An Electronic Bagpipe Chanter for Automatic Recognition of Highland Piping Ornamentation

Duncan W. H. Menzies
Centre for Digital Music
Queen Mary University of London
London, UK
d.w.h.menzies@eecs.qmul.ac.uk

Andrew McPherson
Centre for Digital Music
Queen Mary University of London
London, UK
andrewm@eecs.qmul.ac.uk

ABSTRACT

The Highland piping tradition requires the performer to learn and accurately reproduce a diverse array of ornaments, which can be a daunting prospect to the novice piper. This paper presents a system which analyses a player's technique using sensor data obtained from an electronic bagpipe chanter interface. Automatic recognition of a broad range of piping embellishments allows real-time visual feedback to be generated, enabling the learner to ensure that they are practicing each movement correctly.

The electronic chanter employs a robust and responsive infrared (IR) sensing strategy, and uses audio samples from acoustic recordings to produce a high quality bagpipe sound. Moreover, the continuous nature of the IR sensors offers the controller a considerable degree of flexibility, indicating significant potential for the inclusion of extended and novel techniques for musical expression in the future.

Keywords

Great Highland Bagpipe, continuous infrared sensors, ornament recognition, practice tool, SuperCollider, OSC.

1. INTRODUCTION

For centuries, the Great Highland Bagpipe (GHB) has sustained a deep military connection, and the discipline associated with all aspects of the armed forces is strongly reflected in the Highland piping tradition. Even reasonably simple bagpipe music requires the performer to know a wide variety of distinct and clearly defined techniques, and to execute them with meticulous precision. This formal musical culture, coupled with the level of physical endurance required to produce a consistent sound, can cause the learning curve to be intimidatingly steep to the beginner.

In recent years, there has been a wealth of research into the development of tools to assist musicians in instrumental practice, much of which has focussed on the piano and other keyboard instruments. However, it seems that bagpipe pedagogy, at least in a technological context, remains largely unexplored.

This paper presents a system which exploits the regimented, rule-based nature of Highland piping to analyse a player's performance using sensor data obtained from an electronic bagpipe interface. Real-time visual feedback allows the learner to ensure that they are practicing each technique correctly. This is not intended as a replacement for an experienced human instructor, but rather as a supplementary tool to help the student maximise the efficiency of their practice time, and to avoid the introduction of bad habits between lessons.

NIME'12, May 21 – 23, 2012, University of Michigan, Ann Arbor.

2. BACKGROUND

2.1 The Great Highland Bagpipe (GHB)

The GHB is comprised of five individual pipes (two tenor drones, one bass drone, the blowpipe and the melody pipe or chanter) attached to an airtight bag, which acts as an air reservoir. The bag is held under the player's arm, with the drones resting over the shoulder. By maintaining a constant pressure on the bag with the elbow, the player ensures a steady flow of air through the drone and chanter reeds.

It is worth noting that the level of air flow and pressure on the bag are generally not used as expressive parameters, but rather kept as consistent as possible in order to produce a continuous, unbroken sound. The bass and tenor drones produce a uniform octave accompaniment which remains unchanged regardless of the melody. Therefore, at least in traditional piping, the chanter is the sole means of expressive control of the GHB.

2.1.1 The Chanter

Figure 1 shows the nine notes of the GHB scale and their associated fingerings. This diagram raises two important details. Firstly, the GHB is a transposing instrument, sounding approximately a semi-tone above the notated pitch [3]. For the remainder of this paper, notes will be referred to by their piping names as opposed to concert pitch.

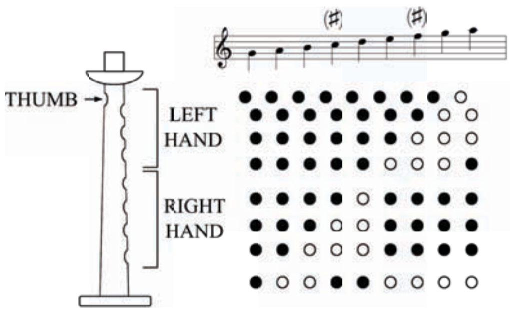

Figure 1: Highland bagpipe scale and fingerings.

Furthermore, although no key signature is given (as is typical of all GHB music), the C and F of the scale are in fact C♯ and F♯. The chanter can therefore reproduce a full *Mixolydian* scale, with an additional 7th (low G) below the tonic (A). "Cross fingerings" for C♮ and F♮ have been developed in recent years, and are common among many contemporary pipers. However, these have not been accepted into traditional piping circles, and hence the standard fingerings are invariably referred to simply as "C" and "F".

2.1.2 Playing Technique and Ornamentation

In addition to its restricted melodic range, the GHB provides no facility for dynamic control, and hence has only one (not inconsiderable) volume level. Moreover, the fact that both drones and chanter produce a constant, uninterrupted sound prevents the use of silences or timbral variations for the purposes of emphasis and articulation.

To address these limitations, pipers employ a variety of ornaments to separate and accentuate the melody notes. These take the form of one or more short gracenotes, performed in a specific order. A wide range of ornaments exist, from individual gracenotes to more elaborate embellishments such as the *birl* and *crunluath* (Figure 2). These are rigorously and formally defined, and their correct execution is an integral aspect of piping technique [11].

Figure 2: Examples of Highland bagpipe ornamentation: (a) G gracenote, (b) crunluath, (c) birl, (d) doubling on C.

The aspiring piper must therefore invest considerable effort to learn and master at least the basic ornaments before attempting all but the simplest tunes. This process can often take six to twelve months of regular and disciplined practice, and is particularly challenging in situations where lessons with a tutor are infrequent.

Moreover, to the untrained ear, it can be somewhat difficult to discern whether or not an ornament was executed correctly. In the absence of an experienced instructor to provide immediate criticism in such instances, the student's inability to evaluate their own technique accurately can lead to wasted practice time and the introduction of bad habits.

The approach presented in this work aims to recognise when a user attempts a particular embellishment, and to provide real-time visual feedback on their performance. This would allow the player to gauge their progress and to correct errors quickly, in order to maximise the efficiency of time spent practicing between lessons.

2.2 Related Work

2.2.1 Electronic Bagpipes

Several brands of electronic Highland bagpipes are commercially available, of which the DegerPipes[1], TechnoPipes[2] and Redpipes[3] are most prominent. These use single capacitive touch-switches in place of the finger-holes, and create the audio output by wavetable synthesis. Such products have found moderate acceptance among modern folk/fusion groups. However, while the synthesised sound produced is often acceptable in the context of a live show by a full band with drum kit and electric guitars, it falls short of the fidelity required for solo performance or recording.

Furthermore, the capacitive sensors are binary in nature; the "holes" are always either fully open or closed. This does not accurately reflect the finger-holes of an acoustic chanter, which can be gradually covered and uncovered to slide between notes. While this practice is absent (and indeed frowned upon) in traditional piping circles, it is nonetheless a popular technique among contemporary performers. Moreover, and perhaps more importantly, such sensors are unable to detect continuous states in the brief intermediate period as a finger is quickly lifted or replaced. This leads to unnaturally abrupt transitions between notes at the micro level, which undoubtedly has a detrimental effect on the perceived authenticity of the audio output.

There have been several attempts within the academic community to develop alternatives to this discrete sensor strategy. The EpipE [2, 8] is a uilleann bagpipe chanter interface, which extends the capacitive sensing approach to include an array of sixteen small binary touch-switches for each hole. The number of electrodes was chosen such that for each individual sensor, the corresponding change in frequency would be less than the Just Noticeable Difference (JND) discernible by the human ear [13], giving the impression of continuous frequency variation.

The FrankenPipe [9] instead uses photoresistors mounted inside the holes of an acoustic GHB chanter. This provides a wide analogue range for each hole, and has the advantage of retaining the physical feel of a traditional chanter.

2.2.2 Musical Practice Tools

Recent studies have investigated the use of visual feedback to assist in the process of practicing a musical instrument. Goebl and Widmer [6] propose an approach to recognise and evaluate certain "well-defined sub-aspects" of piano playing using a MIDI keyboard. This system can detect recurring patterns (e.g. Alberti bass passages) and provide a real-time display of deviations in pitch, synchrony and dynamics.

The *MIDIator* tool [12] also captures piano performance via MIDI, and compares this "expressive score" with an existing template, in order to provide a quantitive summary of the differences in tempo, volume, duration and articulation. Other research has directly examined the raw audio signal produced by an acoustic instrument to provide real-time visualisation of "musical expression" (in terms of variations in tempo and dynamics [4]), and of sonic features such as harmonic content, noisiness, loudness and fine pitch [5].

In addition to musical concepts, the physical aspects of music pedagogy have also been explored. The SYSSOMO system [7] uses a score-following algorithm to align graphical information about a pianist's arm movements (obtained using wireless motion sensors) with a piano roll representation of their performance. The approach presented in [1] employs a similar sensor configuration to detect the physical gestures of a student "conducting" along to a piece of music. These are then compared to pre-recorded patterns and identified in real time using a method based on time warping and Hidden Markov Models.

3. ELECTRONIC CHANTER INTERFACE

3.1 Hardware

The purpose of the chanter interface is to detect the performer's continuous finger movements quickly and accurately, and to send this data to the audio software via USB. The physical playing experience should be as similar as possible to an acoustic chanter, and the sensors should be robust to variations in ambient light and skin dryness (the latter being a common source of error in capacitive touch-sensors).

The sensing system employed in this work consists of an infrared (IR) emitter (LED) and detector (photodiode) for each hole, between which a constant IR beam exists. When the beam is interrupted by the player's finger, a continuous change in the voltage V_{out} can be measured.

The 16MHz ATmega328 microcontroller polls the eight input pins, and transmits the raw sensor data to the host computer via USB at a baud rate of 115200 bit/s. This configuration does not introduce any audible latency, allowing a degree of musical intricacy comparable to that of a traditional bagpipe chanter.

[1] http://www.deger.com/

[2] http://www.fagerstrom.com/technopipes/

[3] http://redpipes.eu/

3.2 Sound Generation

The audio generation software is written in the SuperCollider programming language[4]. While existing electronic chanters produce a synthesised output, this system instead uses sampled bagpipe recordings which have been carefully truncated to $\simeq$ 10 ms such that they can be looped indefinitely.

Since the characteristic waveform produced by an acoustic chanter is by nature distinctly stationary, this approach provides a highly convincing GHB sound, and because no synthesis is involved, the computational workload of the program is modest. Moreover, the samples require only 250 KB of memory in uncompressed .wav format, and hence could easily be stored on an embedded device in the future.

Once the sound has been generated, the current note being played is communicated to the ornament recognition software via the Open Sound Control (OSC) protocol[5]. The sound is produced prior to the analysis and visualisation stage (Figure 3) to avoid the introduction of any unnecessary latency in the audio output.

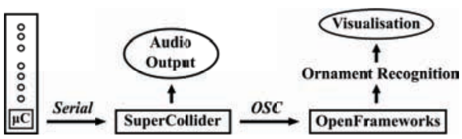

Figure 3: Overview of electronic chanter system.

4. ORNAMENT RECOGNITION AND VISUALISATION

The ornament recognition and visualisation software is written in C++ using the OpenFrameworks toolkit[6]. Pitch data is received from SuperCollider and displayed on a stave in a rolling format from right to left, as shown in Figure 4(a).

For the purposes of this work, an ornament is defined as a series of one or more gracenotes in a specific sequence. A gracenote is taken to be a note whose length is less than or equal to a specified limit L. At present, an experimentally determined value of $L = 17$ samples ($\simeq$ 85ms) is used. Each ornament has its own particular subset of possible previous and following notes. For example, a *throw on D* can come from any note except low G, but must end on D.

Based on these principles, a collection of 54 templates corresponding to the most common ornaments was produced. The algorithm assumes any consecutive string of gracenotes to be a potential embellishment. Upon completion of the sequence (i.e. when a note longer than L is detected) the recognition method is called.

Firstly, the notes immediately before and after the detected ornament are considered. Any template for which either of these is disallowed is removed from the comparison database. The ornament is then compared to the remaining templates, and assigned a similarity rating in each case.

The template with the highest score (provided this is above a predefined minumum value) is returned as the most likely outcome. When an embellishment is executed correctly, the green notes in the pitch visualisation turn blue, and the ornament name is displayed above the stave. Figure 4(b) illustrates this for the G gracenote and crunluath movements (notated in Figure 2 (a) and (b) respectively).

4.1 Detecting Ornamentation Errors

In instances where no exact match is found (i.e. the player has made a mistake), the system searches for the attempted embellishment by considering three possible classes of error. These can be defined by the number of notes in the detected ornament relative to the template, and form the three sections of the comparison algorithm.

4.1.1 Length of ornament = length of template

The abundance of different embellishments in traditional piping can lead to one or more gracenotes in a sequence being remembered wrongly. If the movement is otherwise correctly executed, the performed ornament will have an equal number of notes to the required template. The similarity rating between ornament and template of equal length is therefore simply assigned according to the number of note indices at which the corresponding pitches match.

When an inconsistency is detected, the erroneous notes are coloured red in the visualisation, and the correct pitches displayed beneath the stave. Figure 4(c) illustrates this for the case of a birl (see Figure 2(c)), where the opening high G gracenote has been mistakenly replaced by an F.

4.1.2 Length of ornament < length of template

Perhaps the most common errors in ornamentation, particularly among inexperienced pipers, occur when the player fails to move their fingers with sufficient speed and precision to articulate the movement clearly. This results in the omission of one or more gracenotes, causing the detected ornament to be shorter than the appropriate template.

In such cases, the algorithm compares corresponding note indices until a discrepancy is reached (i.e. a gracenote in the template is missing from the performed embellishment). The pitch and position of the absent note are recorded, and the comparison continues from the subsequent point in the template sequence. Figure 4(d) shows the visualisation of a crunluath with a missing E gracenote.

4.1.3 Length of ornament > length of template

Another recurring problem among novice players is the introduction of unwanted additional notes caused by careless transitions between fingerings. These are referred to as "crossing noises" and are particularly common when changing back and forth between the top hand notes (high A to E) and the bottom hand (D to low G).

The inclusion of crossing noises within a gracenote sequence will result in the detected embellishment having more notes than the correct template. In this situation, the comparison process is complementary to that of 4.1.2, with any pitch not present in the template being deemed erroneous. Figure 4(e) depicts a *doubling on C* ornament (Figure 2(d)) preceded by a crossing noise, marked "X".

However, crossing noises are not only a problem within embellishments, but can also be present between melody notes. This is considered very bad technique, and should be carefully avoided [11]. It is thus beneficial that the algorithm highlight crossing noises whenever they occur.

Crossing noises are generally extremely short, and can be difficult for inexperienced pipers to detect. Experimentation with the electronic chanter suggests that true crossing noises (as opposed to simply wrong notes or errant gracenotes) are usually no longer than $\simeq$ 20ms. Based on this assessment, any individual note whose duration is less than 20ms is marked as a crossing noise, as shown in figure 4(f).

5. DISCUSSION AND FUTURE WORK

Preliminary testing indicates that the algorithm is effective and consistent in the recognition of correctly executed Highland piping ornamentation. Moreover, for a wide range of common errors, it can successfully identify the attempted movement and provide real-time visual feedback on how to

[4]http://supercollider.sourceforge.net/
[5]http://opensoundcontrol.org/
[6]http://www.openframeworks.cc/

(a) Simple arpeggio.

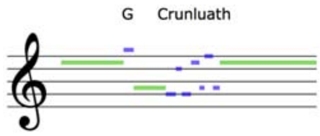

(b) Correctly executed ornaments.

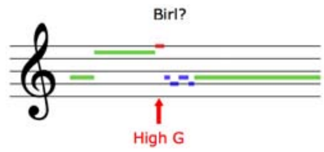

(c) Birl with erroneous F gracenote.

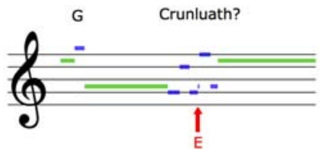

(d) Crunluath with missing E gracenote.

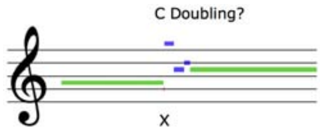

(e) Doubling on C with crossing noise.

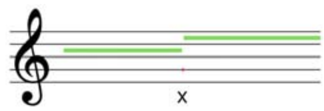

(f) Crossing noise between two melody notes.

Figure 4: Examples of visual feedback system.

achieve the desired outcome.

An extensive user study, involving both beginners and experienced pipers, is intended for the near future. It is hoped that this will provide a deeper understanding of the capabilities and limitations of the system, and the extent to which such a tool could be useful in a pedagogical context.

A number of potential improvements to the ornament recognition algorithm can already be envisaged. For example, there are instances (albeit comparatively few) in which a piper might wish to prolong a particular gracenote within an embellishment for emphasis. To allow for such cases, a template matching approach based on dynamic time warping (DTW) could be employed. Paulus and Klapuri [10] use DTW to determine the similarity between temporal rhythmic patterns. By incorporating a similar method into the existing ornament recognition system, the player could evaluate not only how accurately a given embellishment was executed, but also how their performance compares on a more expressive level to that of a virtuoso piper.

Furthermore, the algorithm could be extended to include the raw sensor data, as opposed to considering only the final pitch. This would allow greater insight in cases where one or more gracenotes in a sequence is missing. By examining the continuous finger positions, the system could distinguish whether the player narrowly failed to articulate each note, or if certain finger movements were omitted altogether.

As regards the electronic chanter itself, the next development will be to produce a standalone controller and sound generation device which could be used in a live performance situation without a host computer. Moreover, to be of practical use in a recording context, the sound produced would need to be virtually indistinguishable from that of an acoustic chanter. For this reason, a listening test could be carried out to compare the chanter interface with existing electronic bagpipes and their acoustic counterparts, in order to determine the perceived authenticity of the audio output.

This work has focussed predominantly on traditional Highland piping practice. However, the continuous infrared sensing strategy offers significant potential for the development of extended and novel techniques to broaden the expressive capabilities of the bagpipe chanter in the future.

6. ACKNOWLEDGMENTS

This work was funded by the Engineering and Physical Sciences Research Council (EPSRC) as part of the Centre for Doctoral Training in Media and Arts Technology at Queen Mary University of London.

7. REFERENCES

[1] F. Bevilacqua, F. Guédy, N. Schnell, E. Fléty, and N. Leroy. Wireless Sensor Interface and Gesture-Follower for Music Pedagogy. In *Proc. NIME 2007*, pages 124–129, New York, NY, USA, June 2007.

[2] C. Cannon, S. Hughes, and S. Ó Modhráin. EpipE: Exploration of the Uilleann Pipes as a Potential Controller for Computer-based Music. In *Proc. NIME 2003*, pages 3–8, Montreal, Canada, May 2003.

[3] J. Dickson. *The Highland Bagpipe: Music, History, Tradition.* Ashgate Publishing, Ltd., 2009.

[4] S. Dixon, W. Goebl, and G. Widmer. Real Time Tracking and Visualisation of Musical Expression. In *Proc. Music and Artificial Intelligence*, pages 63–71, Edinburgh, Scotland, September 2002.

[5] S. Ferguson, A. Vande Moere, and D. Cabrera. Seeing Sound: Real-time Sound Visualisation in Visual Feedback Loops used for Training Musicians. In *Proc. Information Visualisation*, pages 97–102, London, England, July 2005.

[6] W. Goebl and G. Widmer. Unobtrusive Practice Tools for Pianists. In *Proc. Music Perception and Cognition*, Bologna, Italy, August 2006.

[7] A. Hadjakos, E. Aitenbichler, and M. Mühlhäuser. SYSSOMO: A Pedagogical Tool for Analyzing Movement Variants Between Different Pianists. In *Proc. Enactive International Conference*, Pisa, Italy, November 2008.

[8] S. Hughes, C. Cannon, and S. Ó Modhráin. Epipe: A Novel Electronic Woodwind Controller. In *Proc. NIME 2004*, pages 199–200, Hamamatsu, Japan, June 2004.

[9] T. Kirk and C. Leider. The FrankenPipe: A Novel Bagpipe Controller. In *Proc. NIME 2007*, pages 301–304, New York, NY, USA, June 2007.

[10] J. Paulus and A. Klapuri. Measuring the Similarity of Rhythmic Patterns. In *Proc. ISMIR 2002*, pages 150–156, Paris, France, October 2002.

[11] Shepherd, R.T. *Learn to Play the Bagpipe: A New Method of Practice Chanter Tuition.* Mel Bay Publications, 2002.

[12] S. Shirmohammadi, A. Khanafar, and G. Comeau. MIDIATOR: A Tool for Analysing Students' Piano Performance. *Revue de Recherche en Éducation Musicale*, 24:35–48, 2006.

[13] E. Zwicker, G. Flottorp, and S. S. Stevens. Critical Band Width in Loudness Summation. *JASA*, 29:548–557, 1957.