ASSESSING THE SUITABILITY OF THE MAGNITUDE SLOPE DEVIATION DETECTION CRITERION FOR USE IN AUTOMATIC ACOUSTIC FEEDBACK CONTROL

Mr. Marc Ciufo Green
Audio Lab
Dept. of Electronics
University of York
York, UK
marc.c.green@york.ac.uk

Dr. John Szymanski
Audio Lab
Dept. of Electronics
University of York
York, UK
john.szymanski@york.ac.uk

Dr. Matt Speed *
Allen & Heath,
Penryn, UK

ABSTRACT

Acoustic feedback is a recurrent problem in live sound reinforcement scenarios. Many attempts have been made to produce an automated feedback cancellation system, but none have seen widespread use due to concerns over the accuracy and transparency of feedback howl cancellation. This paper investigates the use of the Magnitude Slope Deviation (MSD) algorithm to intelligently identify feedback howl in live sound scenarios. A new variation on this algorithm is developed, tested, and shown to be much more computationally efficient without compromising detection accuracy. The effect of varying the length of the frequency spectrum history buffer available for analysis is evaluated across various live sound scenarios. The MSD algorithm is shown to be very accurate in detecting howl frequencies amongst the speech and classical music stimuli tested here, but inaccurate in the rock music scenario even when a long history buffer is used. Finally, a new algorithm for setting the depth of howl-cancelling notch filters is proposed and investigated. The algorithm shows promise in keeping frequency attenuation to a minimum required level, but the approach has some problems in terms of time taken to cancel howl.

1. INTRODUCTION

In any system where sound captured by a microphone is amplified and reproduced by a nearby loudspeaker, a portion of the amplified sound emanating from the loudspeaker will be received by the microphone. This sound is subsequently re-amplified and fed back to the loudspeaker [1]. In this way, the sound system forms a closed loop (see figure 1). The most apparent effect of this acoustic feedback loop is the screeching sound that can develop (termed 'feedback howl'), which can cause severe limitations to the system's performance.

Formulated in a form relating specifically to acoustic feedback by van Waterschoot and Moonen [2], Nyquist's criterion states that if for a radial frequency ω:

$$\begin{cases} |G(\omega,t)F(\omega,t)| \geqslant 1 \\ \angle G(\omega,t)F(\omega,t) = n2\pi, \qquad n \in \mathbb{Z} \end{cases} \tag{1}$$

where $G(\omega,t)$ and $F(\omega,t)$ represent the short-term frequency responses of the forward and return parts of the loop respectively, then the system is unstable and has potential to feed back at that frequency.

Howl inevitably occurs in sound reinforcement systems as amplification levels are increased beyond a certain level - the system's Maximum Stable Gain (MSG) [2]. Given the system loop magnitude response, it is possible to predict the MSG as follows:

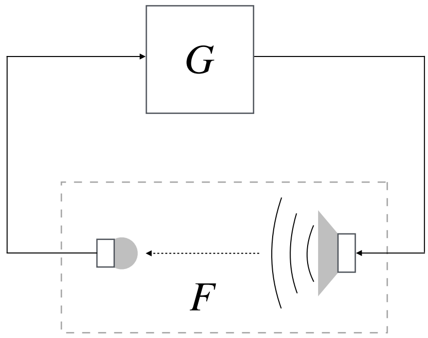

Figure 1: *Schematic for a PA system producing acoustic feedback.* ***G*** *represents the electroacoustic forward signal path, and* ***F*** *represents the acoustic return path.*

$$\text{MSG [dB]} = -10\log_{10}\frac{\max(|G(\omega)F(\omega)|^2)}{\overline{|G(\omega)F(\omega)|^2}},\ \ \omega \in \mathcal{P} \tag{2}$$

where the denominator of the fraction is the Mean Loop Gain (MLG) and $\mathcal{P}$ represents the set of frequencies in the range of interest that fulfil the phase condition of equation (1). The MSG defines the upper limit for usable amplification levels from any given PA system. It is the aim of feedback control systems to increase the MSG of a system, giving more usable gain before feedback howl occurs.

Despite the availability of numerous automatic feedback management systems on the market, the majority of professional sound engineers prefer to manage feedback manually, typically reducing the magnitude of problematic frequency bands using a graphic equaliser [3, 4]. This is largely due to the common perception that automatic feedback control systems are unreliable - there is always the risk of an automatic system falsely identifying a desired sound component as feedback howl and attempting to suppress it, or an actual instance of howl going undetected. Both scenarios could ruin a carefully-constructed live mix.

Setting aside these flaws, a reliable automatic feedback management system would make a desirable addition to live sound technology. The acoustic return path response $F(\omega)$ can change dramatically over time depending on room temperature [5], the addition of an audience into a performance space and particularly

* M Speed is now affiliated with Focusrite Audio Engineering

microphone movement [6, 7]. This can affect the frequencies at which feedback howl is likely to occur.

This paper investigates the Magnitude Slope Deviation (MSD) method for automatic detection of feedback howl, and is organised as follows: Section 2 outlines previously-proposed feedback howl detection methods, details the MSD method and proposes new algorithms based upon this method. Section 3 introduces the software toolkit created for this research and outlines the tests that were undertaken. Results from these tests are presented in Section 4 and discussed in Section 5. Section 6 concludes the paper.

2. HOWL DETECTION

Most modern automatic feedback control systems focus on breaking the gain condition of equation (1) by applying filters to the audio signal in the forward path in order to attenuate problematic frequencies. The Notch filter-based Howling Suppression (NHS) technique is by quite some way the method that has seen the widest use [8, 6, 9]. NHS systems use a bank of narrowband notch filters, reducing the gain at very localised frequency bands to remove howl frequencies.

A key factor in the effectiveness of these systems is the means by which howl is identified from a background of 'desired' musical or speech sound. In order to minimise false positive identifications that could result in incorrect suppression of music or speech, it is important to accurately differentiate howl from desired signal components

The first step in howl identification is spectral analysis of the incoming signal, followed by the application of a standard peak picking algorithm to find local maxima in the spectrum and identify a number of candidate howl frequencies (generally around 10) [2, 10, 9]. Each candidate frequency peak is subsequently analysed to determine whether the peak is caused by a feedback howl or a desired source signal component. Various methods of doing this have been proposed, based upon observed spectral and temporal characteristics of feedback howl that are distinct from music or speech, including:

- higher magnitude than desired components (when allowed to develop) [10]
- sinusoidal in nature - highly localised in frequency [11]
- lack of harmonic components until signal clips [11]
- consistently present across time, very little frequency deviation [2]
- exponential increase in magnitude until signal clips [12]

These features are illustrated in Figure 2, which shows a spectrogram of a simulated microphone signal featuring speech components and a howl component that is clearly visible at just over 700 Hz.

The methods of howl identification broadly fall into two categories based on which howl characteristics they utilise in their analysis: spectral or temporal. Spectral methods compare the magnitude of a candidate howl peak to a reference magnitude that can be set manually, or obtained by a number of means. If the ratio between the candidate howl peak and the reference magnitude exceeds a certain threshold, the candidate peak is flagged by the system as howl. Reference magnitudes can be pre-set absolute values [3, 2, 13, 8, 10], or calculated from the average power across the spectrum [8, 14, 15, 9], neighbouring frequency magnitudes [16, 12] or harmonic frequency magnitudes [4, 8].

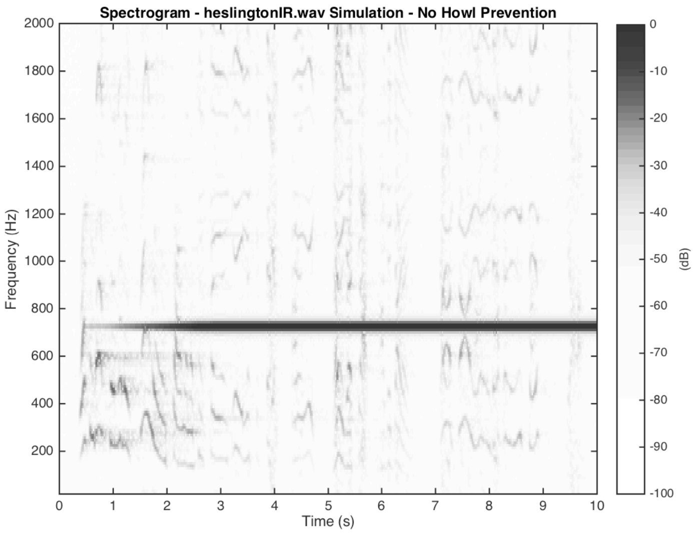

Figure 2: *Spectrogram of feedback simulation using a speech signal and a 'church' acoustic environment. No howl prevention measures have been applied.*

As opposed to the spectral methods, temporal methods of howl identification require several frames of frequency data in order to conduct their analysis. The Peak Magnitude Persistence technique [13, 15, 17] looks for candidate frequencies that are present across large amounts of time relative to the typical duration of speech or music components.

The most recently-proposed temporal method of howl identification is measurement of the so-called Magnitude Slope Deviation (MSD) of candidate frequency bins [10]. This technique, proposed by Osmanovic et al. [12], utilises the fact that howl component power increases linearly over time when plotted on a decibel scale. The system described (intended for use in aviation communication systems) looks at the changes in frequency magnitude of a candidate frequency band over time by storing an 8-frame 'history buffer' of magnitude values. Once peak picking has identified potential howl frequencies, the historic data for those frequencies is analysed and a 'global reference' line gradient between the first and last magnitude values in the memory buffer is calculated. Gradient values for adjacent magnitude values are subsequently calculated and compared against this reference gradient to find deviations in gradient value. If the average deviation between gradients is < 0.05 dB per frame then the frequency band in question is flagged as a probable howl candidate. The MSD technique was verified to work well in aircraft communication scenarios, but has not previously been tested in live music or speech sound reinforcement scenarios.

2.1. The 'Summing' MSD Method

The originally-proposed method to calculate MSD involves numerous gradient calculations for each new frame of frequency spectrum data. For a magnitude history buffer of length N and with K frequency bins, KN gradients must be calculated. For more detailed frequency analyses, the number of calculations required increases significantly, and this computational inefficiency is not ideal for time-critical feedback cancellation scenarios. To mitigate

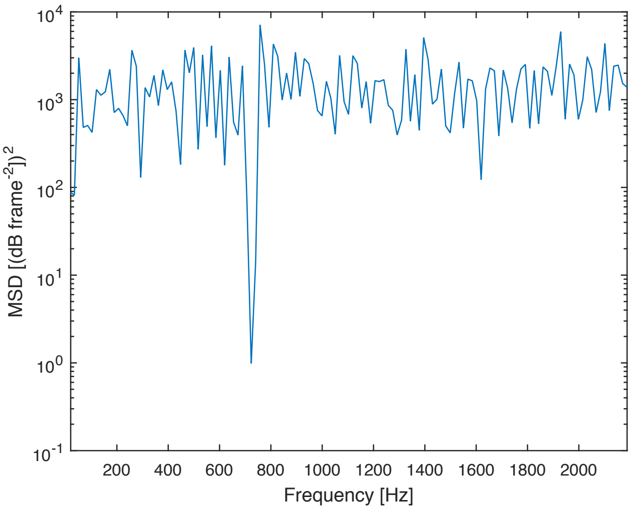

Figure 3: *MSD data calculated with the summing method using a 16-frame FFT history buffer for the howl simulation shown in Figure 2.*

this problem, a new method for assessing the MSD of frequency spectrum data was implemented and tested as part of this research.

In the new 'summing' method, a finite-difference approximation to the second-order derivative (with respect to time) of the decibel-scale magnitude history buffer data is calculated. Any frequency bin that is linearly growing in magnitude over time should have values for gradient change that are consistently close to zero. To calculate the Magnitude Slope Deviation, the absolute values of historical gradient change data for each frequency bin are squared (in order to accentuate the difference between minima and the rest of the data), then summed. The approach can be expressed as:

$$\mathrm{MSD}(k,m) = \sum_{n=(m-N)+1}^{m} |G''(k,n)|^2 \tag{3}$$

where m is the present analysis frame and $G''(k,n)$ is the finite-difference approximation of the second derivative of the dB-scale magnitude history data by frequency bin k and analysis frame n, with respect to n. This produces pronounced minima where magnitude gradients are more consistent over time, and minima that are below a certain threshold are flagged as howl frequencies. Figure 3 shows data calculated using this technique for the howl simulation shown in Figure 2, using a history buffer of 16 frames (N = 16) of FFT analysis data, each calculated using 256 samples (k = 128). The pronounced valley at 725 Hz is a strong indication that howl is present at that frequency.

One problem with this summing approach is that $\mathrm{MSD}(k,m)$ values would approach zero if there were periods of time with no energy in the kth bin. This would give false positive howl identifications. A future version of this algorithm should address this issue with the inclusion of a condition to exclude the identification of frequency bins with zero or very low energy.

2.2. The MSD-Inspired Notch Depth Setting Algorithm

One of the problems faced by Notch filter-based Howling Suppression systems is how best to set the depths of notch filters introduced to cancel howl. It is preferable that the notch filter depths be as shallow as possible so as not to affect the 'desired' sound too much, whilst still being deep enough to ensure that the gain condition of equation (1) is broken. This paper proposes a new method to set notch filter depths known as MSD-Inspired Notch Depth Setting (MINDS) algorithm.

The idea behind MINDS, given that feedback howls are characterised by their growth over time, is that one way to find the minimum filter depth required to cancel the howl is to monitor the candidate frequency's magnitude over time and gradually increase the depth of a notch filter until the magnitude of the howl ceases to increase. At this point, if the howl magnitude remains relatively constant, the filter must be holding the howl in equilibrium (loop gain $\approx$ 0 dB) and only a small additional depth increase should be needed to cancel the howl.

MINDS is implemented by utilising historical frequency magnitude data to compare the latest two magnitude gradients. If both these gradients are above zero (or close, allowing very slowly decaying howls where loop gain is just below unity) then the notch depth is increased by 1 dB, up to a maximum depth of -25 dB. It is reasoned that if both these gradient values are negative, this indicates that any howl present must be in decline, whereas if one gradient is positive and the other negative, this is indicative of a rapidly-changing signal that would not be consistent with the presence of howl. Filter depth is not increased in either case. This behaviour is outlined in Algorithm 1, where k_c is a candidate frequency bin. In this way, even newly-received howl candidate frequencies may not trigger the increase of filter depth. The process repeats every time new frequency magnitude data (and with this, new howl candidate data) becomes available.

At present, this algorithm provides no provision for the removal of notch filters or any reduction in their depth after howling has been suppressed. This was not a problem for the short simulations described here.

if $G'(k_c, m) > -0.5$ **and** $G'(k_c, m-1) > -0.5$ **then**
  if *notch depth* > -25 **then**
    notch depth = notch depth - 1;
  end
end

Algorithm 1: MSD-Inspired Notch Depth Setting

3. TEST METHODOLOGY

3.1. The Feedback Analysis and Cancellation Toolkit

In order to test the effectiveness of the MSD and MINDS algorithms in different scenarios in a repeatable fashion, it was necessary to create a system capable of simulating feedback scenarios using any given 'desired' stimulus sound and acoustic environment. To this end, a system known as the Feedback Analysis and Cancellation Toolkit (FACT) was developed. FACT is split into three subsections, dealing with creating the virtual feedback loop (simulation), detecting any howls that arise (detection) and cancelling them out (notch filtering).

The most important part of the FACT system is the simulation of the acoustic feedback itself. This is initialised using a stimulus sound that represents the 'desired' microphone input (this can be any monaural audio file) and a monaural Room Impulse Response (RIR) representing the loop response of the sound system. All of the simulations conducted for this study use stimulus sounds 20-30 seconds in length, and RIRs gathered using one loudspeaker and one microphone using the swept-sine technique. All simulations were run using a sampling frequency of 44.1 kHz.

Since the development of feedback howl is dependent on the Maximum Stable Gain of the sound system, the MSG of each RIR used was calculated using equation (2). In order to achieve a desired target Mean Loop Gain for the simulation, a factor by which to multiply the RIR values was specified as:

$$2^{(\text{Target Gain [dB]} - \text{MLG [dB]})/6} \qquad (4)$$

In order to facilitate the introduction of filters mid-simulation when modelling the acoustic feedback loop, the stimulus sound is split into frames. Each frame is convolved with the RIR one at a time and the output from the convolution is recorded and simultaneously added back into the input, starting at the beginning point of the next frame. The convolution output is not confined to the next frame, allowing convolution outputs to accumulate over time. This effectively simulates the coupling of the loudspeaker output to the microphone input in addition to the desired sound. FACT makes use of Hann windowing, and each frame is overlapped by 50% with adjacent frames in order to smooth changes in notch filter processing that may occur between frames. Figure 4 shows a diagrammatic representation of this process. Hann windows were used as overlapping these by 50% gives unity gain in the overlapped region. A signal that is split into frames that overlap by 50%, windowed using the Hann function and then recombined can be reconstructed exactly.

To gain frequency spectrum data to use in howl analysis, an FFT process using analysis frames of 256 samples was run concurrently with the howl simulations. Before FFT analysis, the simulation output signal was downsampled by one order of magnitude to 4410 Hz using MATLAB's resample function, which automatically applies an anti-aliasing filter. This was in order to cut down on the amount of frequency data to be analysed by the system (full-spectrum analysis would produce ten times more frequency data). Analysis of RIRs using the criteria in equation (1) showed no feedback howls developing below the upper frequency analysis cutoff of 2205 Hz. This was corroborated in preliminary howl simulations without any notch filtering, so the reduced frequency range was deemed sufficient for these simulations.

3.2. Testing 'Summing' MSD

To test the effectiveness of the new 'summing' MSD algorithm against the original, simulations were run using a short sample of conversational speech as stimulus sound. This stimulus was chosen as the original MSD algorithm has previously been established as working well with speech signals [12]. Simulations were run using four different RIRs in order to assess the effectiveness of the algorithm across several different acoustic environments. These were an RIR of a living room recorded by the author - referred to as 'Small Room' - and three RIRs of larger spaces sourced from openairlib.net [18], which will be referred to as 'Church 1', 'Church 2' and 'Hall'.

The simulations were run using an MLG of 2 dB above the MSG level calculated for each IR. Firstly, each simulation was run with howl detection disabled. This was to allow howl to develop in each case in order to confirm the presence of howl frequencies predicted by analysing the RIRs using equation (1). Figure 2 shows the output of one such simulation. Next, the simulations were re-run twice, once using the 'original' MSD detection method and once using the 'summing' MSD detection method, both using magnitude history buffer lengths of 16 frames. The gradient tolerance for the 'original method' simulations were set to 0.5 dB change per frame and the 'summing method' simulations used a minima threshold of 1 $(\text{dB frame}^{-2})^2$ (see equation (3)). This makes the testing more stringent on the summing method, as calculating the MSD based on the values used here for the original method would equate to using a summing minima threshold of 4 $(\text{dB frame}^{-2})^2$. It should be noted that this also has the potential to make results from the original algorithm less accurate, though this was not observed in the tests presented here.

The effectiveness of each algorithm was assessed by analysing FACT output data. Detection speeds were found by examining timestamp data corresponding to the earliest introduction of filters at correct frequencies. The accuracy of each algorithm was determined by finding the number of unique filter frequency values and calculating the percentage of unique values that correspond to actual howl frequencies. Howl identifications in adjacent frequency bins were treated as identification of a single howl (allowing for spectral leakage). For these tests, in order to evaluate computational efficiency, MATLAB's 'Run and Time' feature was also used to generate a report indicating the processing time of each simulation. The simulations were run on an Apple MacBook Pro, powered by a 2.4 GHz Intel Core 2 Duo CPU.

3.3. The Impact of Varying the History Buffer Length

In the original proposal for the MSD method, the buffer length of magnitude history data used for analysis is 8 frames [12]. Since this system was initially proposed to cancel howl only in speech signals, it is possible that 8 frames might not be adequate data for musical scenarios. A series of tests were therefore run to assess how varying the length of magnitude history data used to calculate gradient deviations by the MSD algorithm affects its accuracy in detecting howl from different stimulus sounds.

Tests were run using the 'Small Room' RIR at a gain 3 dB above the MSG. Loop response analysis of the IR at this gain level showed two frequencies liable to howl. Simulations were then run iteratively, with magnitude history buffer lengths varied between three frames - the minimum required to detect a persisting gradient - and twenty-four frames. The simulations were run using three stimuli in order to assess how using different buffer lengths could yield different levels of detection accuracy in multiple scenarios. The stimuli were the conversational speech sample used previously, an excerpt from 'Jupiter' from Holst's *The Planets* [19] ('Classical Music') and an excerpt from 'The Raven That Refused To Sing' by Steven Wilson [20] ('Rock Music'). Howl detection times and accuracies were calculated as described in section 3.2.

3.4. Testing the MINDS Algorithm

The effectiveness of the MINDS algorithm was assessed using data gathered for the magnitude history buffer length tests described in section 3.3. In these simulations, the same instance of feedback howl was detected at slightly different times and therefore had been allowed to grow to different magnitudes before a filter was introduced. In this way, the final depths reached by the filters introduced to counter the howls could be compared against the magnitude of the howl when they were introduced. Notch filter depth

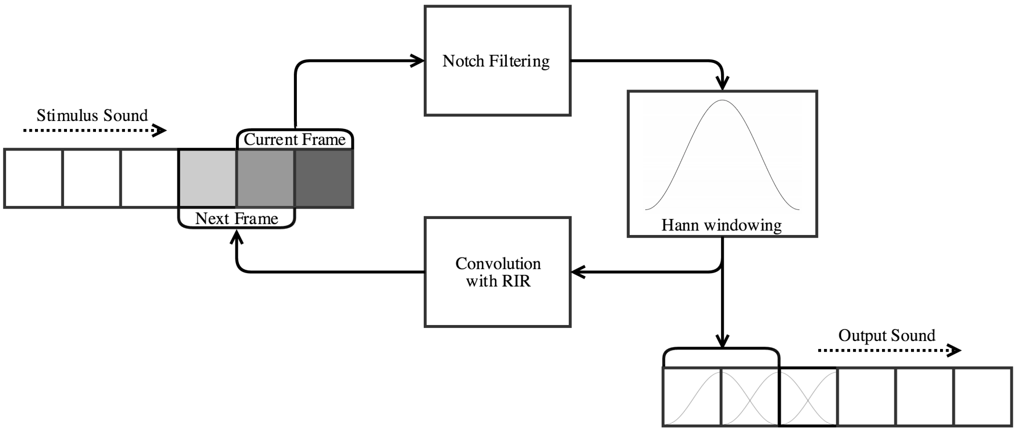

Figure 4: *Convolution process used to create the virtual acoustic feedback loop.*

Table 1: *MSD Algorithm Howl Detection Times [seconds]*

	Howl 1		Howl 2	
IR	Original	Summing	Original	Summing
Small Room	1.074	1.045	5.050	3.846
Hall	2.148	2.017	5.050	5.050
Church 1	7.619	7.619	8.664	8.606
Church 2	1.509	1.451		

Table 2: *MSD Algorithm Processing Times*

Time [s]	Original	Summing
Small Room	48.394	0.262
Hall	46.610	0.339
Church 1	43.841	0.330
Church 2	44.260	0.385

data was only used when the introduced filter remained assigned to the howl frequency by the end of the simulation, not reassigned due to any false positive howl IDs.

4. RESULTS

4.1. MSD Algorithm Types

Both original and summing forms of the algorithm were 100% accurate in howl identification. Table 1 shows the detection times for the first and second howl frequencies in each simulation as detected by both the original and summing algorithms. As can be seen, the summing algorithm detects howl occurrences as quickly or faster than the original algorithm in every case.

Table 2 shows the processing time taken in each simulation to run the MSD algorithm. It can be clearly seen that the summing algorithm is an order of magnitude more efficient to run than the original algorithm. The quoted times represent 1754 calls to the MSD evaluation function, meaning that, on average, the original algorithm takes 26ms to run, whereas the summing algorithm takes 188μs. The summing algorithm is therefore almost 140 times more efficient than the original algorithm on average. It is just as accurate as the original algorithm at detecting howls and always detects howls just as fast or faster than the original algorithm.

Although these tests represent only a small number of scenarios, the results demonstrate the advantages of the summing algorithm very clearly. For this reason, the other simulations in this study used the summing algorithm.

4.2. History Buffer Length

Figure 5 shows how detection accuracy varied with magnitude history buffer length using each of the three stimuli. Accuracy of detecting howl from speech increases rapidly as buffer length is increased, reaching 100% using a buffer of just seven frames. This perhaps illuminates why a buffer of eight frames was originally proposed by Osmanovic *et al.* [12], as the original system was designed for use with speech.

At that same seven-frame mark, accuracy has reached 66% (meaning only one incorrect howl ID) in the 'Classical Music' simulations. The accuracy hits 100% using an eleven-frame buffer, briefly dipping back down to 66% before staying at 100% from a thirteen-frame buffer onwards.

Howl detection using the 'Rock Music' stimulus is significantly less accurate. Accuracy does steadily increase as the buffer length is increased, eventually climbing to 22% as the buffer length reaches twenty-four frames. This provides some indication that a much longer buffer length could remedy the inaccuracy of the MSD algorithm in the rock music scenario. Unfortunately, twenty-four analysis frames represents over a third of a second of audio - already a long time to allow howl to develop. Going by the upward trend, over 100 frames of data would be required in order to approach 100% accuracy (although this has not been tested). This would correspond to almost 1.5 seconds of audio - clearly an unacceptable amount of time to allow potential howls to develop before detection.

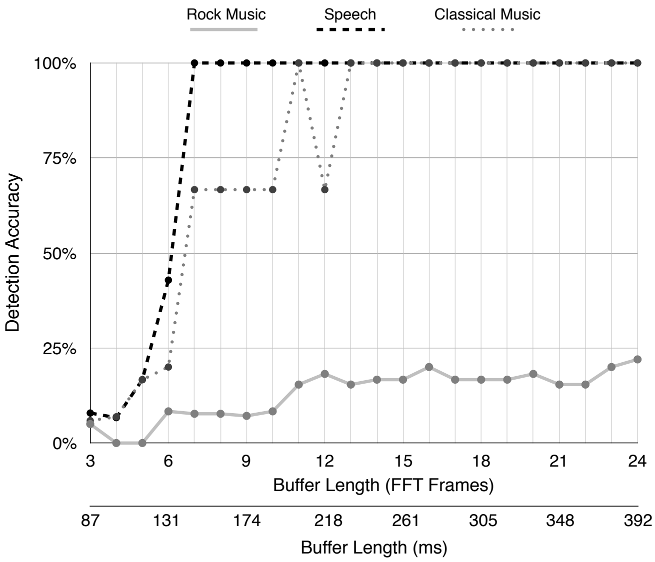

Figure 5: *Detection accuracies of howl from different stimuli when varying the length of magnitude history buffer available for MSD analysis.*

4.3. MINDS Algorithm Evaluation

Figure 6 shows final notch depths plotted against initial howl magnitudes for both first and second howl frequencies. As can be seen, there is a slight trend downwards in filter depths as howl magnitude increases in both cases, but this is not as pronounced as might be expected. In the case of the first howl, for instance, the magnitude of the howls before the filter was added varies from -9.2 dB to over 8 dB - a range of over 17 dB, whilst the final filter depths only vary 3 dB across that entire range. With the second howl, the magnitude varies over a range of over 20 dB whilst the filter depths only vary by 4 dB. These results give some evidence of the effectiveness of the MINDS algorithm at arriving at the minimum notch depth required to cancel a howl peak, regardless of the magnitude of the howl upon detection. As these results are limited to one set of simulations, however, more tests would be required to give conclusive evidence on this point.

Despite the promising nature of the aforementioned results, there is some evidence from other simulations that the algorithm's mandate to keep frequency attenuation to a minimum can sometimes have unintended consequences. Figure 7 shows one such case. The initial howl frequency at 724 Hz is quickly detected and effectively cancelled. The second howl, at 1171 Hz, is detected and a filter is introduced at 2.93 seconds. The depth of the filter reaches its final value of -10 dB by 3.3 seconds. It can be seen from the spectrogram that the howl persists at a significant magnitude for several seconds after the filter depth ceases to increase, remaining a component of the signal for the duration of the ten-second segment depicted in the figure. Since this howl is no longer growing in magnitude and is not significantly higher in magnitude than the desired audio, most detection algorithms are not able to identify it. Since it is not likely to be flagged as a howl candidate again, the howl is allowed to persist, which could affect listener perception of sound quality in a live scenario.

5. DISCUSSION

Results from the testing of the two forms of MSD algorithm reflect favourably on the summing method. The summing method never took longer to catch howl than the original method in any case, and the increase in computational efficiency is vast, which should make it more feasible to implement MSD analysis on embedded processing chips with limited computing power. It could be especially useful in keeping the necessary number of calculations low when applying the MSD algorithm to rock music scenarios, which tests indicated may require much longer buffer lengths (and hence more calculations if the original MSD method was used) than other scenarios to maintain accuracy.

More testing is required to confirm its effectiveness in a wider variety of scenarios, however it seems the summing MSD algorithm can be recommended over the original algorithm despite the limited amount of data comparing the two methods in this study.

The results presented show that in addition to the previously-confirmed effectiveness of the MSD algorithm at identifying howl from speech, the algorithm can also discern howl from classical music, albeit requiring a larger magnitude history buffer to achieve optimal accuracy than for speech in this case. The findings reported here contradict those of van Waterschoot and Moonen [10], who reported results from a similar test indicating that an MSD algorithm using a 16-frame history buffer is 55% likely to trigger in error as gradient deviation threshold is adjusted to give a 100% chance of howl detection. Even more interesting is the fact that their tests used a solo violin piece as stimulus sound. One would expect this kind of stimulus to be relatively spectrally sparse, which the results of this study indicate should make howl detection accurate using the MSD algorithm. One possible reason for this discrepancy is the fact that van Waterschoot and Moonen used a full-spectrum FFT analysis in their test.

The clear exception to the generally excellent accuracy of the MSD algorithm found here is when the rock music stimulus was used. Accuracy ratings in this case were much lower than for other stimuli. The fact that so many false howl identifications were made is probably due to the harmonic richness of the bass guitar sound, clearly visible on the simulation shown in Figure 8. This figure shows the close correspondence of many of the filter frequencies to these harmonics and how a great deal of the low-end of the signal has been attenuated by the five-second mark. Since the stimulus is a polished rock production, the bass guitar likely has dynamic range compression applied. This means that there is a period of time after the instrument's attack phase (when the strings are plucked) where the amplitude of the instrument will be more or less constant, rather than exhibiting a natural amplitude decay. It is this period of constancy that is likely triggering the false howl identifications. This could also explain why the MSD algorithm performed relatively well here when classical music was used as the stimulus, as dynamic range compression is not typically used in classical recordings. Since dynamic range compression is common in rock and pop live performances as well as recordings, this represents a shortcoming of the MSD detection algorithm that needs to be addressed before it can be incorporated into any commercial feedback control systems intended for use in those scenarios.

The problems encountered using MSD with a rock-music stimulus sound recall the early systems using Peak Magnitude Persistence howl detection, which examined the sound signal for correlation at time intervals that had to be "greater than the duration of... a single note in a musical performance" [13] for any level of

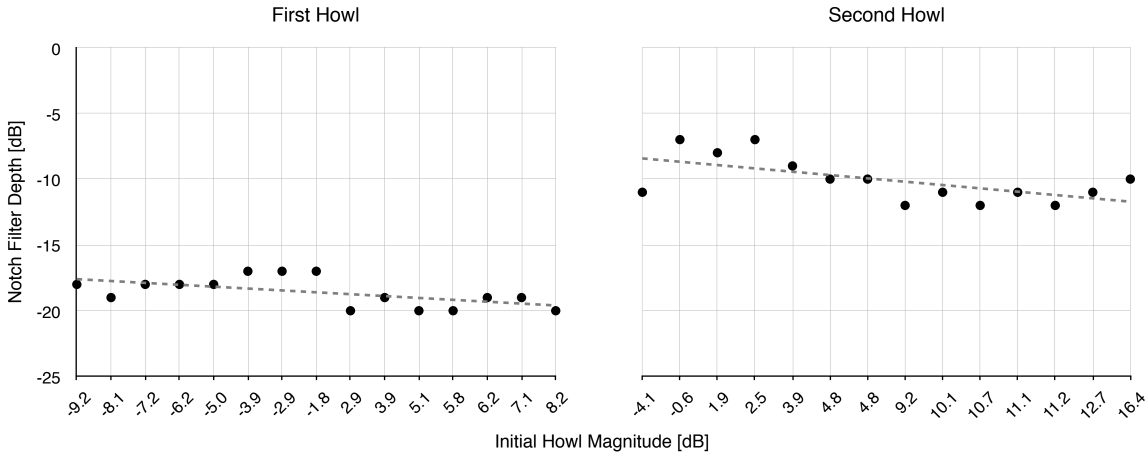

Figure 6: *Final notch filter depths against howl magnitude upon identification. Simulations used 'Classical Music' as stimulus and 'Small Room' as loop response.*

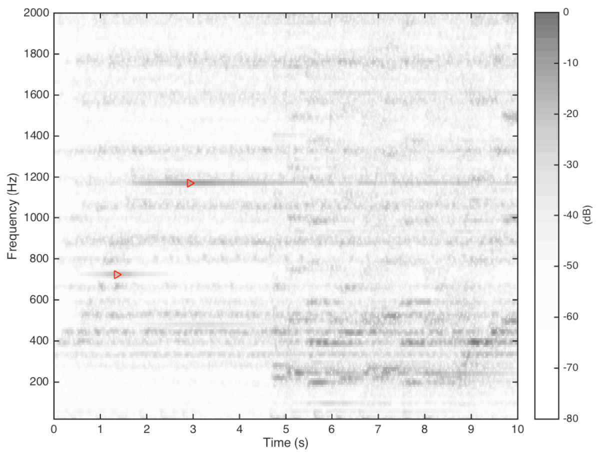

Figure 7: *Spectrogram of simulation output using MINDS to set notch filter depths. Triangles indicate addition of notch filters.*

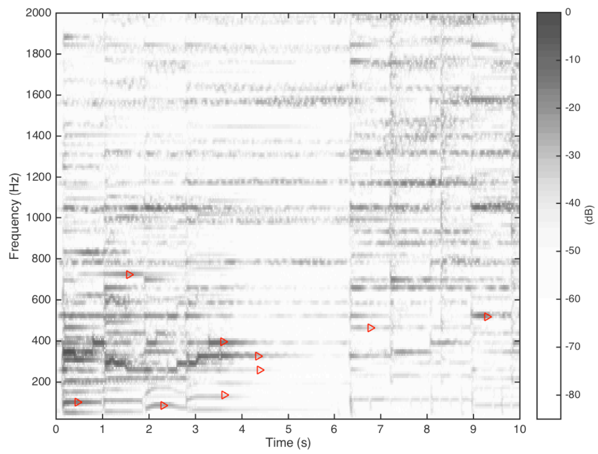

Figure 8: *Spectrogram of simulation output using MSD algorithm on rock music stimulus. Triangles indicate addition of notch filters.*

accuracy to achieved. It seems that the MSD algorithm may face a similar stipulation, at least in situations where dynamic range compression is in use. The typical length of a bass note in 'The Raven That Refused To Sing' is about 1 second. This corresponds closely to the trends shown in Figure 5, which indicate that around 1.5 seconds of audio information would need to be analysed by the MSD algorithm in order for detection accuracy approaching 100% to be achieved with this stimulus.

5.1. MINDS Algorithm

The results presented here show some evidence for the effectiveness of the MINDS algorithm at finding the optimum depth for a notch filter cut, regardless of the magnitude of the howl upon its detection. Since only gradient information is considered, the variance in final filter depths is very small compared with that of the howl magnitudes upon their introduction. Whilst these results are very promising, there is a tendency to make the notch filter depths slightly too shallow. This problem could perhaps be alleviated by introducing a 'final cut' stage to the algorithm. In such a feature, the depth of each notch filter would be increased by an additional fixed amount when the gradient of the howl frequency magnitude turns negative, thus ensuring a fast cancellation of the howl at the expense of perhaps attenuating the problematic frequency slightly more than the absolute minimum required. In its present incarnation the MINDS algorithm simply stops increasing the filter depth at this point, allowing the howl to decay at its own pace, which can sometimes take several seconds.

6. CONCLUSION

The aim of this paper has been to investigate the viability of the MSD and MINDS algorithms for use in automatic acoustic feedback cancellation systems in live-sound scenarios. The new 'summing' method for calculating MSD has been shown to be much more computationally efficient, yet no less accurate or timely, and can be recommended as an MSD implementation of choice moving forward. The algorithm has been shown to work well in the speech and classical music scenarios tested here, but potentially less so in rock music scenarios where extensive use of dynamic range compression can interfere with the functionality of the algorithm, causing it to be 8 less accurate when using comparable history buffer lengths to classical music or speech scenarios.

The MINDS algorithm has been shown to be very promising in terms of its ability to cancel howl instances whilst keeping notch filter depths to a minimum. There are some problems with this approach in terms of the speed at which howl instances can be cancelled, and it would be desirable to modify this algorithm in order to reduce the time taken to calculate optimum notch filter depths.

7. ACKNOWLEDGEMENTS

This research was conducted in conjunction with Allen & Heath Limited, who provided support and supervision throughout the project.

8. REFERENCES

[1] R. O'Donnell, "Prolog to 'fifty years of acoustic feedback control: State of the art and future challenges'," *Proceedings of the IEEE*, vol. 99, no. 2, pp. 285–287, Feb. 2011.

[2] Toon Van Waterschoot and Marc Moonen, "Fifty years of acoustic feedback control: State of the art and future challenges," *Proceedings of the IEEE*, vol. 99, no. 2, pp. 288–327, Feb. 2011.

[3] Joseph O Renauer, "The Feedback Finder: A Simple Solution to a Common Problem," in *Preprints 63rd AES Conv.*, Los Angeles, May 1979.

[4] Michael P Lewis, Timothy J Tucker, and Doran M Oster, "Method and apparatus for adaptive audio resonant frequency filtering," U.S. Patent 5,245,665, Sep. 14, 1993.

[5] K Heinz, "Feedback Elimination Simplified," in *Preprints 54th AES Conv.*, Los Angeles, May 1976.

[6] G Rombouts and Toon van Waterschoot, "Proactive notch filtering for acoustic feedback cancellation," in *Proc. 2nd Annual IEEE Benelux / DSP Valley Signal Process. Symp.*, Antwerp, Mar. 2006, vol. 22, pp. 169–172.

[7] W. K. Connor, "Experimental investigation of sound-system-room feedback," *Journal of the Audio Engineering Society*, vol. 21, no. 1, pp. 27–32, Feb. 1973.

[8] M. H. Er, T. H. Ooi, L. S. Li, and C. J. Liew, "A DSP-based acoustic feedback canceller for public address systems," *Microprocessors and Microsystems*, vol. 18, no. 1, pp. 39–47, Jan. 1994.

[9] Ariel F Rocha and Aníbal J S Ferreira, "An Accurate Method of Detection and Cancellation of Multiple Acoustic Feedbacks," in *Proc. 118th AES Convention*, May 2005.

[10] Toon van Waterschoot and Marc Moonen, "Comparative evaluation of howling detection criteria in notch-filter-based howling suppression," in *Proc. 126th AES Convention*, Munich, May 2009.

[11] C R Boner and C P Boner, "Minimizing Feedback in Sound Systems and Room-Ring Modes with Passive Networks," *Journal of the Acoustical Society of America*, vol. 37, no. 1, pp. 771–775, Jan. 1965.

[12] Nermin Osmanovic, Victor Clarke, and Erich Velandia, "An In-Flight Low Latency Acoustic Feedback Cancellation Algorithm," in *Proc. 123rd AES Convention*, New York, Oct. 2007.

[13] Eugene T. Patronis, "Electronic Detection of Acoustic Feedback and Automatic Sound System Gain Control," *Journal of the Audio Engineering Society*, vol. 26, no. 5, pp. 323 – 326, May 1978.

[14] Mitsuru Hanajima, Michiaki Yoneda, and Toshiyuki Okuma, "Howling Eliminating Apparatus," U.S. Patent 6,125,187, Sep. 26, 2000.

[15] Paul Robert Williams, "System for Elimination of Acoustic Feedback," U.S. Patent 8,634,575, Jan. 21, 2014.

[16] S. Ando, "Howling detection and prevention circuit and a loudspeaker system employing the same," U.S. Patent 6,252,969, Jun. 26, 2001.

[17] J. B. Foley, "Adaptive Periodic Noise Cancellation for the Control of Acoustic Howling," in *Proc. IEEE Colloq. Adaptive Filters*, London, Mar. 1989, pp. 7/1 – 7/4.

[18] "OpenAIR | The Open Acoustic Impulse Response Library." [Online]. Available: http://www.openairlib.net/ [Accessed: Jul. 13, 2015].

[19] G. Holst, Composer, C. Dutoit, Conductor, Holst: The Planets (Classic FM: The Full Works) [Sound recording]., Decca (UMO), 2012.

[20] S. Wilson, Artist, The Raven That Refused To Sing (And Other Stories) [Sound recording]., Kscope, 2013.