Fourth International Meeting on Wind Turbine Noise

Rome Italy 12-14 April 2011

Measurement of amplitude modulation frequency spectrum

Dave McLaughlin
SgurrEnergy, 225 Bath Street, GLASGOW G2 4GZ, UK
david.mclaughlin@sgurrenergy.com

Abstract

Allegations of amplitude modulation (AM) of broadband noise from wind turbines are increasingly being used as an argument against planning consent for wind farms. While AM is commonplace in the near field of a turbine, it is rarely observed in the far field.

The measurement of the modulation spectrum of broadband noise is well known in the sonar community, and appeared in a text-book as early as 1991. A simplified description is:

1. filter the sampled data to the frequency range of the broadband noise being modulated;
2. square the filtered data;
3. conduct a spectral analysis of the output by Fast Fourier Transform (FFT).

The outcome of this process is the modulation spectrum: a graph of modulation amplitude against frequency.

This procedure has many advantages. Aperiodic modulation, such as caused by traffic noise, will be automatically excluded. Periodic modulation will be detected as a peak or peaks in the modulation spectrum. Modulated noise from a wind turbine will be expected to occur at the blade passing frequency and its harmonics. If shaft rotational frequency is present, that suggests the source may be due to one blade more than the others.

If the amplitude over time does not follow a smooth sinusoid, but varies in some other periodic fashion, that will be reflected in a modulation spectrum with multiple peaks at harmonics of the fundamental frequency. Types of AM such as “swish” and “thump” may be distinguished by different harmonics in the modulation spectrum, or by differences in the frequency spectrum of the underlying noise being modulated.

A better understanding of the different characteristics of different kinds of AM will aid wind farm operators in managing their assets and wind turbine manufacturers in refining their product.

Introduction

Allegations of amplitude modulation (AM) of broadband noise from wind turbines are increasingly being used as an argument against onshore wind farms. While AM is commonplace in the near field of a turbine, it is only rarely observed in the far field.

When AM is observed in the far field, it is important to have a means of quantifying it. This can help:

- operators to manage their turbines to avoid complaints,
- and manufacturers to design their turbines and control systems to prevent AM occurring.

Some methods have been proposed to measure the modulation spectrum of broadband noise, for instance Lee *et al* [1]. But an algorithm already exists and has been in sonar signal processing text-books since 1991 [2]. It can be implemented largely in the time domain using Finite Impulse Response (FIR) filters.

Other work is in progress [3, 4] to quantify modulation. While a simple metric may prove effective in predicting the perceived amplitude of the modulation, there are a range of characteristics that can assist in describing the noise, and which may help in diagnosing the cause and, ultimately, in curing it.

Influences on far-field amplitude modulation

Some work has been carried out to advance the understanding of the generation of amplitude modulation in the near field of a turbine, notably by Oerlemans [5]. The reasons why AM is only rarely observed in the far field are not so well understood. Wind shear is often modelled as being logarithmic with height, or sometimes as a power law, but measurements have sometimes shown much more complex variation with height. Figure 1 shows an example of a low-level jet measured with a lidar. This shows wind speed increasing up to a maximum between 100 m and 200 m above ground, and decreasing at higher altitudes.



Figure 1: Low-level jet

Figure 2 shows some short-time samples from the jet of Figure 1. In each case, the grey curve is the result of fitting a power law curve to data from a 50 m met mast; the red dashed curve is fitted to the shear measured across the disc of a hypothetical turbine of 100 m diameter and 100 m hub height, and the red dotted curve is the actual wind speed measured by lidar.

It is clear that a mast of half hub height can dramatically underestimate wind shear across the turbine disc; it is also clear that the power law itself is sometimes inadequate to describe wind shear. It could be that unusual wind shear conditions occasionally lead to a dramatic increase in the level of modulation at source.

Although the wind speed decreases above this jet, it will increase again at greater height. This could lead to a channel, analogous to those encountered underwater [6], which could carry sound a significant distance downwind. It is even possible that caustics or convergence zones could exist, focussing the sound at specific ranges downwind.



Figure 2: Wind shear across the rotor disc

Essentials of the algorithm

Consider broadband noise whose peak amplitude varies between $\hat{p}_{min}$ and $\hat{p}_{max}$. The sound pressure variation over time can be written

$$p = \left[\frac{1}{2}\left(\hat{p}_{\max} + \hat{p}_{\min}\right) + \frac{1}{2}\left(\hat{p}_{\max} - \hat{p}_{\min}\right) e^{i\omega t} \right] \times noise \,. \quad (1)$$

Squaring the input signal results in

$$p^2 = \left[\frac{1}{4}(\hat{p}_{\max} + \hat{p}_{\min})^2 + \frac{1}{2}\left(\hat{p}^2_{\max} - \hat{p}^2_{\min}\right)e^{i\omega t} + \frac{1}{4}(\hat{p}_{\max} - \hat{p}_{\min})^2 e^{i2\omega t}\right] \times noise^2 . \quad (2)$$

The terms in brackets are:

a DC term of magnitude $\hat{p}_{DC} = \frac{1}{4}(\hat{p}_{\max} + \hat{p}_{\min})^2$, (4)

a tone of angular frequency ω and peak amplitude $\hat{p}_{\omega} = \frac{1}{2}\left(\hat{p}^2_{\max} - \hat{p}^2_{\min}\right)$, (5)

a tone of angular frequency 2ω and peak amplitude $\hat{p}_{2\omega} = \frac{1}{4}(\hat{p}_{\max} - \hat{p}_{\min})^2$. (6)

Expressions (4) and (5) can be combined to give a measure of the variation in amplitude

$$\frac{\hat{p}_{\max}}{\hat{p}_{\min}} = \frac{2\hat{p}_{DC} + \hat{p}_{\omega}}{2\hat{p}_{DC} - \hat{p}_{\omega}} . \quad (7)$$

Practical implementation

A microphone signal is digitally sampled and divided into time segments. The length of each time segment must be at least the reciprocal of the desired frequency resolution in the modulation spectrum. So, for example, to obtain a resolution of 0.1 Hz, the time segments should be at least 10 seconds long. It is convenient if the number of samples in each time segment is a power of two. If the initial sampling frequency is 44.1 kHz (a standard sampling frequency used in CDs), a segment of 2^{19} samples is about 12 seconds long.

The processing is outlined schematically in Figure 3, and discussed below.



Figure 3: Flow chart of the algorithm

Much of the broadband noise from wind turbines is in the range from a few hundred Hz to a few kHz; this is the noise whose modulation is to be measured. If higher frequencies are to be discarded to speed up processing, the signal can be passed through a low-pass anti-aliasing filter and down-sampled to a lower sampling frequency. This down-sampling is often known as “decimation”. As with all low-pass filters in this implementation, a Finite Impulse Response (FIR) filter can be used.

A band-pass filter can then be used to select a frequency band for further processing. Octave bands are suggested. It is a remarkable consequence of the convolution theorem [7] that the multiplication of a time history by $e^{2\pi i f}$ has the effect of shifting all frequencies in that time history up by a value f. By setting $f = -f_C$, where f_C is the centre of the frequency band of interest, that band can be shifted left so that it is centred on 0. This process is known as “mixing”. The band-pass filter can then be implemented as a low-pass filter, which is much more computationally efficient.

Figure 4 illustrates this process. The blue curve is the decimated input signal. The green curve is the same, but shifted left by about 1061 Hz (the arithmetic centre of the octave with geometric centre 1 kHz). Note that negative frequencies wrap around to appear at the top of the frequency scale: this is conventional with the Cooley-Tukey FFT algorithm. The red curve then shows the effect of a one-octave band-pass filter implemented as a low-pass filter.



Figure 4: Implementing a band-pass filter as a low-pass filter

The filtered signal is then squared.

At this point, it is possible to conduct a spectral analysis of the squared signal by Fast Fourier Transform (FFT). However, the result would be a spectrum running up to hundreds or thousands of Hz, depending on the bandwidth of the filter just applied. Yet the frequencies of interest in the modulation spectrum are likely to be of the order of 1 Hz. So, once again, it is sensible to apply an anti-alias filter and down-sample the signal. This is shown in Figure 5. In this case, the length of the filter may be very large, and significant performance gains could be achieved by applying a series of much shorter filters.



Figure 5: Effect of modulation spectrum anti-alias filter

Before applying the FFT, it is common practice in signal analysis to apply a window function to improve the ability to discriminate between tones that are close together in frequency, but this is of dubious merit in the present application as the tones are relatively widely spaced.

The output of this process is the modulation spectrum: a graph of modulation amplitude against frequency.

Results

Figure 6 shows the time history of a signal sampled at 44.1 kHz and chopped into segments of 2^{18} = 262,144 samples. Subjectively, this sample begins with a gradual “swish” sound, which towards the end changes to a more abrupt “thump”.



Figure 6: Signal sampled at 44.1 kHz and chopped into segments of 2^{18} samples

Figure 7 to Figure 10 show the results of processing data in octaves from 250 Hz to 1 kHz, for those four time segments.

It might be expected that the higher harmonics would increase in shifting from swish to thump. In fact, the most striking change is in the relative amplitudes of the fundamentals in the different octaves. The swish is predominantly in the 1 kHz octave (Figure 7), while the thump is predominantly in the 250 Hz (Figure 9).



Figure 7: Modulation spectrum, 0 to 5.94 seconds



Figure 8: Modulation spectrum, 5.94 to 11.89 seconds



Figure 9: Modulation spectrum, 11.89 to 17.83 seconds



Figure 10: Modulation spectrum, 17.83 to 23.78 seconds

Conclusions

This procedure has many advantages. By transforming the noise into the frequency domain, many features become immediately apparent:

- Aperiodic modulation, such as caused by traffic noise or fluctuations in wind speed, will be automatically excluded.
- Periodic modulation will be detected as a peak or peaks in the modulation spectrum.
- Peaks due to periodic phenomena not associated with wind turbines will appear at different frequencies.
- Modulated noise from a wind turbine will be expected to occur at the blade passing frequency and its harmonics (multiples).
- If shaft rotational frequency is present, that suggests the source may be due to one blade more than the others.

The method proposed here is not new: it has been proved over many years. It can run in real-time. It provides an objective measure of amplitude modulation, subject to agreement on a number of parameters, such as the noise bandwidth and the roll-off of the various filters. The method can be used by environmental health professionals to assess complaints, by wind farm operators to manage their assets and by wind turbine manufacturers to refine their product.

The difference between swish and thump appears to be mainly in the frequencies of the noise being modulated, with swish being predominantly in the 1 kHz octave and thump being in the 250 Hz octave

Recommendations

If multiple harmonics are present, there is an issue of how best to quantify the modulation. It is hoped that RenewableUK's research programme will address this.

Amplitude modulation is commonplace in the near field of wind turbines, yet is rarely observed in the far field. On the rare occasions when it is observed in the far field, the cause is not well understood. It may be due to an unusually high amplitude at source, perhaps due to wind shear or wind veer across the turbine disc. Alternatively, it may be due to focussing of the sound by wind shear and temperature gradients. A detailed mapping of the wind flow by a steerable lidar could provide valuable insights into the origin of this phenomenon.

References

1 Lee, S. et al, An estimation method of the amplitude modulation in wind turbine noise for community response measurement, Third International Meeting on Wind Turbine Noise, Aalborg, Denmark, 17-19 June 2009

2 Nielsen, Richard O., Sonar Signal Processing, Artech House, 1991

3 Fundamental research in amplitude modulation - a project by RenewableUK. Bass J, Bowdler D, McCaffery M and Grimes G (UK), Wind Turbine Noise 2011

4 Wind turbine amplitude modulation: Research to improve understanding as to its cause and effect. Bullmore A, Jiggins M, Cand M, Smith M, Wright M, Davis R and Oerlemans S (UK), Wind Turbine Noise 2011

5 Oerlemans, Stefan, Detection of aeroacoustic sound sources on aircraft and wind turbines, Thesis University of Twente, Enschede, 2009

6 Urick, Robert J.,Principles of Underwater Sound, 3rd Edition, Peninsula Publishing, 1997

7 http://mathworld.wolfram.com/ConvolutionTheorem.html