7th International Conference
on
Wind Turbine Noise
Rotterdam – 2nd to 5th May 2017

Application of the UK IOA Method for Rating Amplitude Modulation

David Coles, 24 Acoustics Ltd, davecoles@24acoustics.co.uk
Tom Levet, Hayes McKenzie Partnership Ltd, tom@hayesmckenzie.co.uk
Matthew Cand, Hoare Lea, matthewcand@hoarelea.com

Summary

Following the work of the Amplitude Modulation Working Group (AMWG) for the UK Institute of Acoustics (IOA), a method for the quantification of amplitude modulation from wind turbines has been proposed. The method was developed to obtain a consistent and repeatable measure of the modulation depth characteristics of wind farm noise, which can be related to the psycho-acoustic response people experience. Details of the method are discussed and pertinent aspects highlighted.

Results are presented, discussing the analysis of noise measurements undertaken at residential receptor locations near wind farm sites.

1. Introduction

The Amplitude Modulation Working Group (AMWG) was established by the UK Institute of Acoustics (IOA) to derive a method for measuring and rating amplitude modulation (AM) in wind turbine noise.

Amplitude modulation (in this context) is a regular fluctuation in the level of noise, the period of fluctuation being related to the rotational speed of the turbine. AM is considered an inherent characteristic of wind turbine noise. However, a number of factors can give rise to an increase in modulation depth, which can cause specific complaints from residents neighbouring a wind farm. The characteristic of the sound might be described by a listener as a regular 'swish', 'whoomph' or 'thump', depending on the cause and severity of the modulation.

Given the varying severity and perceived annoyance of AM, it is vital to be able to rate AM in a robust and repeatable manner. This in turn allows policy makers to consider penalising levels of AM that are considered unacceptable.

The AMWG has developed a method to reliably identify the presence of amplitude modulated wind turbine noise within a sample of data and rate its magnitude.

The AMWG published a Discussion Document in April 2015 (IOA AMWG, 2015). Following publication, comments, observations and criticisms were received from interested parties. Taking input from the responses, a final 'Reference Method' was developed for adoption.

2. Method Strengths

The method proposed by the AMWG addresses a number of key issues that should be considered when assessing AM. The primary strengths of the method are described below.

2.1 Quantification of Amplitude Modulation

The ability to quantify the magnitude of AM is crucial. The metric obtained should be meaningful and representative. Where it is considered that levels of AM at a site are unacceptable, quantification of AM also allows the effect of any mitigation to be measured and therefore determine whether sufficient mitigation has been applied.

2.2 Repeatability – Minimal User Input

The method proposed by the AMWG is repeatable as it requires very little input from the practitioner prior to running the algorithm. The only required input for the processing algorithm is to define an allowable range for the fundamental frequency of modulation. This is straightforward to determine, as it is directly related to the rotor speed, and can therefore be calculated from the turbine specification.

2.3 Resistant to Extraneous Noise Sources

The method implements a number of techniques to minimise the effect of extraneous noise sources. These range from band-filtering the input data, to assessing the prominence of spectral peaks in the frequency domain. As such, many samples that are corrupted by extraneous noise (and would usually result in large false-positive values) are rejected by the method and not assigned a value for AM. This significant and effective reduction in false-positives comes with minimal introduction of false-negatives.

2.4 Meaningful Results Can be Obtained Quickly

The assessment of noise from wind turbine sites usually involves analysis of large datasets, spanning weeks or months. Since the AMWG method rejects corrupted noise samples (along with those samples containing no sustained modulation), it allows the practitioner to process large datasets and obtain meaningful results quickly. This enables issues to be addressed and resolved more efficiently.

Notwithstanding the strengths outlined above, it is still essential for the practitioner to exercise professional judgement and review any dataset with an appropriate level of scrutiny. Following the processing of the data, user input is required in the form of a verification process, to ensure identified periods are wind farm related and not affected by other modulating sources. It is possible that other sources in the local environment may be modulating at frequencies similar to the blade passing frequency, and in the same acoustic range, e.g. a dog barking, or a pigeon cooing.

3. Method Overview

The proposed method is a 'hybrid' approach. The modulation depth is calculated in the time domain, while the frequency domain is used to discriminate wind turbine AM and reject samples corrupted by extraneous noise sources (or those containing no obvious modulation).

An overview of the method is presented here, and some key aspects are highlighted. Full details of the method are described in the report published by the IOA AMWG (2016), which should be read by anyone considering implementing this method.

The principal output from the method is a series of 10-minute values representing modulation depth. The 10-minute values are calculated from a sequence of 10-second results. Analysis of each 10-second block comprises the following:

- Band-filtering the input data to focus the analysis on frequencies associated with wind turbine AM;
- Using Fourier analysis to assess the power spectrum and remove energy not associated with the fundamental modulation frequency (which itself should be related to the wind turbine(s));
- Performing an inverse Fourier transform to provide a 'clean' time-series containing energy only at the fundamental modulation frequency (and associated harmonics);
- Calculating the modulation depth by subtracting the L_{95} from the L_5 of the reconstructed time-series.

A key strength of this method is its ability to reject samples corrupted by extraneous noise sources. The techniques used to achieve this, along with other pertinent details, are described below in Section 4.

4. Pertinent Details of the Method

Section 3 provides an overview of the method proposed by the AMWG and highlights the simple principles upon which the method is based. However, the sophistication of the method is contained within the details. Key aspects of the method are highlighted below, however, as mentioned above, the AMWG report should be referred to for full details of the procedure. The reader will note that 10-second samples can be rejected at various stages of the analysis, as described below. The effect of these rejections is realised when calculating the 10-minute values, and forms a fundamental role in the method's ability to discriminate genuine AM. This is detailed further in Section 4.6.

4.1 Band-Filtering Input Data

The input signal for the method is a time-series of band-limited, A-weighted, 1/3 octave L_{eq} data in 100 millisecond samples. The following three frequency ranges (which each encompass seven 1/3 octave bands) are defined:

- 50 to 200 Hz
- 100 to 400 Hz (reference)
- 200 to 800 Hz

The seven 1/3 octave bands should be A-weighted and then summed logarithmically into a single band-passed stream of data for input to the method.

Focussing on a limited frequency range dominated by modulation, assists in both the identification of AM and in excluding spurious data. It also results in higher levels of AM compared to those obtained from broadband (A-weighted) analysis. In fact, the band-limited data can detect AM which might have been masked using a broadband analysis based on overall L_{Aeq} values.

4.2 Fourier Transform

A standard Fourier transform is applied to the input time-series to transform the data into the frequency domain and obtain a modulation spectrum. An important distinction from some frequency-domain based methods, such as that proposed by RUK (2013), is that both the real and imaginary parts of the Fourier output are retained. The full complex output contains phase information and is used later in the analysis to transform the data back into the time-domain.

The input to the Fourier transform is a 10 second block of 100 ms L_{eq} samples. This results in a frequency resolution of 0.1 Hz, and a maximum resolved frequency of 5 Hz. This places a limit on the maximum modulation frequency that can be assessed using this method – since three harmonics are considered, a maximum fundamental frequency of 1.6 Hz can be assessed. This translates to a rotor speed of 32 RPM for a 3-bladed turbine.

The output of the Fourier transform is converted to a power spectrum using equation 1:

$$S = \frac{|F\{x\}|^2}{n^2} \qquad [1]$$

where $F\{x\}$ is the output from the Fourier transform, and n is the length of input data (100, in this case). Analysis of the power spectrum is performed to determine whether the sample contains valid AM. Pertinent details of this analysis are described below in Section 4.3.

4.3 Analysis of the Power Spectrum

A typical power spectrum for a sample containing AM is shown below in Figure 1.

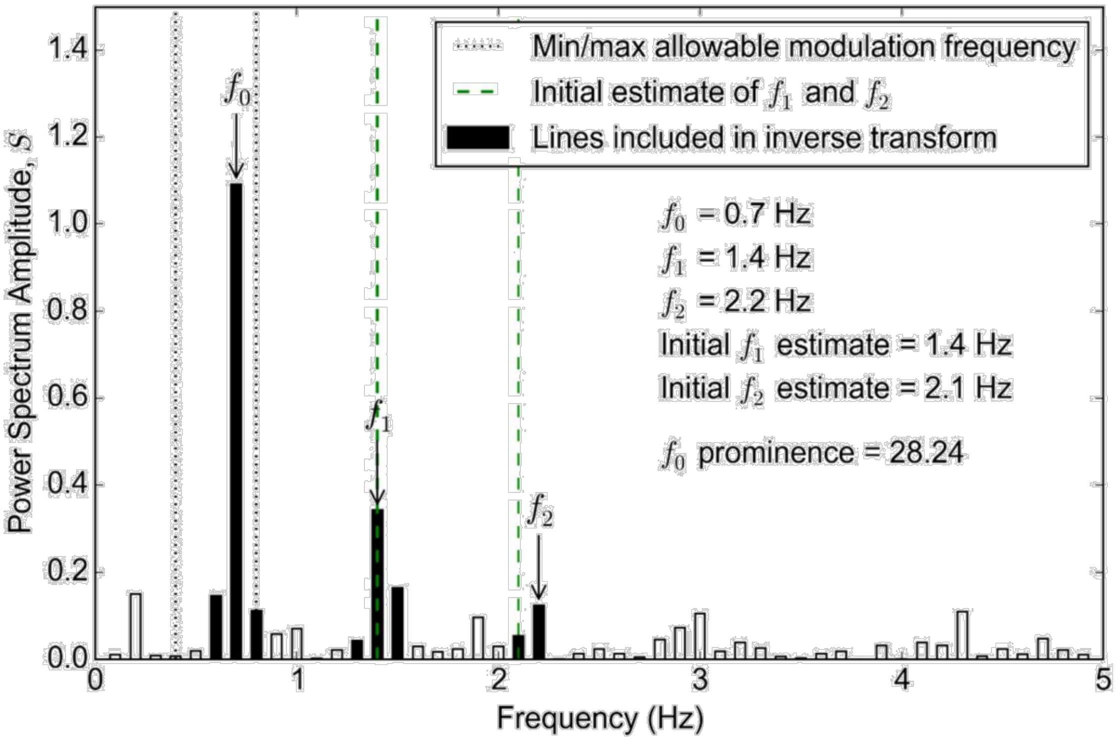

Figure 1 – Power spectrum of a sample containing AM. The positions of the first three harmonics (f_0, f_1, and f_2) are shown, along with the prominence of the fundamental peak.

The first stage of analysis is to find the highest peak within the allowable range for the fundamental modulation frequency (as set by the practitioner). A peak is simply defined here as a local maximum. If a peak is not found within the allowable range, this is a clear indication that the sample has been corrupted (or doesn't contain any notable modulation) and the 10-second sample is rejected from the analysis.

Once a fundamental frequency of modulation has been found, the location of associated harmonics is determined close to the multiples of the fundamental frequency. The method for doing this is described in the AMWG report.

The identification of a peak in the allowable range is not necessarily an indication that the sample contains genuine wind turbine AM. It is possible to greatly reduce the number of false positives by assessing the prominence of the peaks in the power spectrum. This exploits the fact that genuine wind turbine AM produces pronounced peaks in the power spectrum. Figure 1 shows a sample containing high modulation, which produces a very clear peak at the fundamental frequency of modulation (0.7 Hz). Figure 2 shows the power spectrum of a sample containing no notable modulation. There are clearly local maxima within the allowable range of fundamental modulation frequencies. However, none of the identified peaks identified are ‘prominent’ relative to the neighbouring spectral frequencies and this sample should not be considered further in the analysis.

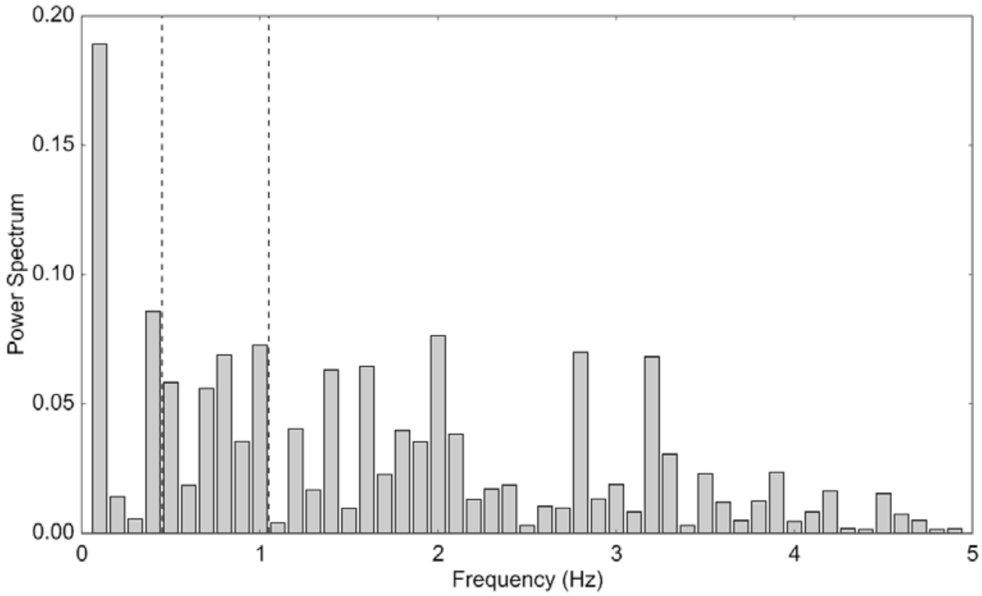

Figure 2 – Power spectrum of a sample containing no AM. Although local maxima are found within the allowable range of fundamental modulation frequencies (marked by the dashed vertical lines), none of the peaks are considered prominent.

The AMWG has proposed a means of determining the prominence of a peak within a power spectrum. This forms a critical part of the analysis and greatly reduces the number of false positives. The method is described below:

1. The magnitude of the fundamental peak, L_{pk}, is taken as the amplitude of a single line in the power spectrum at the frequency of the peak;
2. The two lines either side of the peak are ignored;
3. The masking level, L_m, is taken as the linear average of two lines each side of the peak (beyond those lines immediately adjacent to the peak);
4. The prominence, p, of the peak is calculated using:

$$p = \frac{L_{pk}}{L_m} \qquad [2]$$

An example clarifying the classification of masking lines in the power spectrum is shown below in Figure 3. The lines adjacent to the peak are ignored. The masking lines are the two lines beyond the adjacent lines either side of the peak.

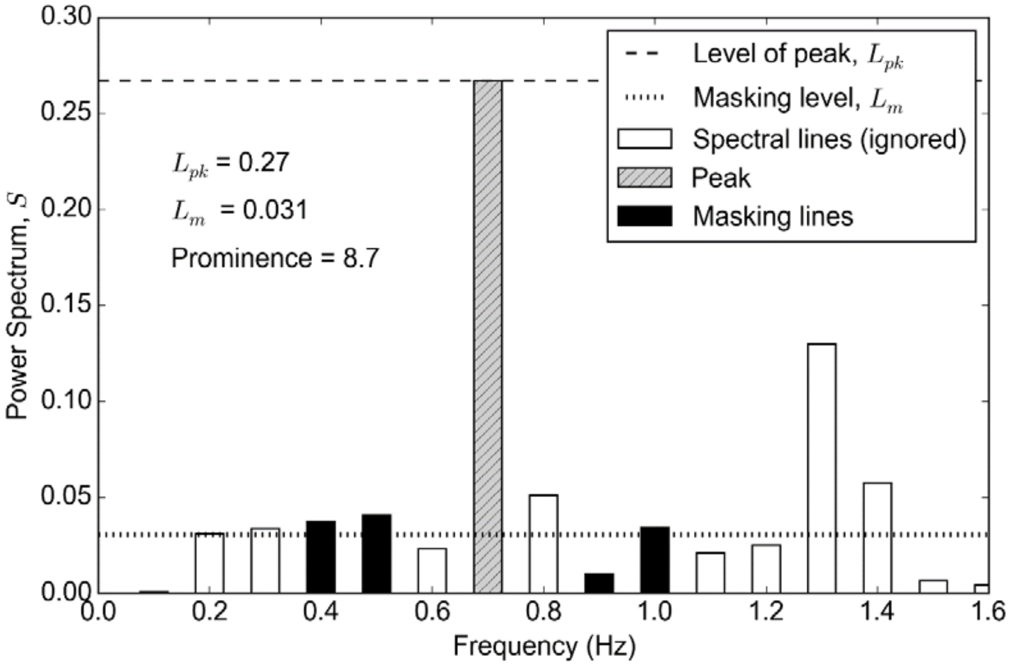

Figure 3 – An example calculation of the peak prominence.

If the prominence of the peak is less than a value of four, the 10-second sample is rejected from the analysis.

4.4 Inverse Fourier Transform

Analysis of the power spectrum is used to identify the frequencies of interest (namely the fundamental and the next two harmonics). However, after these frequencies have been identified, the rest of the analysis is performed on the original output of the Fourier transform (containing real and imaginary components) rather than the power spectrum. For each harmonic identified, three lines in the Fourier output are retained (the centre line, and one line either side). Lines for the corresponding negative frequencies are also retained. All other values in the Fourier output are set to zero. The inverse Fourier transform is then performed on this array (which should only contain energy associated with the fundamental frequency of modulation and its main harmonics). This is clarified in Figure 4 below.

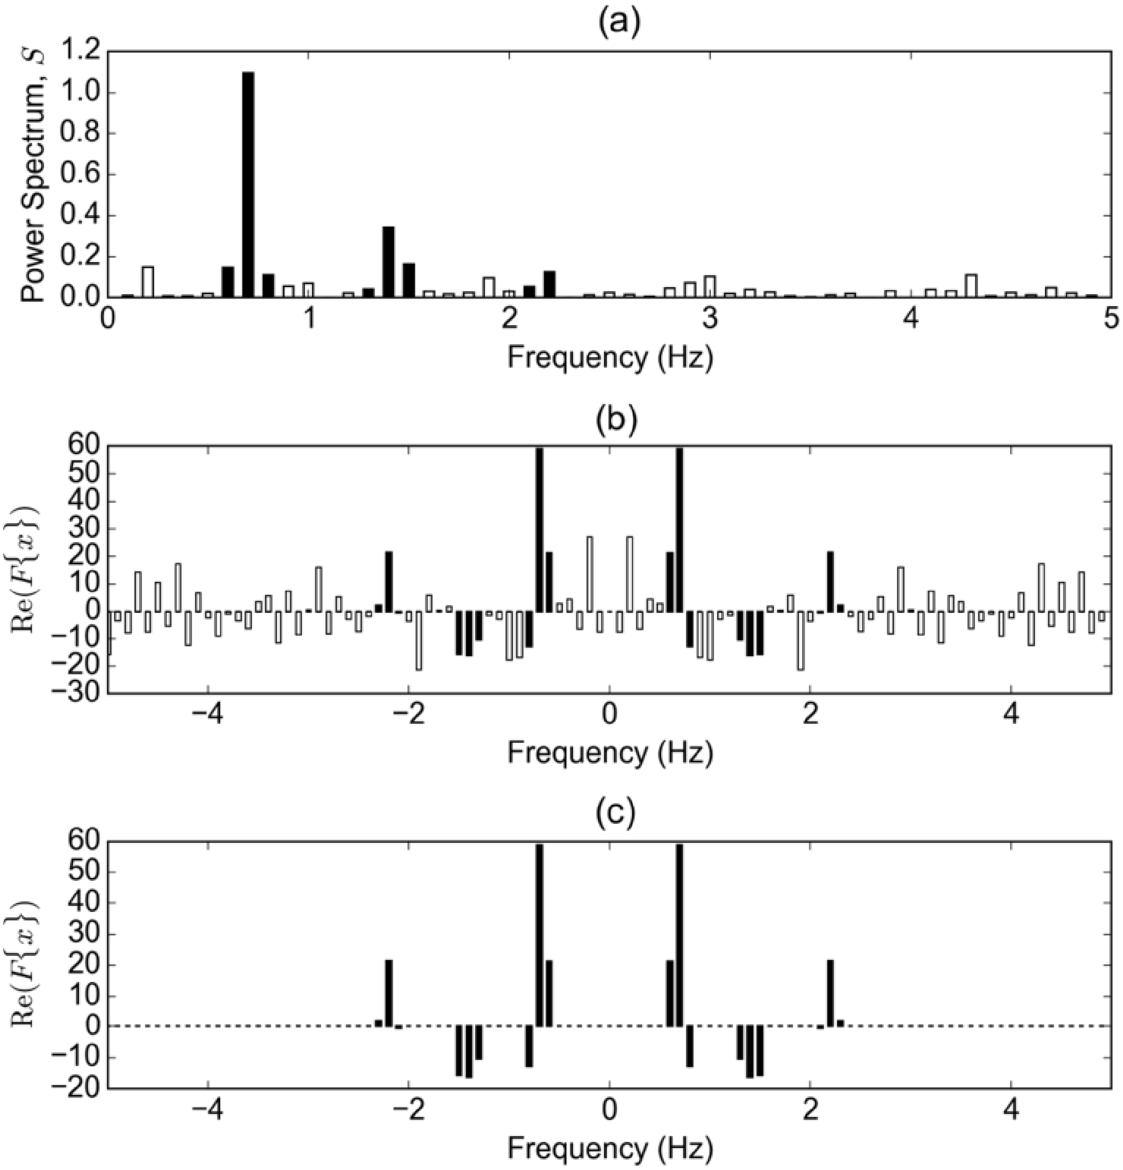

Figure 4 – A clarification on indices to include in the inverse Fourier transform. Panel (a) shows the power spectrum and the identification of indices to include. Panel (b) shows the original output from the Fourier transform (only the real part is shown here) with the identified indices shown as black lines. Note that the complex conjugates are also shown as black lines (the negative frequency components). Panel (c) shows the array with the identified indices included, and zeros at all other values. The inverse Fourier transform is performed on this array (note that the full array, including imaginary components, should be used – only the real part is shown here).

The result of the inverse Fourier transform should be a 'clean' version of the original time-series, containing only energy related to the fundamental frequency of modulation (and its main harmonics). An example is shown in Figure 5.

4.5 Determination of Modulation Depth

Once the reconstructed time-series has been generated, the modulation depth for the 10-second period is calculated simply by subtracting the L_{95} from the L_5, in a similar manner to Fukushima, Yamamoto et al. (2013). Calculating the modulation depth in this manner has the effect of weighting the value towards the highest modulation within the 10-second period.

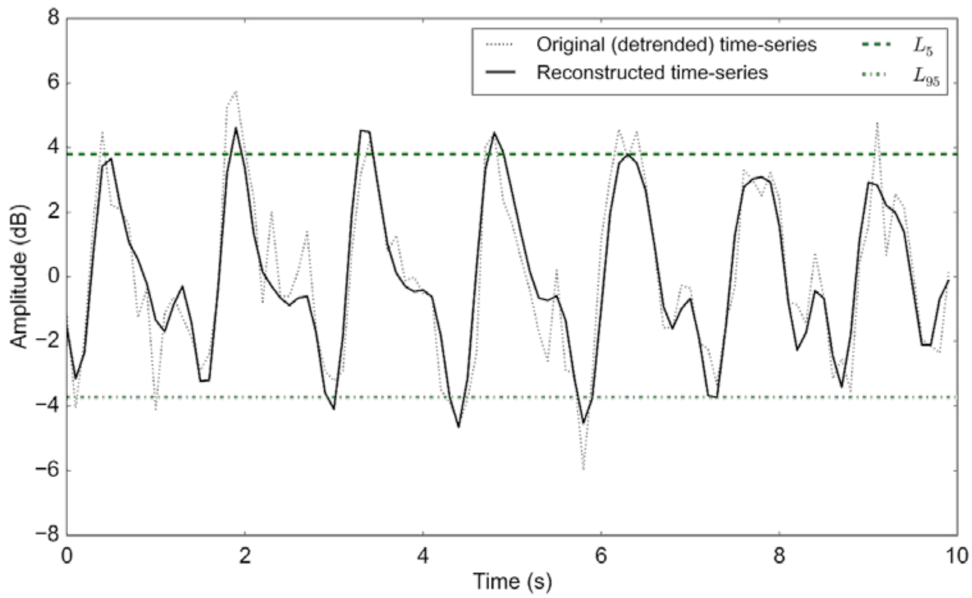

Figure 5 – The reconstructed time-series compared to the original (detrended) time-series. The modulation depth is calculated from the difference between the L_5 and L_{95} (both shown on the chart).

4.6 Calculation of 10-Minute Value

There are a number of possible ways to calculate a value for a 10-minute period from a sequence of (up to 60) 10-second results. One method would be to take the linear average, however, since the modulation within each 10-second period is averaged, averaging these results may undervalue the impact of AM within a 10-minute period. Another option is to take the maximum 10-second result within a 10-minute period. However, this would be very prone to spurious results and result in a value that is not robust from one 10-minute period to another. The AMWG method uses the 90th percentile (L_{10}) of the valid 10-second results. This is considered to represent the typical worst-case instances of AM within a 10-minute interval, without being excessively sensitive to possibly spurious extreme values. Figure 6 shows 10-second and 10-minute results for a 100 minute period, and gives an indication of where the 10-minute values sit within the spread of 10-second results.

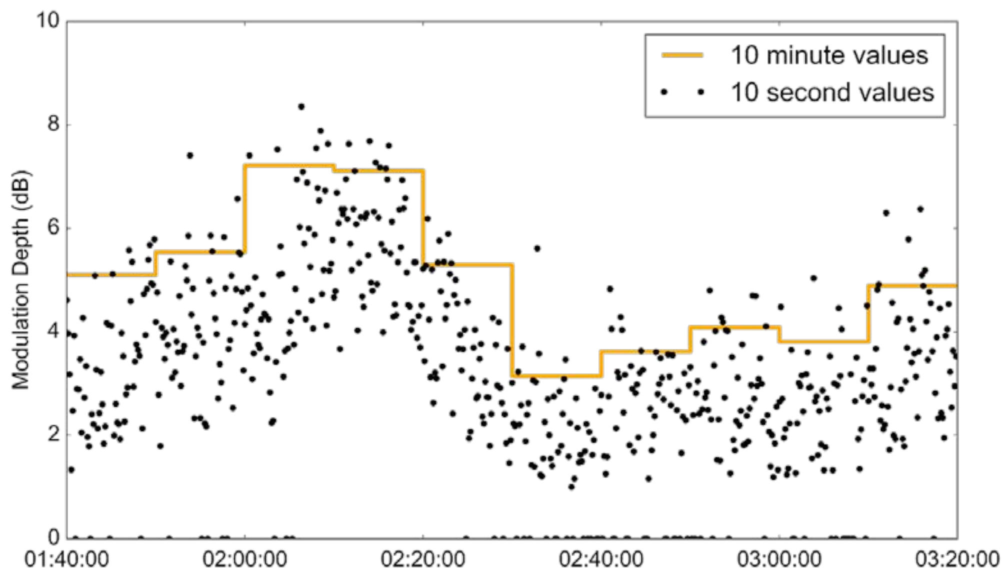

Figure 6 – A series of 10-second results and the corresponding 10-minute values.

It is important to note that only valid 10-second samples are used in the determination of the 10-minute value (some will have been discarded as detailed above). Furthermore, and critically, a value for AM is only assigned to a 10-minute period if there are at least 30 (i.e. 50%) valid 10-second results within that period. This criterion has been found to be a very effective indicator to exclude spurious data where little continuous AM attributable to wind turbines could be detected. In other words, this is an objective indicator of the presence of sustained wind turbine AM with varying magnitude. This criterion was chosen to be conservative, to minimise the risk of false exclusion of valid data, and so it is possible that some samples, i.e. 10-minute periods with more than 50% valid 10-second blocks still represent erroneous data (i.e. false positives). Conversely the 50% criterion will exclude isolated periods of sporadic/brief AM.

The effectiveness of the method to identify and quantify wind turbine AM (even in the presence of extraneous noise sources) is demonstrated in Section 5.

5. Application to Real-World Data

A method for rating wind turbine AM can only be considered fit for purpose if it produces meaningful results when applied to real-world data. The examples presented below demonstrate the effectiveness of the AMWG method in quantifying wind turbine AM and enabling a meaningful assessment to be undertaken efficiently.

5.1 Detecting Amplitude Modulation in the Presence of Noise

The 50% criterion, described above, is a very effective means of suppressing extraneous noise in a dataset. This is demonstrated in Figure 7, which shows a 24-hour dataset corrupted by sources of extraneous noise. Panels (b) and (c) show the difference made by applying the 50% criterion – in Panel (c), the extraneous noise is suppressed (no 10-minute AM values are reported) and ratings are assigned to 10-minute samples within the only period of the day in which the turbines were operational. This illustrates the effectiveness of the method.

5.2 Determining Prevalence of Amplitude Modulation

Determining the prevalence of AM is simplified by the AMWG method since results are not reported for periods which do not contain sustained modulation. It is possible to review data from longer noise surveys quickly and ascertain whether AM has occurred. Figures 8 and 9 demonstrate this – Figure 8 shows a one week period with a relatively high occurrence of AM, while Figure 9 shows a one week period with a relatively low occurrence of AM.

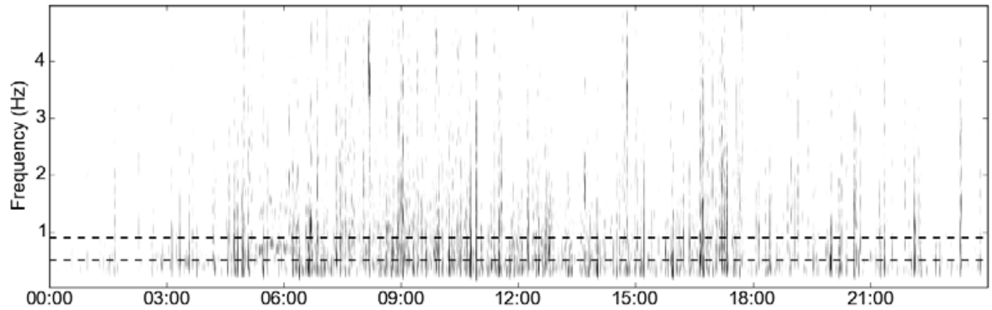

(a)

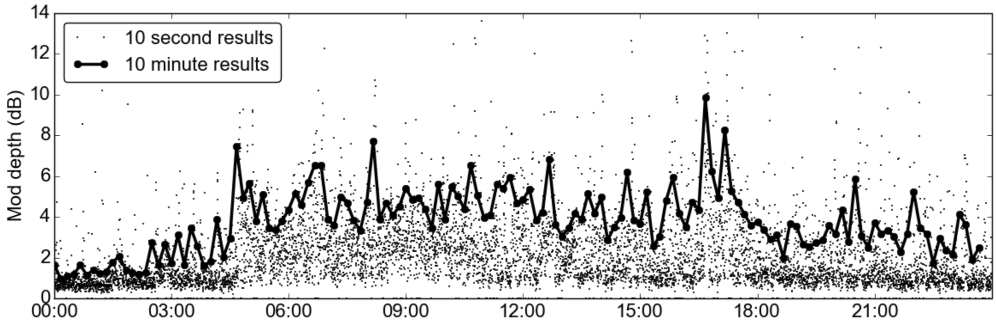

(b)

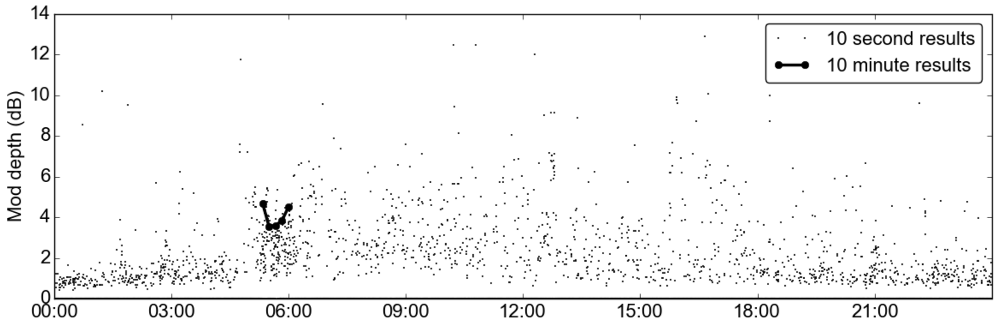

(c)

Figure 7 – Example of 50% criterion applied to data with a relatively large amount of corruption from non-turbine sources (birds, trees, etc.). Panel (a) shows a waterfall plot, which shows that there is only a consistent trend of modulation apparent in the expected modulation frequency range (shown by dashed lines) around 06:00. The 10-minute results are shown both without (b) and with (c) the 50% criterion applied. It was verified in this case that the only valid period in which 10-minute results are presented in (c) corresponds to the only period in which the turbines operated on that day.

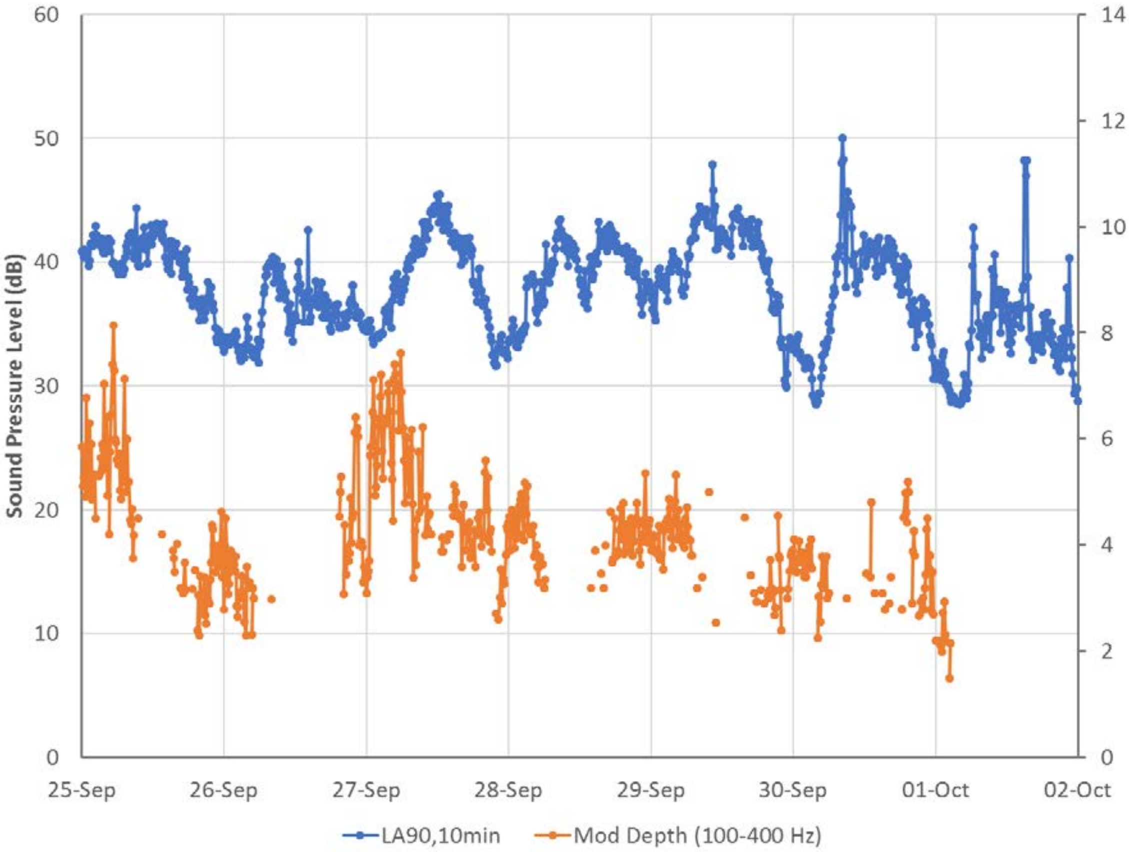

Figure 8 – A one week period with a relatively high occurrence of wind turbine AM. The chart shows data from the same survey location as Figure 9.

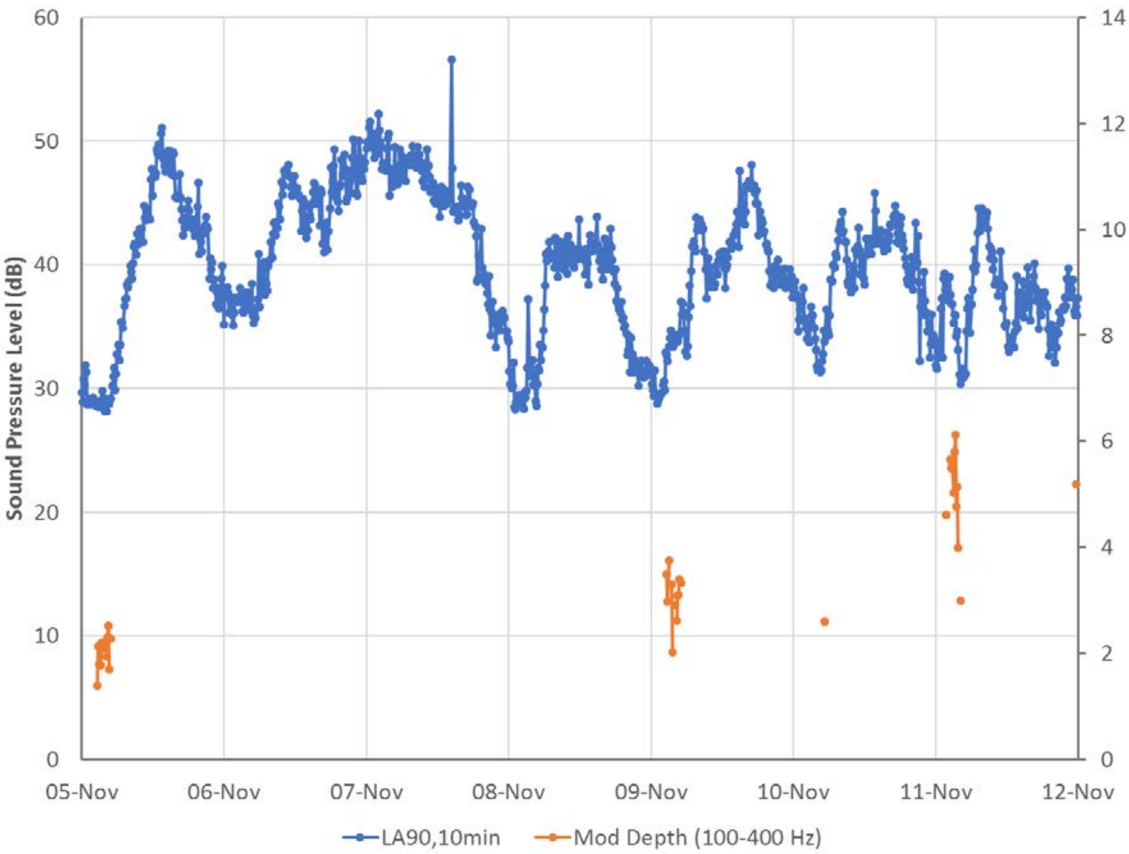

Figure 9 – A one week period with a relatively low occurrence of wind turbine AM. The chart shows data from the same survey location as Figure 8.

5.3 Determining Conditions For Mitigation

The full potential of the AMWG method is realised when noise data are combined with meteorological data to identify conditions under which AM occurs at any given site. This could form the basis of a mitigation scheme to reduce AM levels at relevant times. Figure 10 shows data analysed from a site where AM is found to occur. Panels (a) and (b) show the drastic effect of applying the 50% data filtering criterion – a large number of false positives are removed and two distinct regions are identified showing the conditions under which significant AM is occurring. This provides a substantial increase in efficiency when analysing large datasets, and enables conditions under which AM occurs to be identified quickly.

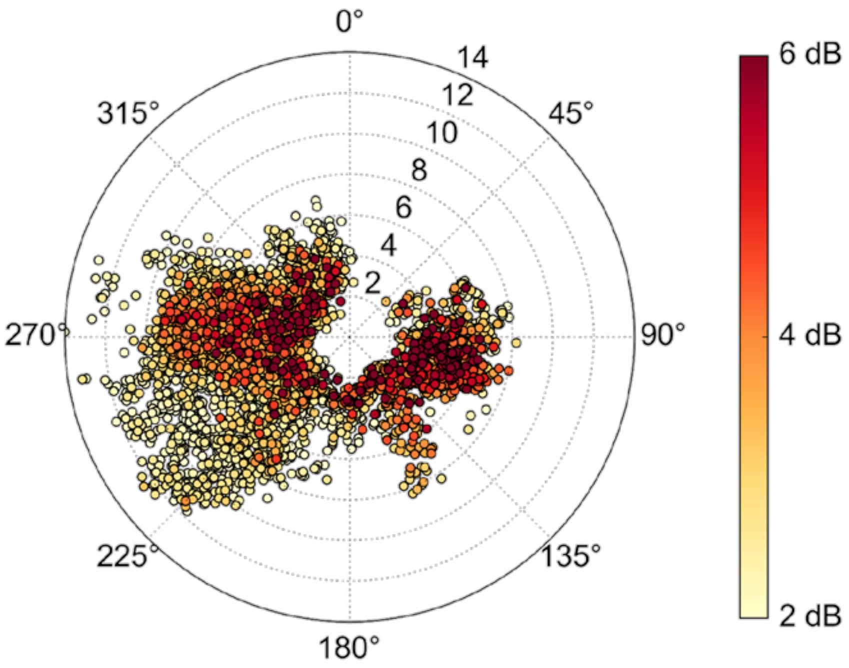

(a)

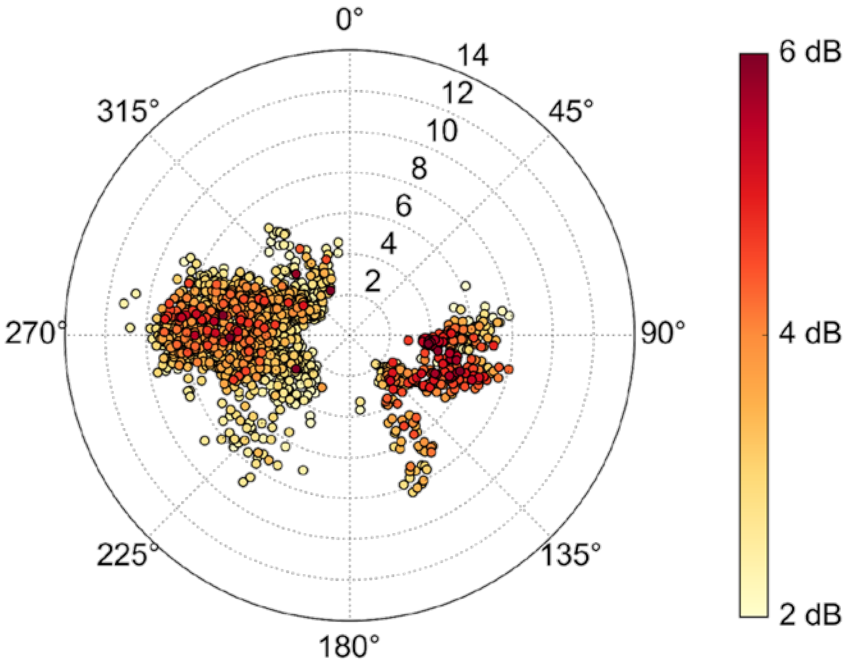

(b)

Figure 10 – A demonstration of the 50% criterion when applied to a dataset comprising noise and meteorological data. Panels (a) and (b) show the data with and without the 50% criterion respectively.

6. Conclusions

The IOA AMWG has published a method for rating wind turbine AM. It provides a meaningful and representative value of the modulation in measured signals, but requires 1/3 octave band measurements at 100 ms resolution. The method utilises a 'hybrid' approach, with the modulation depth being calculated in the time-domain, while filtering of extraneous noise sources is conducted in the frequency-domain. Numerous techniques are employed to minimise false positives and remove samples that are either corrupted or do not contain sustained modulation. The result is a robust and repeatable method for rating AM, which performs well when applied to real-world data and has the potential to significantly increase the efficiency of analysing large datasets.

References

Fukushima, A., Yamamoto et al. (2013) *Study on the amplitude modulation of wind turbine noise: Part 1 – Physical investigation.* Proc. Internoise 2013

IOA Amplitude Modulation Working Group (2015) *Discussion Document, Methods for Rating Amplitude Modulation in Wind Turbine Noise.*

IOA Amplitude Modulation Working Group (2016) *A Method for Rating Amplitude Modulation in Wind Turbine Noise.*

RenewableUK (2013) *Wind Turbine Amplitude Modulation: Research to Improve Understanding as to its Cause and Effect.*