIS IN TRADITIONAL SHETLAND FIDDLE MUSIC

Estefanía Cano
Fraunhofer IDMT
cano@idmt.fraunhofer.de

Scott Beveridge
Songquito
scott.beveridge@songquito.com

ABSTRACT

This work aims to characterize microtiming variations in traditional Shetland fiddle music. These microtiming variations dictate the rhythmic flow of a performed melody, and contribute, among other things, to the suitability of this music as an accompaniment to dancing. In the context of Shetland fiddle music, these microtiming variations are often referred to as *lilt*. Using a corpus of 27 traditional fiddle tunes from the Shetland Isles, we examine inter-beat timing deviations, as well as inter-onset timing deviations of eighth note sequences. Results show a number of distinct inter-beat and inter-onset rhythmic patterns that may characterize lilt, as well as idiosyncratic patterns for each performer. This paper presents a first step towards the use of Music Information Retrieval (MIR) techniques for modelling lilt in traditional Scottish fiddle music, and highlights its implications in the field of ethnomusicology.

1. INTRODUCTION

The Shetland Isles is an archipelago situated 100 miles north of the Scottish mainland (Figure 1). The distinctive music-culture of the Isles reflects a wide range of influences from around the North-Atlantic, including mainland Scottish, Scandinavian, and North-American, in particular. The earliest evidence of a performance tradition on violin, or fiddle, in the Shetland Isles dates from the eighteenth century, but there is also evidence of a pre-violin string-instrument tradition that dates to a much earlier time. However, it was over the course of the nineteenth century and into the twentieth century that the modern instrument established itself as an important component of ceremonies and rituals where dancing often played an important role [3]. The fact that solo fiddle music was used for dancing had a great influence on performance practice, with commentators often referring to the 'lilt' of a musician's performance when describing its suitability as an accompaniment to dancing. Even though a strict definition of lilt does not exist, the term often refers to the rhythmic flow imparted to the music by the performers. The term lilt bears a resemblance to the concept of swing in Jazz performance, where

Figure 1: The Shetland Isles

notated consecutive eighth notes are not played with equal duration, often approximating a tied triplet notation with a 2:1 ratio [4]. Given that an exact definition of lilt does not exist, and that many rhythmic elements could potentially play a role in imparting this dance feel to the music, this paper focuses on understanding general trends in microtiming variations in Shetland fiddle music. With this work we hope to bring a better understanding of what "*playing with lilt*" might entail.

2. RELATED WORK

The first systematic analysis of lilt in Shetland fiddle music was carried out by Peter Cooke in 1986 [3]. In this work, an electrokymograph, an early device used to measure variations in pressure, was used to capture the pitch contour and envelope of sound from a live fiddle performance. This was then used to manually calculate note onsets and approximate note durations. Anecdotal observations made by Cooke show eighth note durations that are rarely equal, with ratios of consecutive eighth notes in the range of 4:3 to 2:1. This *long-short* duration pattern is reminiscent of

Performer	Region	Title	Tr. ID
Willie Hunter	Mainland	Aandowin at da Bow	81352
		Caald Nights o Winter	81349
		Come Again Ye'r Welcome	89928
		Da Aald Hill Grind	81386
		Da Forfeit o da Ship	81351
		Fram Upon Him	81353
		Garster's Dream	81354
Bobby Peterson	Mainland	The Flowers of Edinburgh	80293
		Clean Pease Strae	36653
		Glen Grant	36656
		Kebister Head	80295
		MacDonald's Reel	80294
		The Fishers Hornpipe	36646
		The Millers Hornpipe	36647
		Willafjord	36644
John Jamieson Irvine	Whalsay	Auld Reel o Whalsay	91259
Robert Bairnson	Mainland	Da Burn o'Finnigirth	50890
		Da Hill o'Finnigirth	50890
		Mangaster Voe	50888
		Oliver Jack	50889
		The Headlands	50888
		Willafjord	50889
Andrew Poleson	Whalsay	Greigs Pipes	96869
		Lady Mary Ramsay	96896
		Supple Sandy	100412
		Unknown Reel	96861
		Up da Stroods da Sailor Goes	100411

Table 1: The Shetland fiddle corpus. All the recordings were made between 1955 - 1977 and are available in the Tobar an Dulchais collection.

swing in jazz performances [4, 6]. The seminal work carried out by Cooke in the 80s has seen a recent resurgence of interest, and has been the focus of recent analysis using state-of-the art computational analysis methods [1].

Swing research in fiddle music has also been carried out in the context of Irish traditional music [9]. In this work, a simple pitch detection and autocorrelation approach was used to estimate note durations and calculate 'swing percentages'. This method was proposed to detect the presence of swing in both synthesized and audio recordings.

Given its central role in the Jazz repertoire, swing has received considerable attention in this domain [2, 4, 6]. In [6], Friberg and Sundström examine the relationship between swing ratio and tempo in a number of popular jazz recordings. Using spectrograms obtained from the cymbal, the swing ratio was calculated based on the timing of successive eighth notes. The reported swing ratio ranged between 3.5:1 at low tempi to 1:1 at high tempi. In a later reexamination of this work, [4] applied automatic techniques to estimate swing ratios from a larger corpus of jazz recordings. The results support the initial findings of Friberg and Sundström with the added benefit that they can be applied at scale. Microtiming variations play an important role in a number of other music traditions including Ragtime syncopation [10], Samba rhythms [7], and Jembe music from Mali [8].

3. METHOD

This work represents an initial step towards the computational analysis of microtiming variations in Shetland fiddle music. As a starting point, it was necessary to compile a relevant fiddle data set. Once the recordings had been selected, all the tunes were transcribed to traditional music notation by an expert musicologist (Sect. 3.1). To be able to extract accurate timing information, each transcription was aligned in time with its corresponding audio track. For the alignment, the transcriptions were exported in MIDI format, and a MIDI-to-audio synchronization algorithm was used (Sect. 3.2). Finally, the time-aligned transcriptions were analyzed to extract timing information in the symbolic domain (Sect. 3.3). In the following sections, we provide a detailed description of all the steps in our analysis.

3.1 Fiddle Corpus

The fiddle corpus compiled in this work comprises 27 recordings of solo fiddle tunes available in the Tobar an Dualchais website.[1] Tobar an Dualchais is a collaborative project aimed at the digitization and preservation of Gaelic and Scottish recordings. The recordings in the corpus are field recordings made on a Nagra III reel-to-reel tape recorder with Sennheiser microphones. These were captured as monophonic wave files at 22.1kHz, 16 bit resolution. Table 1 provides a description of the corpus. Five performers are included, three from Mainland Shetland and two from the island of Whalsay (see Fig.1 for a map of the Shetland Isles). Track IDs from the Tobar and Dualchais collection are also provided to allow future research on the same corpus. For more information on the corpus see our accompanying website.[2]

The entire corpus was transcribed by a musicologist from the University of Aberdeen in Scotland (with expertise on Shetland fiddle traditions), providing us with transcriptions in conventional music notation for each of the 27 tunes.

3.2 MIDI-to-audio Synchronization

To calculate microtiming variations, the 27 transcriptions were exported in MIDI format and synchronized to their corresponding audio tracks. Synchronized MIDI transcriptions allow us to find the exact location in time of note onsets. For the MIDI-to-audio alignment, the method proposed in [5] was used. This method represents both the audio and MIDI streams as a combination of chroma and chroma onset features. The optimal alignment between audio and MIDI is obtained through dynamic time warping (DTW) using dynamic programming. Two difficulties were observed with the automatic synchronization. First, synchronization of the initial bars was inaccurate for some tunes, resulting in noticeably asynchronous audio and MIDI notes. Second, given that the tuning of certain files (or segments of the files) often lays between two consecutive semitones, the algorithm struggled to deal with these pitch variations and resulted in inaccurate alignments. To account for these inaccuracies in the automatic synchronization, the annotations of the complete corpus

[1] `http://www.tobarandualchais.co.uk/`
[2] `https://github.com/ecanoc/ShetlandFiddles`

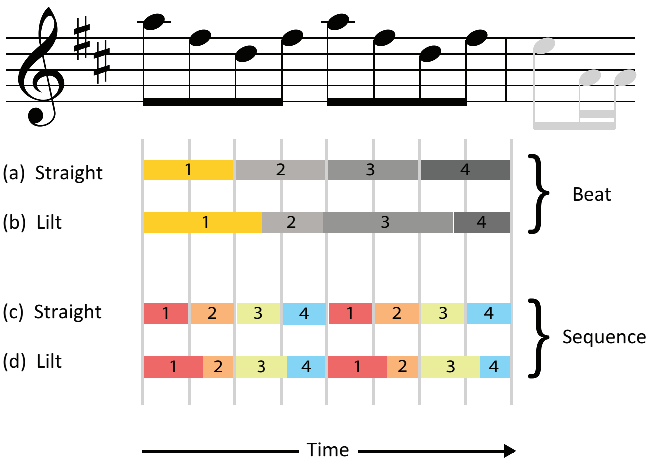

Figure 2: Schematic depiction of our onset-based analysis on a segment of the tune *Clean Pease Strae*. **Beat-level analysis:** (a) Beats of equal length (straight version) (b) Beats with timing deviations (lilt version). **Sequence analysis:** (c) Eighth notes of equal length (straight version) (d) Eighth notes with timing deviations (lilt version).

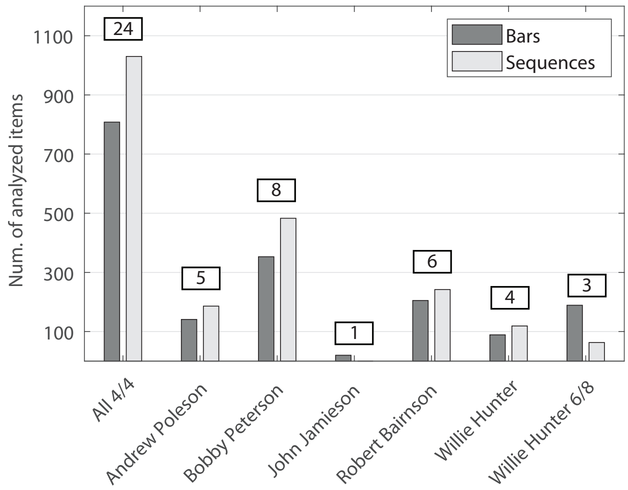

Figure 3: Total number of bars, note sequences and tunes (shown in boxed numbers above) per performer in the Shetland fiddle corpus.

were then manually fine-tuned by two professional musicians using Sonic Visualiser [3] .

3.3 Onset-based Analysis

The synchronized MIDI transcriptions are used to analyze microtiming variation in our corpus. Two different analyses were conducted: (a1) Microtiming variations on the beat level, and (a2) Microtiming variations in eighth note sequences. Given that accurate annotations of the note offsets are not available or easily extracted, we focus on measuring time intervals between onsets in this study. For the beat-level analysis (a1), we calculate time intervals between consecutive beats in the bar. We refer to these intervals as inter-beat-intervals (IBI). For the eighth-note-sequence analysis (a2), we calculate time intervals between consecutive eighth notes in the sequence. We refer to these intervals as inter-onset-interval (IOI).

Our analysis can be better understood by looking at Fig. 2, where a segment of the tune *Clean Pease Strae* is displayed. For the beat level analysis, Fig. 2a shows inter-beat-intervals (IBI) of equal length. This is referred to as the *straight* version of a performance. Our study seeks to understand whether deviations from the straight version exist, and whether performers slightly shorten or stretch the beats when playing with lilt. This scenario is depicted in Fig. 2b, where beats 1 and 3 are slightly longer than beats 2 and 4. Similarly, we analyze sequences of eighth notes in all tunes, and focus on measuring inter-onset-intervals (IOI) for sequences of four eighth notes (for $\frac{4}{4}$ meter), and sequences of three eighth notes (for $\frac{6}{8}$ meter). This can be seen in Figs. 2c and 2d, where the straight and lilt versions of the two eighth note sequences in the bar are shown, respectively. As with the beat-level analysis, we seek to understand whether performers slightly shorten or stretch the eighth notes when playing with lilt. In Fig 2d, the eighth notes 1 and 3 in the sequences are slightly longer than notes 2 and 4.

3.3.1 *Practical Considerations*

To facilitate the analysis of our corpus, we implemented an algorithm for pattern detection in MIDI files. The algorithm includes methods to detect relevant information such as downbeats, rhythm sequences (eighth notes, quarters, etc), double-stops, pitch sequences, etc. Supported by our pattern detection algorithm, the following considerations were taken: (1) Given that our analysis uses note onsets to detect inter-beat-intervals (IBI), only those bars where an onset was present (a note was played) on each of the beats in the bar were considered. While this means that parts of the performance were discarded in our analysis, it also completely removes the need to perform beat annotations, and relies entirely on the performers' onsets to calculate timing. (2) For the eighth-note-sequence analysis, only complete sequences of four eighth notes in $\frac{4}{4}$ bars, and three eighth notes in $\frac{6}{8}$ bars were considered. We restrict our analysis to sequences whose first eighth note coincides with an on-beat. Here again, we completely rely on the performer's onsets to calculate inter-onset-intervals (IOI) and timing variations. (3) A song-wise mean normalization of IBI and IOI was performed for the beat-level analysis and for the sequence analysis, respectively. This normalization allows us to remove differences in beat and note duration dictated by the different tempi of the tunes, and allows us to make direct comparisons between tunes. (4) In our corpus, only three tunes out of the 27 are written in $\frac{6}{8}$. These three tunes are all performed by Willie Hunter, and hence, conclusions with respect to the usage of timing in $\frac{6}{8}$ bars are restricted to the ones revealed by Willie Hunter's performances. (5) In all tunes where double-stops

[3] https://www.sonicvisualiser.org/

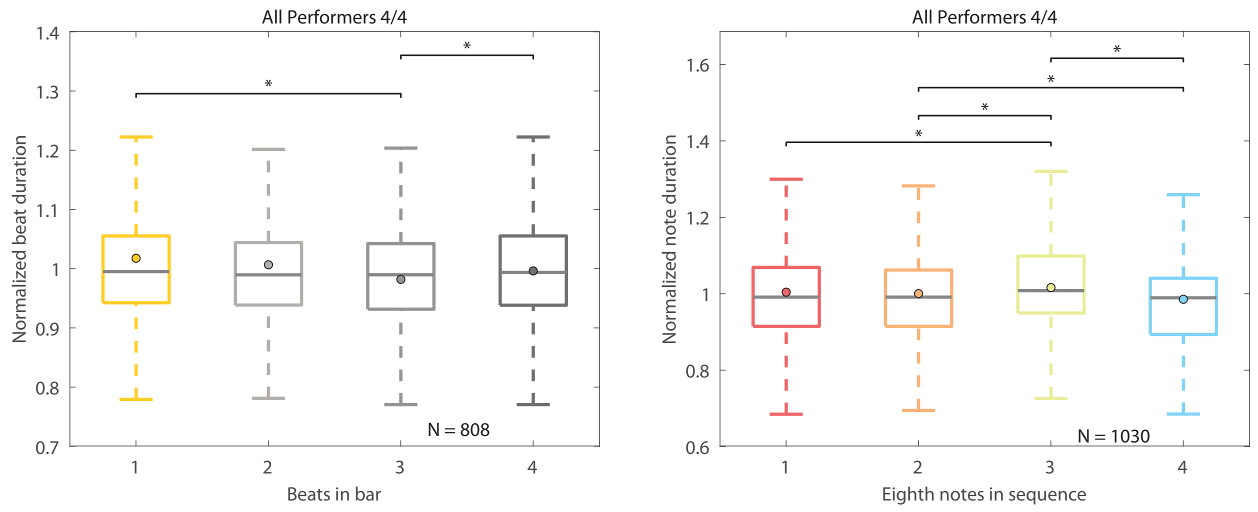

(a) Corpus beat-level results

(b) Corpus note-sequence results

Figure 4: Results for all performers and all tunes in $\frac{4}{4}$ bars in the corpus. (a) Beat-level analysis, and (b) Note-sequence analysis. Significant differences are highlighted (* p-value < 0.05: Wilcoxon signed-rank test for pairwise comparisons)

were detected, only the onset of the lower note was used in our calculations.

Figure 3 shows a summary of the number of tunes, complete bars found and complete sequences found per performers in the corpus using our pattern detection algorithm.

4. RESULTS

4.1 Annotation reliability

To assess the reliability of the annotations, five of the tracks in the corpus were annotated by the two expert annotators. As a measure of annotator agreement, we use Chronbach's alpha coefficient α. For simplicity, we refer to the annotators as a1 and a2 in the following. The agreement between annotators was extremely high with $\alpha(a1, a2) = 0.9995$. It is important to consider, that the starting point for both annotators was the MIDI synchronization obtained with the method proposed in [5]. Hence, we also calculate the Chronbach's coefficient between the automatic MIDI synchronization M, $a1$, and $a2$, resulting in $\alpha(a1, M) = 0.9999$, and $\alpha(a2, M) = 0.9995$. These results validate both the reliability of our final annotations, and the high accuracy obtained by the automatic synchronization in the first place. Only minor modifications were required in order to obtain the final annotations from the automatically synchronized ones. These results also validate the usage of the method in [5] as an efficient approach for MIDI-to-audio synchronization of fiddle recordings, and opens the possibility to further enlarge our corpus in the future.

4.2 Onset-based analysis

In the following sections, we present our results in a top-down manner, starting with trends observed over the entire corpus, followed by a performer-wise analysis, and concluding with a comparative analysis of the tune *Willafjord* (the only tune performed by two perfomers in our corpus). To summarize the results, boxplots are presented both for the beat-level and sequence analyses. In all the figures, the boxplots represent the minimum and maximum values (whiskers), and the first, second (median) and third quartiles (box). In addition, the mean values for each distribution are displayed with colored circles. We present normalized note and beat durations in our plots (see Sect. 3.3.1).Values smaller than 1 show IOIs and IBIs shorter than that of the mean. Similarly, values larger than 1 show IOIs and IBIs longer than that of the mean.

4.2.1 Corpus analysis

Pairwise differences of the mean IOIs and mean IOBs of the entire corpus are calculated using the Wilcoxon signed-rank test at a 5% significance level (p-value < 0.05). Results for the corpus analysis are presented in Fig. 4 where the contribution of all performers and all tunes in $\frac{4}{4}$ meter are displayed.

Figure 4a shows the results of the beat-level analysis for all performers and tunes in $\frac{4}{4}$ bars. The pairwise comparisons revealed the normalized beat duration of beat 3 to be significantly shorter than 1 and 4. The slight shortening of beat 3 across all performers might indicate an attempt to delineate the on-beats in the $\frac{4}{4}$ bar.

Figure 4b shows the results of the sequence analysis. Note 3 in the sequence is significantly longer than all other eighth notes in a sequence. In addition, we see a significant difference between note 2 and note 4, with note 4 slightly shorter than note 2. In contrast, no significant difference was found between notes 1 and 2 in the sequence, appearing to be performed mostly in a *straight* manner .

4.2.2 Performer-wise analysis

In this section, we analyze possible performer idiosyncrasies both on the beat level and on the note sequence

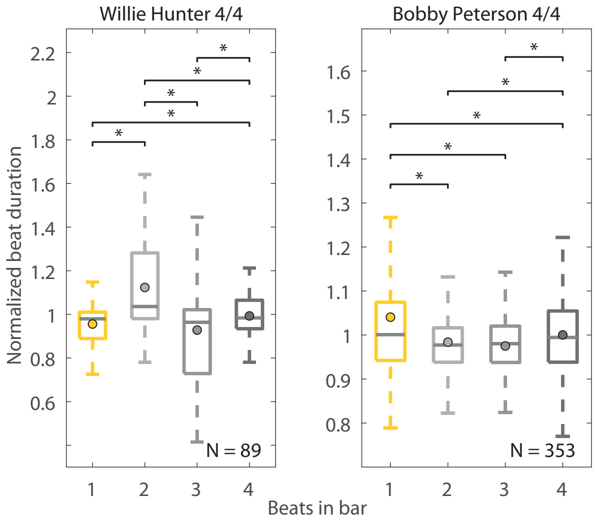

Figure 5: Performer-wise beat-level results. Significant differences are highlighted (* p-value < 0.05: Wilcoxon signed-rank test for pairwise comparisons)

level.

On the beat level, significant differences are evident in only two of the five performers, Willie Hunter and Bobby Peterson, with the other performers mostly showing beats of equal length. As seen in Fig. 5, the general trend in both Hunter and Peterson is to slightly shorten the third beat. These results support the trend observed in the corpus analysis of beats in Fig. 4a. In Willie Hunter, beat 3 is significantly shorter than beats 2 and 4, but beats 1 and 3 appear to be mostly played with the same duration. In contrast, beat 3 in Bobby Peterson is significant shorter than beats 1 and 4, but not than beat 2. Beat 1 in Bobby Peterson appears to be slightly prolonged, showing significant difference with all other beats.

On the note-sequence level, significant differences are evident in five performers as displayed in Fig. 6. John Jamieson only performs one tune in our corpus, and no eighth note sequences were found in it. In this analysis, we also include the results for the three tunes in $\frac{6}{8}$ performed by Willie Hunter. For the $\frac{6}{8}$ tunes, we consider sequences of three consecutive eighth notes. In Willie Hunter's $\frac{6}{8}$ sequences, we see a tendency to slightly shorten note 2. These differences are significant with respect to note 1 and 3. In Hunter's $\frac{4}{4}$ sequences, note 3 is significantly longer than all the others, a pattern that we also observe in Bobby Peterson. However, for Bobby Peterson, we additionally see a slightly shortened note 4 (significant with respect to note 2 and 3), and a tendency to lengthen note 2 with respect to 1. In contrast, Robert Bairnson appears to play mostly straight sequences, with a slightly shorter note 3 with respect to 1. As mentioned by Cooke in [3], the shortening of the 3rd note may be related to bowing: if notes 2, 3, and 4 are played with one bow, the middle note if often shortened and accented. The influence of bowing on lilt falls beyond the scope of this study, but calls for further future investigation. Andrew Poleson on the other hand, shows a consistent shortening of note 2, with significant differences with respect to all other notes.

It appears that a tendency to play pairs of eighth notes (either 1-2 or 3-4) in a long-short (LS) pattern might exist, with Willie Hunter ($\frac{4}{4}$) and Bobby Peterson showing a LS pattern in notes 3-4, and Andrew Poleson in notes 1-2. These results also go in hand with observations made by Cooke in [3].

4.2.3 Comparative analysis of the tune Willafjord

To analyze differences between performers, we compare the two versions of the tune *Willafjord* in our corpus; one performed by Bobby Peterson, the other by Robert Bairnson. Here we focus specifically on the differences at the note sequence level, and perform pairwise comparisons of the four notes sequence. Given that there are differences between the two performances, a direct correspondence between the sequences that we compare cannot be established, and hence, we treat them as unmatched pairs. To do so, the Wilcoxon sum-rank test is performed at a 5% significance level (p-value < 0.05). The tests reveal significant differences in IOIs for all of the eighth note positions, suggesting completely different performing practices by the two fiddlers for the tune *Willafjord*.

Our comparative analysis of the tune *Willafjord* gives a tantalizing insight into the source of such performer idiosyncrasies. These can be further elucidated in the biographies of Bobby Peterson and Robert Bairnson in Peter Cooke's original ethnographic account. Although both men lived on the main island of Shetland (approximately 23 miles apart), Bobby Peterson spent a great deal of his young life as sailor on a whaling ship. There, he was subject to a great many influences including hearing and learning tunes from other sailors from Scotland and Scandinavia. He learned the fiddle mostly be ear, in fact the first tune he ever learned was the focus of our analysis, the tune *Willafjord*. In contrast, Robert Bairnson, although only a short distance away, was fairly isolated as infrastructure on the Shetland Isles in the early 20th century made travelling very difficult. Bairnson also lived in a district where religious attitudes deemed music and dancing sinful. He was taught formally by a local minister, learning notation at the same time. Therefore he lacked the 'indigenous repertory' of other Shetlanders [3, p. 20]. These biographies go some way to explaining why such clear difference can be observed between performers, and in particular, the straight playing of Bairnson versus the more lilting performance of Peterson.

5. CONCLUSIONS AND FUTURE WORK

In this paper we presented microtiming analysis of a corpus of traditional Shetland fiddle music. The corpus, created specifically for this study, was first annotated by expert annotators using automatic MIDI-to-audio synchronization as a starting point. We examined inter-beat-interval timing deviations, as well as inter-onset-interval timing deviations of eighth note sequences to find the elusive qualities of *lilt*. We first analyzed the corpus as a whole, moving to

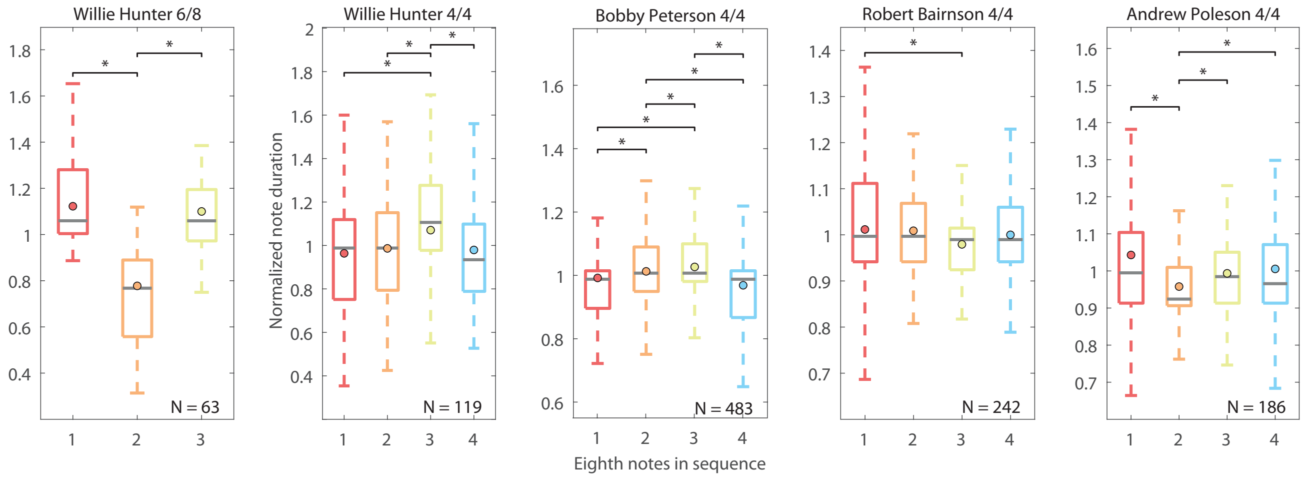

Figure 6: Performer note-sequence results. Significant differences are highlighted (* p-value < 0.05: Wilcoxon signed-rank test for pairwise comparisons)

performer-based analysis, and finalized with a comparative study of a tune played by two performers.

In general, the beat-level analysis indicates a tendency to slight shorten beat 3 in the bar. The note-sequence analysis, suggests the use of long-short patterns in pairs of eighth notes (either 1-2 or 3-4). Even though these emerging patterns in the data may suggest what playing with lilt entails, it is clear that lilt also encompasses idiosyncratic aspect of performance practice. Further investigation is required in order to understand how these patterns relate to other aspects of music performance such as bowing, accentuation, repertoire-specific performing styles, tempo, etc.

With this work, we have reached three important milestones in our research: (1) We have produced high-quality annotations for our corpus that will serve as the foundation and ground-truth for more powerful computational models, (2) we have validated the use of a state-of-the-art MIDI-to-audio synchronization algorithm for the task at hand, and (3) we have revealed microtiming variation patterns in fiddle music from the Shetland Isles.

6. ACKNOWLEDGMENTS

The authors would like to thank the panel of annotators Juan Fernando Alzate and Elissa Sayampanathan, as well as Mr. Ronnie Gibson and Dr. Frances Wilkins at the University of Aberdeen, and Dr. Cathlin Macaulay at the Scottish Studies Archive, School of Literature, Languages and Cultures Celtic & Scottish Studies, University of Edinburgh.

7. REFERENCES

[1] Scott Beveridge, Ronnie Gibson, and Estefanía Cano. Performer profiling as a method of examining the transmission of Scottish traditional music. In *Proc of the 4th International Workshop on Folk Music Analysis (FMA)*, Istanbul, Turkey, 2014.

[2] Matthew W. Butterfield. Why do jazz musicians swing their eighth notes? *Music Theory Spectrum*, 33(1):3–26, 2011.

[3] Peter Cooke. *The Fiddle Tradition of the Shetland Isles*. Cambridge University Press, 1986.

[4] Christian Dittmar, Martin Pfleiderer, and Meinard Müller. Automated estimation of ride cymbal swing ratios in jazz recordings. In *Proc. of the International Conference on Music Information Retrieval (ISMIR)*, pages 271–277, Malaga, Spain, 2015.

[5] Sebastian Ewert, Meinard Müller, and Peter Grosche. High resolution audio synchronization using chroma onset features. In *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 1869–1872, Taipei, Taiwan, 2009.

[6] Anders Friberg and Andreas Sundströöm. Swing ratios and ensemble timing in jazz performance: Evidence for a common rhythmic pattern. *Music Perception: An Interdisciplinary Journal*, 19(3):333–349, 2002.

[7] Fabien Gouyon. Microtiming in "samba de roda"—preliminary experiments with polyphonic audio. In *Proc. of the XII Simpósio da Sociedade Brasileira de Computação Musical São Paulo*, pages 197–203, Sao Paulo, Brazil, 2007.

[8] Rainer Polak. Rhythmic feel as meter: Non-isochronous beat subdivision in Jembe music from Mali. *Music Theory Online*, 16(4), 2010.

[9] Vincent Rosinach and Caroline Traube. Measuring swing in irish traditional fiddle music. In *Proc. of the International Conference on Music Perception and Cognition*, pages 1168–1171, Bologna, Italy, 2006.

[10] Anja Volk and Bas de Haas. A corpus-based study on ragtime syncopation. In *Proc. of the International Conference on Music Information Retrieval (ISMIR)*, pages 163–168, Curitiba, Brazil, 2013.