A corpus-based method for controlling guitar feedback

Sam Ferguson, Andrew Johnston
Creativity and Cognition Studios
Faculty of Engineering & IT
University of Technology, Sydney
{samuel.ferguson,andrew.johnston}@uts.edu.au

Aengus Martin
Faculty of Arts and Social Sciences
University of Technology, Sydney
aengus.martin@uts.edu.au

ABSTRACT

The use of feedback created by electric guitars and amplifiers is problematic in musical settings. For example, it is difficult for a performer to accurately obtain specific pitch and loudness qualities. This is due to the complex relationship between these quantities and other variables such as the string being fretted and the positions and orientations of the guitar and amplifier. This research investigates corpus-based methods for controlling the level and pitch of the feedback produced by a guitar and amplifier. A guitar-amplifier feedback system was built in which the feedback is manipulated using (i) a simple automatic gain control system, and (ii) a band-pass filter placed in the signal path. A corpus of sounds was created by recording the sound produced for various combinations of the parameters controlling these two components. Each sound in the corpus was analysed so that the control parameter values required to obtain particular sound qualities can be recalled in the manner of concatenative sound synthesis. As a demonstration, a recorded musical target phrase is recreated on the feedback system.

1. INTRODUCTION

Electric guitar feedback provides a rich textural palette for music composition and it has been used in a wide variety of musical contexts. Commonly cited examples of users of electric guitar feedback are Jimi Hendrix in popular music, and Glenn Branca in contemporary art music. Feedback which is not guitar-related has also been by composers such as Behrman and and Tudor [1].

There are a number of different mechanisms by which electric guitar feedback occurs. The most common involves a guitar string vibration being transduced into an electrical signal, which is then amplified and emitted by a loudspeaker as an acoustic signal, which in turn vibrates the guitar string. If the net gain in this closed circuit is greater than unity, the signal will grow in amplitude until some part of the signal path saturates and the signal reaches a steady state.

Common to all musical applications is the need for a method for manipulating the characteristics of the feedback sound. Guitar players can coarsely control the feedback system by (i) adjusting the distance between the guitar and the amplifier, (ii) changing the effective lengths of the guitar strings, and (iii) using digital audio effects. However, it is not common for guitarists to use feedback in contexts requiring precise control, for instance, when accurate pitches are required.

Feedback is difficult to control for the following reasons:

NIME'13, May 27 – 30, 2013, KAIST, Daejeon, Korea.

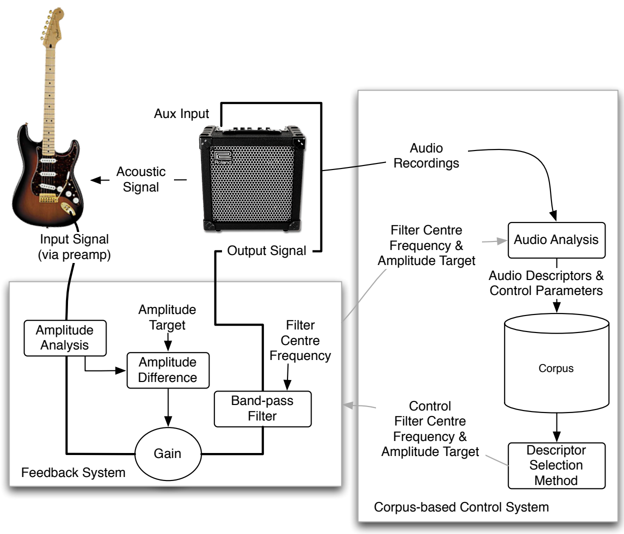

Figure 1: Overview of feedback control system described in this paper.

- Multiple feedback pitches are possible with one string and fret position, due to resonance of higher harmonics.
- The gain in the feedback loop (per string) changes according to the positions and orientations of both the amplifier and the guitar, and is therefore difficult to control precisely as the guitar is often strapped to a moving performer.
- The pitch produced by the system is rarely independent of the gain of the system, meaning that higher pitches are more likely at higher gains.

This paper will introduce a method for improving control of the loudness and pitch of feedback created by an electric guitar and amplifier.

1.1 Previous Research

Models of musical feedback systems have been studied by Morris [2]. His formalisation describes four main elements of a feedback instrument – the feedback:

loop : which can be electrical, acoustic, or digital;

intervention : which is used to transform the feedback and gain some kind of control over feedback;

interruption : which is the method for discontinuing the feedback resonance;

excitation : which is the source of energy used to excite the system to begin the feedback.

Berdahl et al. have developed a series of feedback systems that combine stringed instruments with actuators that apply force to the string without physical contact. Applications of this method include damping a string without contact [3] and the application of 'physical audio effects' to the string output [4]. Later work relates to the dislocation of the musician from the instrument, and involves 'mechanical sound synthesis' [5]. Finally, in recent work Berdahl et al. describe an entire framework for 'feedback control of acoustic musical instruments' [6], giving details of methods and approaches that can be taken for a variety of goals.

Other researchers have also investigated feedback control methods. For example, Burns reports on a variety of idiosyncratic, controlled feedback systems in the context of contemporary music performance [7, 8]. In addition, Di Scipio introduced the idea of 'eco-systemic' signal processing, where a musical system becomes an integral part of a musical eco-system with its surroundings, as opposed to the more regular notion of 'interactive' musical systems where the relationship is one of cause and effect [9].

In previous work, we have investigated the use of an electromechanical guitar control system to precisely control feedback sounds, in conjunction with a corpus-based approach by which a library of available sounds was created and used for later recall [10]. The control system could be used to precisely set effective string lengths and to selectively dampen the strings (thus preventing them from resonating). We systematically audited the sounds which could be produced, by making a recording of the feedback produced for different configurations of the electromechanical control system. A database, or *corpus* of available sounds was created by calculating a set of numerical, psychoacoustically-informed descriptors for each sound and storing the mappings between the control configurations and descriptor vectors. While this system was capable of producing feedback sounds which could be recalled by their descriptors, it was limited by the lack of dynamic control over the sounds available. In addition, a considerable amount of noise arose from the movement of a motorised slide on the strings. This noise was undesirable in itself (because it could not be controlled), and it also gave rise to hysteresis effects, meaning that the feedback sound produced was, in some circumstances, dependant on the noise of the slide movement which was itself dependent on the starting position of the slide.

1.2 Overview

In this paper we investigate an alternative approach for producing and controlling feedback, using a signal processing stage to control the feedback loop, and obtaining the signal processing parameters required for a particular sound quality by referencing a pre-calculated look-up table. We use a standard electric guitar and amplifier, but inserted into the feedback signal chain is a simple DSP unit which manipulates the feedback signal to achieve different feedback pitches and steady-state amplitudes (see Figure 1). The DSP used is related to the automatic equalisation methods used for feedback suppression in public address systems (see, e.g. [11] for a review), but we use it for the purposes of control and production, rather than for suppression. The system is easy to assemble, requiring only an unmodified electric guitar, a standard guitar amplifier, a computer and a low-latency audio interface. It is significantly more robust than our previous electromechanical controller.

This remainder of paper is structured as follows. In Section 2, we describe the guitar feedback system. In Section 3, we describe the creation and analysis of a sound corpus required to obtain a look-up table that maps sound qualities (pitch and loudness) to DSP control parameters. We also describe the automatic re-creation of a target musical phrase from electric guitar feedback and discuss the benefits in general of a corpus-based approach. Finally, we discuss the system in the context of related work and conclude.

2. FEEDBACK SYSTEM DESCRIPTION

In this section, we present the details of the electric guitar feedback system, comprising a standard electric guitar and guitar amplifier; and a computer with an audio interface. We focus on controlling the pitch and steady-state amplitude of the feedback sound, as

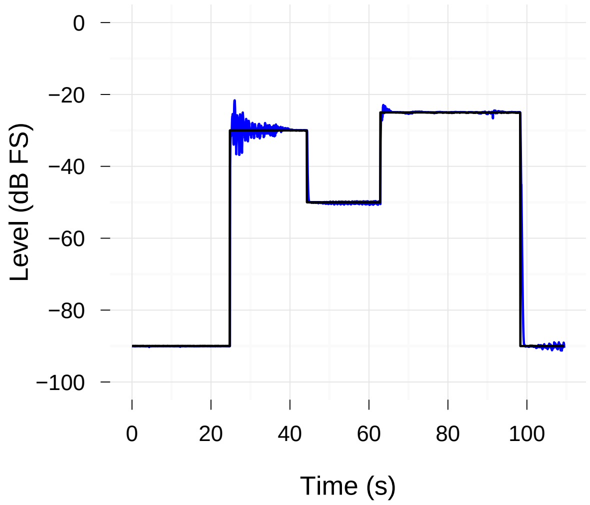

Figure 2: *Controlling the amplitude of the feedback signal using automatic gain control. The black line is the output gain parameter that is sought, and the blue line shows the response of the system.*

these are, arguably, the two most important attributes for musical applications. However, the system may be adapted for controlling other characteristics of the feedback signal (see later sections).

2.1 Construction

First, since electric guitar feedback is sensitive to the relative positions and orientations of the guitar and amplifier, the two were firmly mounted into a metal frame (see Figure 3). A standard instrument lead connects the guitar output to the pre-amplified input of the audio interface, and a second lead connects the output of the audio interface to the amplifier (see Figure 1). This particular amplifier (Roland Cube-30X), has an auxiliary input, allowing a line level output from the audio interface to be used.

2.2 Signal processing

Usually, when feedback occurs, there is a transient stage to begin during which the signal grows in amplitude until it reaches a steady state. The fundamental frequency of the steady-state signal is determined by the frequency component that increased in amplitude most rapidly during the transient stage. In order to control the fundamental frequency of the steady-state signal, we incorporate a bandpass filter into the feedback loop. This increases the gain of a target frequency (the centre-frequency of the filter) with respect to all other frequencies, so that the target frequency is more likely to be the one that grows fastest in amplitude during the transient stage. Note that while this does achive a certain level of pitch control, it does not guarantee that the target frequency will become dominant since the string may not resonate naturally at this frequency.

To control the overall amplitude of the steady-state signal, we introduce an automatic gain control (AGC) into the feedback loop. The automatic gain control works by measuring the difference between the incoming RMS amplitude of the signal and a target RMS amplitude, and adjusting the gain accordingly in a feedback loop. This signal processing is implemented using the Max/MSP interactive audio processing platform[1].

In Figure 2, we show how the level of the feedback signal varies as the target level is changed, while the centre frequency of the bandpass filter is kept constant. After changes in the target level, the feedback signal can sometimes go through a transient oscillatory stage before reaching the target amplitude, seen especially at

[1] http://www.cycling74.com

Figure 3: *The metal frame used for mounting the electric guitar and amplifier. A second guitar and amplifier system is mounted behind the first, but is not used in this study.*

approximately 25 s. In two regions of the plot, between approximately 45 s and 65 s, and after 100 s, the amplitude of the feedback signal oscillates around the target. This type of amplitude oscillation is a characteristic of electric guitar feedback, and is one of its interesting timbral features.

The characteristics of these oscillations (both transient and steady-state) can be controlled by limiting the range of gains that the AGC can apply. Allowing high gains can result in very fast and direct gain change (as seen in Figure 2), but can also result in oscillation after rapid change. In contrast, limiting the AGC to small gains (e.g. ± 6 dB) results in a slower ramp to the target level and only small oscillations when it is reached. In certain musical contexts, a range of gain change limits may be of creative interest.

3. CORPUS ANALYSIS PROCESS

A basic requirement of a musical instrument is that it provides a predictable relationship between the control parameters and the output sound. This requirement is satisfied by the guitar feedback system described in the previous section. However, the relationship between the control parameters and output sound is not a simple one since, as mentioned above, the signal processing components give only limited control over the fundamental frequency and amplitude of the steady-state feedback signal. In fact, the mapping between the control parameters and the characteristics of the output sound is non-linear and discontinuous.

While a performer could in theory learn the complex mapping between control parameters and the characteristics of the output sound (as is necessary for many traditional musical instruments), the difficulty in doing so is exacerbated by the sensitivity of the mapping to the physical characteristics of the feedback system; slight changes to the configuration of the system or the acoustic space in which it is used may result in large changes to the mapping. This constitutes a significant barrier to the use of the guitar feedback system in practise.

To address this issue, we create an interface for the performer that does not require a complex mapping to be learnt. To begin, we enumerate a set of target gain and filter frequency combinations (control parameter settings). We then record the sound output for each one. Each sound is analysed, to derive a set of *acoustic descriptors*, which describe the acoustic results of each control parameter setting. Thus, a look-up table can be created that maps each sound that can be produced by the system (as represented by its acoustic descriptors) back to the control parameter setting required to produce it.

This approach is based on research by Schwarz et al. on concatenative synthesis [12, 13, 14], and we have used a similar method in previous research [10]. We follow Schwarz' use of the term 'unit' to describe a segment of audio that is associated with metadata (i.e. it has been analysed and has acoustic descriptors attached). Our units have a duration of 3 seconds and along with the acoustic descriptors we include the control parameters used to create the sound in the associated metadata.

The approach described is summarised in Figure 4. The corpus creation process is as follows.

- Step through a set of discrete filter centre frequencies and target levels, and for each one, make a 3-second recording of the output audio. Leave a 6 second gap after each recording to allow the feedback to decay.

- Import both the audio files and the control parameter values into the `catart` software [13] (Version 1.2), and analyse each one as a unit. (For the analysis, omit the first second of each recording so that only the steady-state feedback signal is analysed.)

While in this paper we focus mainly on pitch and level descriptors, this method allows a wide variety of descriptors to be calculated for each unit. This includes pitch, loudness, spectral centroid, high-frequency energy and others (again see [13] for details).

During performance, a musician can select sounds by their acoustic descriptors. For each selection, the required control parameters (target gain and filter centre frequency) are recalled and applied so that the sound is produced. Specifically,

- A boolean query is used to select a unit from the look-up table that matches a particular acoustic descriptor value (eg. pitch between 430 and 450 Hz). Then,

- The control parameters used to create this unit are retrieved, and are then

- Used to recreate the sound.

This process can be automated to produce sequences of feedback sounds to match a 'target' sound, using the `catart` software (again, see [13]). Note that we do not use the concatenative synthesis capabilities of this software (which uses concatenation of audio samples from the corpus), but just the corpus analysis, storage, data retrieval, and unit selection algorithm features.

We have implemented the two procedures described in this section, as well as the signal processing required to control the feedback, in a single piece of software. An important feature of the software is that corpus creation, including sound recording and analysis, is automated. This is essential to the use of the system in practise, since a new corpus must be created each time the feedback system is modified (either by being moved to a new acoustic environment or by using a different guitar or amplifier).

3.1 Corpus Example

In this section, we describe the creation and analysis of a corpus for a particular guitar and amplifier combination. The control parameter space used extended across a level range from -42 dBFS to -12 dBFS in 6 dB steps (hence 6 levels), and band-pass filter centre frequencies from 110 Hz to 1661 Hz in semitone steps across the frequency range (48 frequencies), creating 288 control parameter settings. The corpus of 288 3-second recordings took approximately 48 minutes to record with a 6-second gap between recordings.

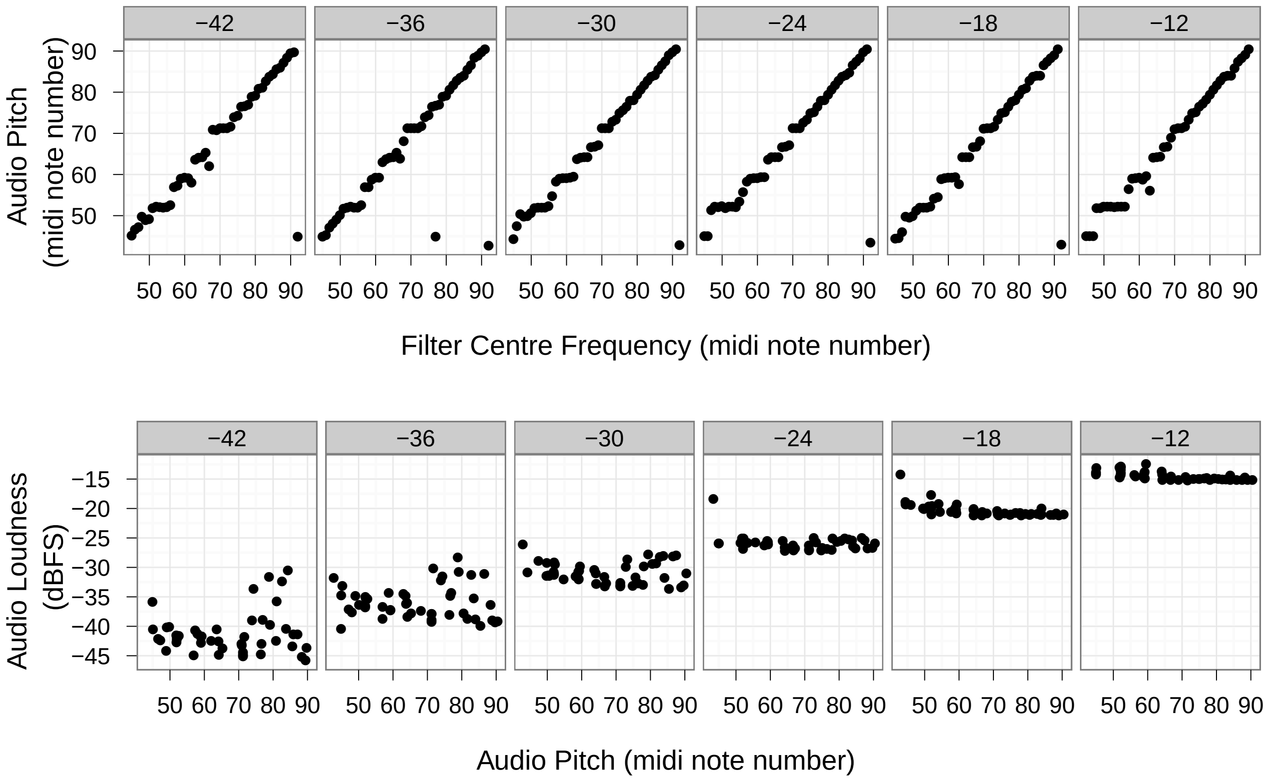

Figure 5: *Corpus audio descriptors compared to control parameters. The graphs are broken into each of the 6 level steps. While there is a trend towards correspondence between target and result, a precise specific filter frequency can be associated with many neighbouring output audio pitches, especially at low pitch ranges. This coarseness means that clearly, the use of a bandpass filter alone is not sufficient to create particular target tones and give a musician the necessary control over the instrument's pitch.*

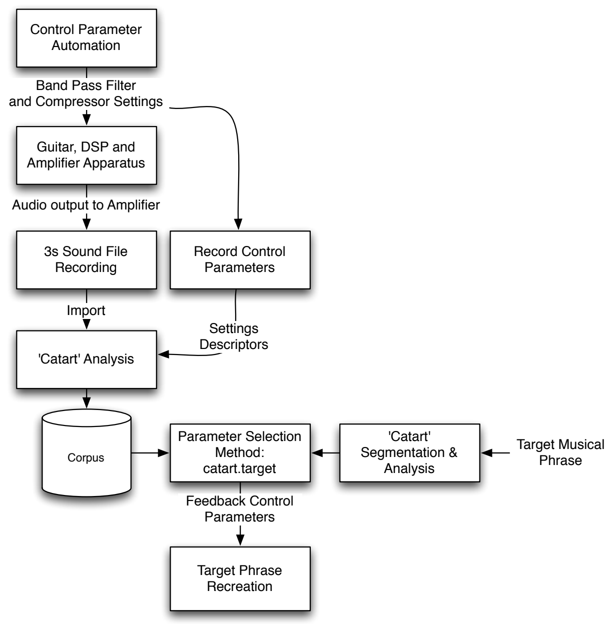

Figure 4: *The creation of a corpus of acoustic parameters and the resulting selection of control parameters.*

The relationship between output audio pitch and the centre frequency of the filter is shown in Figure 5 (top row of plots). Each plot shows the relationship between pitch and filter centre frequency for a given target gain level. It can be seen that in the lower portion of the pitch range, despite the filter frequency being scanned across the pitch space, the output sound's pitches were clustered at particular pitches only. In the upper pitch range they spread more evenly across the pitch space. The selectable tones in the corpus are therefore relatively limited for the lower range.

The relationship between the AGC gain setting and the output level of the sound is shown in Figure 5 (bottom pane). For low target gain levels (plots towards the left), there is considerable variability between the target level and the output level. However, for high target gain levels (plots towards the right) the output gain is much closer to the target gain.

The corpus can also be analysed for other acoustic characteristics. Figure 6 shows a plot of the available combinations of pitch and spectral centroid for the analysed feedback system. This means that the timbral space of feedback sounds can be explored in a performance context.

3.2 Target Musical Phrase Recreation

As a demonstration of the system in use, we recorded a target musical phrase and recreated it using the feedback system. Two bars of the opening to a popular guitar-based song were recorded as audio (not as MIDI), and passed into `catart`'s target phrase matching module, `catart.target`, where it is segmented and analysed. For each of the segments in the target phrase, matching audio descriptor units from the main corpus are found through a unit selection algorithm, and then the units are temporally concatenated by the concatenative synthesis system (see [15] for details).

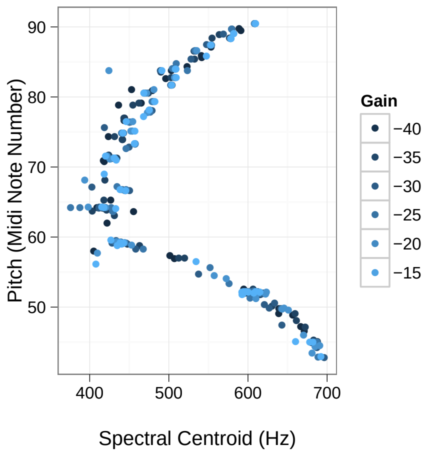

Figure 6: *Corpus analysis also allows inspection of the corpus along audio description dimensions that aren't directly related to control parameters. This graph shows how timbral aspects (the spectral centroid) of the corpus interact with pitch specification.*

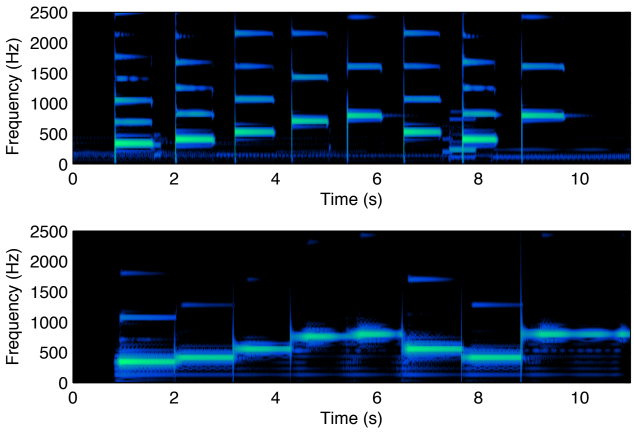

Figure 7: *Two bars of the opening to 'Stairway to Heaven', as performed by a) a guitarist (in the top pane), b) a feedback system with control parameters selected to resynthesize the above sound controlled from the corpus.*

The results of this process are shown in Figure 7. The top pane shows a spectrogram of the original target audio signal recorded directly from the output of the guitar being played in the traditional manner. The lower pane shows a spectrogram of the recorded recreation of this musical phrase, as produced by the feedback system. A range of greater than one octave can be seen in use. The spectrogram shows the strong fundamental frequencies of the feedback, which is consistent with the use of the band-pass filters.

For the recording, some extra damping was required to make sure that strong open resonances (of which there were only a few problematic ones) did not continue to sound between note changes. This was achieved by placing a tissue lightly over the strings and around the fingerboard.

3.3 Benefits of Corpus Usage

To conclude this section, we highlight the advantages of using a combined feedback control and corpus-based method, over the use of a feedback control system alone.

1. Feedback from a guitar doesn't occur at every pitch continuously, only at particular pitches based on the harmonics of the tuned string frequencies, and affected by other real-world acoustical parameters of resonance. Simply specifying the band-pass filter frequency alone cannot give a user information about whether the feedback will occur, or how close to the expected pitch the result will be. As has been shown in Figure 5, the feedback system has quite good coverage at high pitches, but only selective coverage at low ranges. To find out which pitches can be selected from the set, a corpus must be created and analysed. The results of this analysis can be used in a target sound resynthesis system (as in Section 7), or can be used for devising a user interface for performance that shows the performer which pitch options will produce sounds.
2. Not only are there gaps in pitch coverage in a coarse sense, the pitch of feedback usually isn't located at the centre frequency of the filter in a fine sense either. Some notes may 'be there' but sound out of tune. Where a precise correspondence between the expected pitch and the output pitch is required, a corpus-based approach is one way of ensuring the tolerance with which this correspondence may occur.
3. The analysis process allows other acoustic characteristics of the output tones to be incorporated into the selection process and used for mapping. In Figure 6 the tone's spectral centroid is shown in relationship to the pitch. We can see that there appears to be an equivalency, in terms of spectral centroid, between tones of low pitch and high pitch, while many mid-pitch tones have low values of spectral centroid. This parameter is possibly an aspect that can be used as an interaction mapping target. Pitch is the main element considered as a target in this paper, due to its crucial role in most common musical forms, but other acoustic characteristics can be targeted in much the same way that pitch has been.

4. DISCUSSION

In this paper a system for controlling feedback from the combination of an electric guitar, an amplifier and an audio interface acting as a control loop, was described. This system was investigated as an alternative solution for using feedback in musical contexts. The concatenative synthesis system described by Schwarz [16, 13, 15] was used, and the contribution made in this research was to investigate the adaptation of concatenative synthesis methods to the creation of a corpus of audio recordings associated with control parameters for the feedback system. We enumerated the sounds that could be created with this particular apparatus, and have shown as a proof of concept that the pitches in a target musical phrase could be recreated. Corpus based methods also bring with them many other opportunities for specifying control based on other audio parameters (such as timbral descriptors).

While it is difficult to play multiple feedback guitars simultaneously with traditional musical methods, the automation of the aforementioned process means that multiple systems like that described above may be used simultaneously, with their own corpi and within a single control system. This is a possible solution to the question of pitch coverage, as each guitar may be tuned differently, resulting in different harmonics being available. For musical contexts this allows a musical performance with these systems to take the role of an ensemble, with musical roles associated to different guitar and amplifier combinations. Furthermore, spatial sound can be achieved by spatially separating the performing feedback systems, in much the same way that spatial sound is created by a band performing in separate locations.

For live performance contexts, the input from a typical musical instrument could be used as a target sound, with processing and corpus query happening in realtime. This allows the possibility that a musician could control a feedback system in real-time from another, more familiar, musical instrument.

Some limitations exist - one of these is the reliance on the descriptor algorithms to reliably pick the pitch/characteristics of the input sounds. If there are multiple tones, or any type of transition between states, these descriptor algorithms may fail to adequately measure or describe the predominant sound quantitatively. Coupled with the noise in typical guitar systems of this nature, this difficulty can be quite important. The use of high gain levels, however, does ameliorate this problem somewhat. A possible solution is to use a set of different pitch algorithms and a voting mechanism to select a pitch robustly, or alternatively may be based on psychoacoustics.

4.1 Further research directions

The use of a band pass filter may result in much of the characteristics outside the the pass-band to be filtered from the output sound. There is therefore a tradeoff between the narrowness of the filter (with its associated effectiveness at determining the pitch) and the spectral characteristics of the feedback sound being controlled. Research into different filter types (such as comb filters with high Q-values for instance) is a useful approach to be explored, and dynamic filtering may also be a useful approach to employ.

This research has not delved into the questions surrounding multiple pitch output from a single feedback system. Given that two bandpass filters and automatic gain control systems can be included in separate loops within the one instrument, the possibility of two sounds emanating from the same feedback system is not without a theoretical basis (and indeed informal experiments have shown great promise for this technique). However, the control method described in this paper relies on single tones to be analysed by the audio descriptor algorithms for the corpus creation. Two separate control parameter settings taken from a corpus of this type are likely to interact in non-orthogonal ways, but it may be that systems can be created which account for the interactions effectively.

The pitch coverage of the corpus is based on the string's tuning system of the strings, which is therefore quite important for allowing a large range of pitches. The corpus analysed in this study used a tuning system based around fifths and octaves, which therefore reinforced particular notes and resulted in other gaps in the pitch coverage. Investigating a statistical measure of the corpus coverage and target difference, along with a variety of different tunings, may lead to an optimal tuning system being uncovered.

5. CONCLUSION

The use of feedback systems in musical performance is hindered by the requirement that the performer learn a complex mapping between the control parameters of the system and its sound output. We have introduced a corpus-based component to a feedback system, which allows a performer to select sounds based their acoustic properties. This removes a major barrier to the use of feedback systems in musical settings where accurate pitch production is required, and also opens a wide variety of musical possibilities in which the timbral characteristics of feedback systems may be explored.

6. ACKNOWLEDGEMENTS

This research was supported by a Research Development Grant from the Faculty of Arts and Social Sciences, University of Technology, Sydney.

7. REFERENCES

[1] Michael Nyman, *Experimental Music: Cage and Beyond*, Cambridge University Press, Cambridge, UK, 2nd edition, 1999.

[2] Jeffrey M. Morris, "Feedback instruments: Generating musical sounds, gestures, and textures in real time with complex feedback systems," in *Int. Computer Music Conference*, Copenhagen, Denmark, 2007, pp. 469–476.

[3] E. J. Berdahl, Julius O. Smith III, and Adrian Freed, "Active damping of a vibrating string," in *6th Int. Symposium on Active Noise and Vibration Control*, Adelaide, Australia, 2006.

[4] E Berdahl and Julius O. Smith III, "Some physical audio effects," in *9th Int. Conference on Digital Audio Effects*, Montreal, Canada, 2006.

[5] E Berdahl, Julius O. Smith III, and Günter Niemeyer, "Mechanical sound synthesis: And the new application of force-feedback teleoperation of acoustic musical instruments," in *Proceedings of the 13th Int. Conference on Digital Audio Effects*, Graz, Austria, 2010.

[6] E Berdahl, Julius O. Smith III, and Günter Niemeyer, "Feedback control of acoustic musical instruments: Collocated control using physical analogs," *J Acoust Soc Am*, vol. 131, no. 1, pp. 963–973, 2012.

[7] Christopher Burns, "Emergent behavior from idiosyncratic feedback networks," in *Int. Computer Music Conference*, Singapore, 2003.

[8] Christopher Burns, "Designing for emergent behavior: a john cage realization," in *Int. Computer Music Conference*, Miami, USA, 2004, pp. 193–196.

[9] Agostino di Scipio, "'Sound is the interface': from interactive to ecosystemic signal processing,"' *Organised Sound*, vol. 8, no. 3, pp. 269-277, 2003.

[10] Aengus Martin, Sam Ferguson, and Kirsty Beilharz, "Mechanisms for controlling complex sound sources: Applications to guitar feedback control," in *New Interfaces for Musical Expression*, Sydney, Australia, 2010.

[11] T. van Waterschoot and M. Moonen, "Comparative evaluation of howling detection criteria in notch-filter-based howling suppression," in *126th Convention of the Audio Eng. Soc.*, Munich, Germany, 2009, vol. Preprint 7752.

[12] Diemo Schwarz, "Concatenative sound synthesis: The early years," *Journal of New Music Research*, vol. 35, no. 1, pp. 3–22, 2006.

[13] Diemo Schwarz, *Data-driven Concatenative Sound Synthesis*, Ph.D. thesis, 2004.

[14] Diemo Schwarz, Roland Cahen, and Sam Britton, "Principles and applications of interactive corpus-based concatenative synthesis," in *Journées d'Informatique Musicale (JIM'08)*, Albi, 2008.

[15] Diemo Schwarz, Gregory Beller, Bruno Verbrugghe, and Sam Britton, "Real-time corpus-based concatenative synthesis with catart," in *9th Int. Conference on Digital Audio Effects (DAFx-06)*, Montreal, Canada, 2006.

[16] Diemo Schwarz, "A system for data-driven concatenative sound synthesis," in *COST G-6 Conference on Digital Audio Effects (DAFX-00)*, Verona, Italy, 2000.