Measuring Wind Turbine Coherent Infrasound

John Vanderkooy[1] and Richard Mann[2]

Department of Physics and Astronomy[1], Department of Computer Science[2]

University of Waterloo, Waterloo, ON, Canada, N2L 3Gl

jv@uwaterloo.ca,   mannr@uwaterloo.ca

Abstract

To extract the optimum coherent infrasound signal from a wind turbine whose rotation is not precisely periodic, we use an optical telescope fitted with a photodetector to obtain reference blade passage periods, recording these together with the microphone infrasound signal. Signal processing of the quasi-periodic microphone signal is then used to obtain periodic data, which are analyzed by an appropriate length DFT to extract optimum values for the fundamental and harmonics of the coherent signal. The general procedure is similar to order domain analysis for rotating machines and is thoroughly explained and illustrated with measurements and analysis from a number of different wind farms. If several turbines are measured by a single microphone with blade passage periods obtained from several separate reference tracks, it may be possible to retrieve separate useful coherent signals from multiple turbines by appropriate processing.

Introduction

The original impetus for this paper was the infrasound measurement of some Siemens model SWT−2.3−101 2.3 MW wind turbines (WT) which have a rather wide speed range. Manufacturer's specifications give a speed range of 6−16 rpm. We measured some units that varied in speed by over 50% during a 30-minute measurement. Such WTs apparently produce DC and electronically convert this to AC to feed the power grid, thus their rotational speed is not related to the grid frequency. Other WT machines have a much steadier pace, since they have geared AC generators directly attached to the grid. These generators typically have some slip, so their speed is not precisely that of the grid either, and is affected by the strength of the wind. For that matter, the electrical power grid itself is also not quite periodic. All of these factors make our analysis more pertinent.

There is much literature dealing with WT infrasound. The original pioneering work on spinning fan modes was by Tyler and Sofrin [1], and Hubbard and Shepherd [2] focused on wind turbines. The latter also showed aeroacoustic effects of the blade-pylon interaction. Recent work on this interaction by Dooley and Metelka [3] considers it to be a major cause of the infrasound. We do not consider any comparison with theories in this paper, but single out *only the measurement* of the coherent part of the WT noise, that is, the component related to its rotation. The wind itself makes wideband acoustic noise, especially in the infrasound region, and it is not known if the WT influences this background noise to any significant extent. We may study that issue in future work. Our infrasound measurements indicate that the random noise component is often similar whether one is near or far from a WT, and we therefore attribute this to wind noise. Thus we regard the meaningful part of the infrasonic noise from the WT as being mainly the coherent component.

Many repetitive processes have relatively stable but not exactly equal periods. Signals obtained from such quasi-periodic processes may show narrow lines when a DFT is applied to a long dataset, but inevitably the lines will contain at least a few frequency bins, often many more, even though the relative width of the line is very small. If the periodic signal is accompanied by some other spuriae or random noise, it will be difficult to obtain the coherent part of the waveform by using the spectral peaks of the fundamental and harmonics and applying an inverse DFT. In a typical acoustic situation, the microphones may be responding to a number of sources. The methods in this paper offer considerable hope of sorting out such sources.

Analysis Method

It is only when perfect periodicity obtains in a dataset, that a DFT can be used to best display its repetitive nature. In addition, if we want the fundamental and each harmonic to lie precisely on single frequency bins, then the length of the data must be exactly an integral number of periods. This may be difficult, since the data acquisition sampling control is usually independent of the incoming data. In our case this was not a limitation, since we sampled the data at 12.8 kHz, in order to record the audible output of the WT as well. The blade passage frequency of the WT is usually below 1 Hz, so that each period of the data will have many samples. The problem remains that the data contain many periods of varied lengths. The analysis that we use will surmount this difficulty.

There is a tacit assumption in our approach; we assume that the underlying coherent or repetitive waveform has the same relative shape with respect to each blade passage period of the system being measured. This assumption is common sense for most processes such as rotating machines, but may become weaker as the period shows more relative variation. Later we show some measurements that bear on this assumption. We are forced to accept it for any process that is obfuscated by significant noise, since only then can we mitigate the effects of noise. We can think of the recovered final coherent signal as having been constrained to represent a fixed average frequency. From the point of view of rotational machine measurements, we might describe the data now as being in the revolution domain, rather than the time domain. The approach is similar to order tracking [4], which normally employs a sensor to detect the rotational position of the machine, but there are differences. We only have a single pulse in each period, so we must resample the many samples in each period to conform to the revolution domain.

The analysis to follow allows even rather irregular WTs to be analyzed so that they appear to have perfect periodicity. In order to optimize recovery of the coherent output, we used a small telescope with a CdSe photoresistor mounted at the eyepiece to obtain a signal locked to the blade passage occurrences. The photocell was in series with a resistor and biased by a 1.5V battery. The telescope was actually a 20X−60X Bushnell spotting scope. Each time a blade passes the telescopic field of view, the CdSe detector would produce either a pulse of positive or negative polarity, depending on whether the blade was brighter than the sky, or the inverse. This signal was recorded along with the infrasound and other data, and was used to regularize the blade passage and other signals. This greatly increases the contrast of spectral lines in the presence of other random noise signals, which are mostly wind noise for the measurement of WTs. We outline our signal analysis procedure in the sequel to demonstrate its effectiveness.
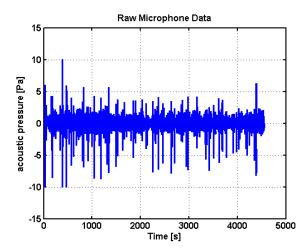
Our telescope was designed for daytime use, and may suffer signal confusion depending on the background sky condition. A clear blue sky gives very clean signals, since it is darker than the blades. To minimize environmental factors, the assembly could be fitted with an infrared or visible laser to produce a narrow outgoing interrogation beam and a receiving telescope with appropriate optical filters, and used under almost any conditions,. Our procedure simply needs a reliable fiducial signal that paces the quasi-periodic process. It is clear that the concepts of this paper could be gainfully applied in any analogous situation, to a variety of devices.

Our acoustic measurements were taken using GRAS 40AZ microphones with constant current preamps. We used a DC coupled National Instruments 9234 data acquisition interface, using Labview® software. We verified the microphone response by measuring the microphone in a sealed box with a driver having a plastic cone and rubber surround. The combination was down 3dB at 0.27 Hz. Since our WT data has harmonics that go from about 0.7 Hz upwards, we have simply used the microphone data without correction. We could have used the measured response to equalize the actual measurements, but there would have been little difference.

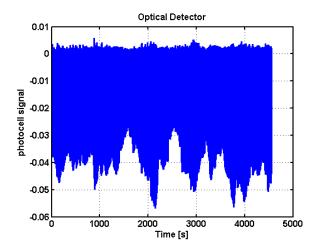**A Turbine Measurement and Its Analysis**

In the details of the example to follow, we will use typical numbers from measurements of the Amaranth Wind Farm, also known as the *Melancthon EcoPower Centre*. This wind farm has 133 GE 1.5MW turbines located in Melancthon township near Shelburne, Ontario, Canada. A location about

420m downwind from a nearest turbine was chosen, and data were captured over about 4500 seconds, at a sampling frequency of 12.8 kHz, exactly 58,307,840 samples. The blade passage period normally was about 1.08 seconds. We decimated the acoustic and the optical blade passage data to 800 samples/second using proper anti-aliasing filters with MATLAB®, leaving 3,644,240 samples in the acoustic data, d[n], and the same number in the blade passage reference array, r[n]. Fig.1 shows the raw microphone data d[n]. There are a few places where the system overloaded, and we simply zeroed these overloads in further analysis. The thick lines in the plot make it seem that the spikes have a significant energy, but in fact they are very sparse relative to the millions of samples.
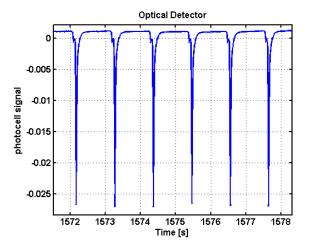

Figure 2. Photocell signal from the telescope trained on a WT to record blade passage times.

Figure 1. The acoustic microphone data, downsampled from an original 12800 samples/sec to 800 samples/sec.

The corresponding reference blade passage signal from the telescope, r[n], is shown in Fig.2, showing some variation in amplitude over the measurement period. Fig.3 shows an expanded portion of the blade passage signal where the signal was weakest, illustrating that even for this portion the reference pulses are very clear. The period in this region is about 1.1 seconds, but it varies throughout the record, as we shall show later.
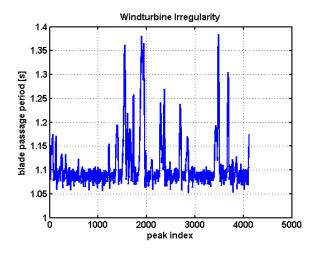

Figure 3. An expanded view of a portion of the photocell signal shown in Fig.2, showing very clear blade passage times.

We usually rely on a sufficiently-high sampling rate of the reference signal r[n] (and also of the data d[n]) that the period markings can be deemed to occur at the samples nearest the fiducial points of the blade passage signal, although this is not necessary in principle. If the reference array is of high quality, it could be interpolated to obtain a fractional sample value for each blade passage period marker. We proceed with the full fractional sample mathematical approach, and comment on the approximations that we have used.

An important detail of the operation of the WT is its irregularity or variation of blade passage periods. To determine this, we analyze the r[n] array, and detect the value of n at each peak of the blade passage

signal. We could compute the fractional value of n by interpolation, but with a sampling rate much greater (often 800 samples/sec or more) than the period (~1.1 sec) we would normally simply use the integer value. We call the array of index values at the peaks R[k], where k is the peak index (here ranging from 1 to 4116). There are R[p]−R[p-1] samples in the p-th period. The array R[k] can have fractional index values, and we maintain this generality in what follows, although in practice integer values are often sufficient. Fig.4 shows the time spacing between the peaks for the whole reference array r[n]. Although the dominant period is about 1.08 seconds, there are excursions up to 1.38 seconds, which may represent partial stalling of the WT. It is clear that this particular unit is fairly erratic, and if we simply did standard DFT studies of the data array d[n], the results would be far from optimal.
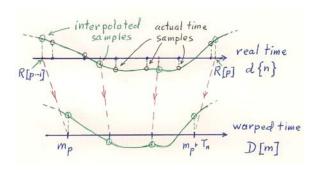

Figure 4. Period irregularity of the turbine observed by the telescope. The normal period is about 1.08 seconds, but the machine slows down significantly on a number of occasions.

What can now be determined is the average period, or here perhaps the dominant period, in the signal r[n]. This could be done by counting the P pulses, here 4116, that represent P−1 periods, noting the total number N of samples involved, representing N−1 sample steps. This gives about 885 samples in this example, which we round to an integer, Tn. This represents 1.1 seconds for the period. From Fig.4, though, we see that 1.08 seconds, or Tn=864 samples, would be a better choice. The actual integer that we pick is not really that important, but sticking

to sensible values will keep our final analysis numbers close to the actual ones. For each measurement, a plot such as Fig.4 will give us a good idea of the average or dominant period, or we could use the calculated one.

We now define a *regularized* output data array, D[m], in which m represents an integer sample index that is meant to be like n, but with *exactly* Tn samples between reference events. D[m] will also have size close to 3,644,240 samples. The regularity of D[m] is essential to make the output data perfectly harmonic. Because we are forcing the output array to be processed in strictly periodic portions, the number of samples in D[m] will not be quite the same as in d[n], but it will be close.

The algorithm's basic approach is to associate the samples in each period of the irregular data, d[n], with the corresponding period in the regularized output array, D[m]. The heart of the algorithm is that R[p]−R[p−1] interpolated samples in real time (about 885 on average), in the real-time p-th period of d[n], will be resampled into exactly Tn warped-time samples (here 885), which constitute the p-th period of D[m]. Fig.5 illustrates how the p-th period of the array d[n] shown on the upper line is resampled to produce the p-th period of D[m] in the lower line. The upper green line is the bandlimited interpolation of the actual time samples, d[n], and the green interpolated samples at prescribed fractional indices given by R[p···] are transferred to the output array, D[m], at integer sample positions.

Figure 5. Illustrating how interpolated variable-length periods of data, d[n], at fractional indices, are mapped to equal-period integer index values constituting the warped-time output D[m].

R[p−1] is the sample index of d[n] at the start of the p-th period (which is fractional in general), and m_P is the integer starting index of the p-th period in D[m]. The sampling rate ratio between the data array d[n] and the output array D[m] will normally be very close to unity. For this particular period, the ratio is Np/Tn=(R[p]−R[p−1])/Tn, and the j-th local integer sample of D[m] in the p-th period must be associated with the local fractional sample j (R[p]−R[p−1])/Tn in d[n]. R[p] is the real-time end index of the p-th period (which again is fractional in general) of d[n], and m_P+Tn is the end index of the p-th period of the warped-time array D[m]. We will use {} brackets for fractional indices, so for fractional index, x, the interpolated data values are denoted d{x}. The care we take for the interpolation can be decided by the nature of the reference data. If there are many samples per period, an interpolated sample might just be the nearest actual sample. Fig.5 shows 4 equally-spaced interpolated fractional-index real-time samples of d{n}, which are resampled into 4 warped-time integer index samples of D[m]. If there are many samples in each period, then we might just use nearest index samples, and in Fig.5 this would mean that 2 or 3 samples in real time might be ignored.

If we start our analysis at a reference point representing the start of a period, we can call this the 0-th sample in both d{n} and D[m]. Then the first period will go from sample R[0] to R[1] for the data array, d{n}, and from sample 0 to Tn for the output array, D[m]. The reference samples are at fractional time indices, in the general case of an interpolated version of r[n]. The index values of the array R[p] are then fractional, as already indicated. If j represents the local index in the p-th period of the regularized output data D[m] (a convenient aspect for the coding), then the j-th sample of D[m] in the p-th period of the data is calculated from d[n] by interpolation as:

$$D[(p−1)Tn+j]=d\{R[p−1]+j(R[p]−R[p−1])/Tn\}.$$

$$(1)$$

The index j goes from 0 to Tn. Index 0 belongs to the starting sample or the last sample of the previous period. Our use of the {} brackets implies interpolation of d[n] for fractional sample numbers. R[p−1] is the starting index in d{n} of the p-th period, starting the count at p=0, and the fractional local index is j (R[p]−R[p−1])/Tn.
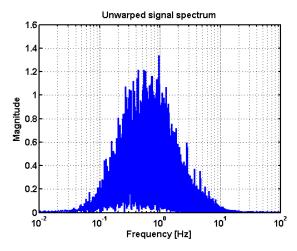
The sophistication of the interpolation that we use can be set by accuracy requirements, oversampling factor, and computing time. To repeat, we might simply use integer values for the array R[p], together with the nearest sample in d[n], given by rounding the index calculation:

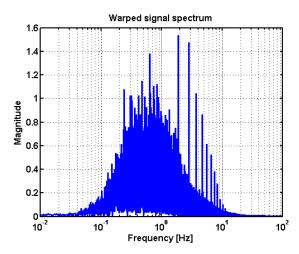$$D[(p−1)Tn+j]=d[round(R[p−1]+j(R[p]−R[p−1])/Tn)]$$

$$(2)$$

If the data are highly oversampled relative to the repetition period, then this nearest sample interpolation is justifiable. The RESAMPLE command in Matlab calls this *nearest neighbor interpolation*. It amounts typically to a very slight time requantization error of the samples. In this approximation, either some samples of d[n] will be omitted in D[m] if there are more samples in the d[n] period than Tn, or some samples of d[n] will be used twice in the output array D[m] if Tn is larger than the number of samples in the period of d[n].

To show the effectiveness of the resampling procedure, Fig.6 shows the magnitude spectrum of the original acoustic data array, d[n]. This plot contains over 450,000 frequencies, since it covers ¼ of the range of the Nyquist frequency of 400Hz, and it represents the DFT (discrete Fourier transform) of just over 3.6 million samples. There are a few lines on the high-frequency side of the main peak that may represent harmonics of the blade passage period, but they are not clear. The spectrum peaks at about 0.5 Hz, which is close to the frequency limitation of the GRAS 40AZ microphone. We can correct for the falling microphone response, and then the spectrum tends to stay flat down to well below 0.1 Hz, but that is not the focus of this paper. Suffice it to say that the wind itself seems to be the cause of these very long-term fluctuations.
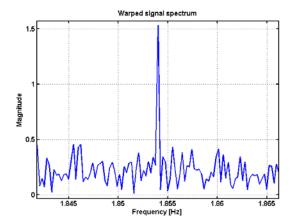
Figure 6. Spectrum of the recorded microphone data as measured with its original time samples.

Fig.7 shows the DFT spectrum of the regularized array, D[m]. Note that there is now a very clear harmonic series with fundamental frequency just below 1 Hz. The lines are extremely sharp, and are each essentially in only one frequency bin. This is shown in Fig. 8, which isolates a portion of the highest 2nd harmonic line.

Figure 7. Spectrum of the microphone data, resampled to be *exactly periodic* by using the reference blade passage signal which was recorded simultaneously with the original microphone signal.
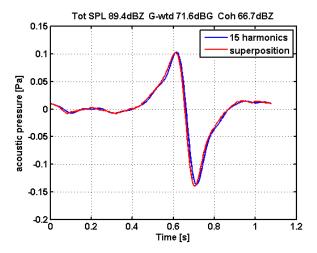

Figure 8. An expanded view of the 2nd harmonic line of the resampled, regularized data, showing that it is only one frequency bin wide.

The fundamental frequency, approximately 0.9Hz, does not show an amplitude distinct from other spectral lines near it, but it is there nonetheless. Other measurements [2] and theory [3] also show weaker fundamental, but stronger harmonic lines.

The harmonic lines are so prominent that they can be used to determine the shape of the acoustic pulse, thereby removing the effect of all the other spectral lines, which represent noise. Fig.9 shows the result of extracting the harmonic lines, making sure that the resulting complex spectrum is properly conjugate even, inverse transforming the result, and displaying exactly one period of length Tn samples. Also shown in the figure is the result of averaging all the periods of D[m] to reduce noise. The plot follows the harmonic inverse very well. It is surprising that the noise has averaged down so much.
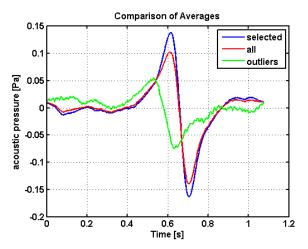
The total unweighted SPL of the microphone data is 89.4 dB, often denoted as 89.4 dBZ. When G-weighted it reduces to 71.6 dBG. This weighting is intended to express the threshold of infrasound perception [5], and may be well below the unweighted value. The unweighted SPL of the coherent pulse signal is 66.7 dBZ. We have averaged about 4100 periods, which reduces the noise amplitude by sqrt(4100) or 36.1 dB. Thus the noise level falls to about 89.4−36.1 or 53.3 dBZ, well below the coherent SPL of 66.7 dBZ.

Figure 9. Showing the coherent infrasonic pulse (blue) obtained by using 15 harmonics of the spectrum of Fig.7. The red curve is the result of averaging all the periods of the resampled data. Although the noise level is considerably larger than the signal shown in this figure, averaging over 4000 periods removes most of this noise.

Each of the coherent pulse waveforms that we show in the sequel have unweighted total SPL, G-weighted total SPL, and unweighted SPL of the coherent pulse included in the title bar. G-weighting enhances spectral components between 10 and 28 Hz, but diminishes those outside that range. In each case G-weighting is lower than the unweighted SPL, since a significant part of the infrasonic energy comes from very low frequencies, which are suppressed by such weighting. G-weighted signals are typically much higher than A-weighted ones, but A-weighting does not apply to infrasound at all. It is also questionable to apply G-weighting to the coherent pulse signal. That would reduce the amplitude of the pulse, but it is perhaps the peak acoustic amplitudes that matter.
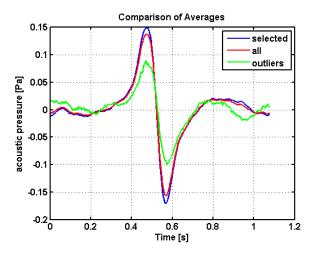
It is worth pointing out that the DFT of the period-averaged data will have exactly the same harmonic components of the periodic pulse as the DFT of the full array D[m]. This is because the full array has precisely an integer number of pulse periods, and the Fourier sums over each period in the full array repeat the same multiplicative functions, adding to produce the same result as averaging all the periods first and then taking the smaller DFT. We do not require a DFT of the complete sequence, unless we are interested in components that have a different periodicity.

We have assumed that even those periods, which deviate from the dominant 1.08 seconds, have the same pulse shape as the normal ones. This may not be true, and to test this, we have separated out periods greater than 1.12 seconds from those that are regarded as normal, and given an average for each group. Fig. 10 shows the average of 3099 selected normal periods, together with 697 abnormal outliers, and also the result when all periods are included. Note that the largest signal occurs for the normal periods, while the longer abnormal periods have a pulse waveform that is both smaller and occurs earlier. This time advance should actually be expected! The distance to the turbine was about 420m, and thus it takes about 420/343 seconds for the acoustic signal to travel to the microphone. If a period is longer, then its acoustic response will occur relatively earlier since the travel time to the microphone is a smaller fraction of that period. Rough estimates of the period lengthening show that the pulse advance is about right. Note that because all the outlier periods are averaged together, that bipolar pulse is weakened due to the distribution of pulse timings.

Figure 10. The coherent output signal for normal periods (blue) is somewhat larger than the average of the normal ones plus outliers (red), while the outliers themselves (green) are somewhat wider and also time shifted.

Now that we understand the effect of the acoustic delay, we can correct for it, now and in the sequel. Fig.11 shows the extracted infrasonic pulse with the correction. Notice that the pulses now all line up, and the amplitude for all the pulses is closer to the

selected ones. The time delay caused by the 420m distance to the WT is almost one blade passage period, so the corrected pulse position in Fig.11 is in about the same position as the original data of Fig.10. The outliers are still lower in amplitude than the selected ones, but that is reasonable, since we would expect the strength of the pulse to be related to the pressure difference across the blades, and this should be less at lower blade speeds. The position of the pulses relative to the reference pulses is not controlled in our plots, so it may not appear centred in later plots.
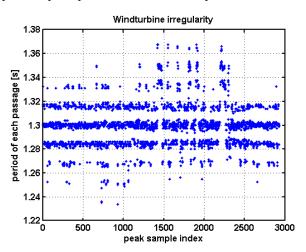

Figure 11. Coherent infrasound pulse when the time of flight from turbine to microphone is taken into account. All the periods including the outliers are time aligned and more alike in amplitude.

One drawback to the above procedure is that the DFT required to analyze the data will have N=P×Tn points, and that number may be poorly conditioned for an efficient FFT algorithm. Although N is already composite, it helps if P and Tn are each factorable into smaller factors. Since our final repetition period is simply an average of the real, varying-length periods, we are at liberty to modify Tn somewhat so that the algorithms are efficient! In fact if there are many periods, P, then we could throw away a few in order to optimize P as well! Changing P or Tn by even a few counts may vastly improve the FFT execution time. Such considerations are not trite, since we may be using many millions of sample points, and patience at the computer is not necessarily a virtue. Although not necessary, in the example presented above, we chose Tn=863, and the DFT was chosen with $N=2^{17} \, 5^2 = 3276800$, which covers most

of the original data and represents 3796.98725 periods, very close to an integer. Typically the exact DFT length has very long execution times. In those instances for which we do take the time to do an exact DFT of all the periods, the results look the same.
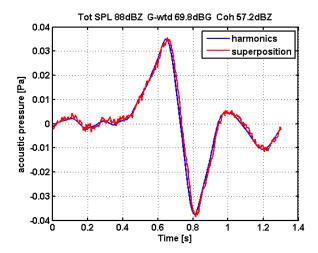
Other Measurements and Analyses

Each of the 5 wind farms that we have measured displayed quite different operating characteristics. Fig.12 shows the WT irregularity over more than an hour of a unit in the Summerhaven Wind Energy Center of 56 Siemens 2.3MW turbines in Haldimand County near Fisherville, Ontario. The most common period is 1.3 seconds, but interestingly the unit 'jumps' between 8 or 9 different periods, with some fluctuations in each one. It almost seems that the turbine is altering its gearing ratio. On several occasions the turbine seems to slow down over many periods, perhaps due to lulls in wind speed.

Figure 12. Period irregularity for a turbine that exhibits "jumps" in its operation.
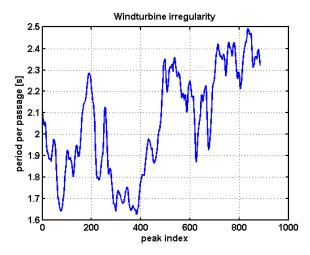
The microphone was generally downwind from the nearest turbine whose blade passages were recorded, about 500-700m away. Again a straightforward DFT of the raw data yielded rather poor spectral lines with significant width, while the warped revolution-domain data showed excellent lines. It proved easy to extract the coherent infrasound pulse from the data, shown in Fig.13. There was somewhat more wind in this measurement, and the total infrasound SPL was about 31dB below the coherent sound. Although the time-of-flight was corrected, it did not

improve things very much, because the relative period variation is quite low.

Figure 13. Infrasonic pulse of the Haldimand County turbine. The noise is again quite low due to averaging of the long data record.

One of our earliest rather enigmatic measurements was from the Conestogo Wind Farm near Arthur, Ontario that has 9 Siemens 2.3MW turbines, and a single 2.1MW unit. The period irregularity shown in Fig.14 is so weird that we did tests to make sure that our equipment was not malfunctioning. These turbines obviously are not geared to generators that are directly grid connected; we surmise that they generate DC power optimally and inject it into the grid electronically

Figure 14. Period irregularity of a Siemens WT showing 50% change in period over a measurement interval of 30 minutes. Such characteristics make our time-warping analysis essential.

The microphone signal seems normal, but these units give a somewhat weaker infrasonic pulse even though the distance to the turbine was only 200m. Fig.15 shows the recovered pulse. Perhaps these units are changing blade pitch or speed under a control mechanism that minimizes infrasonic signals. The microphone may not have been favourably located.
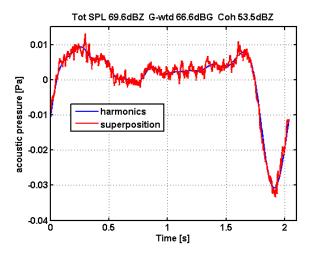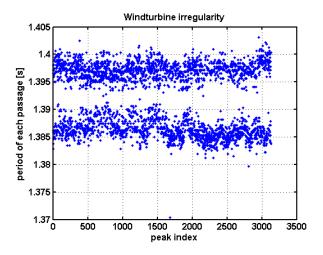

Figure 15. Coherent infrasonic pulse from a very irregularly rotating Siemens turbine. It would be impossible to obtain these data without the time-warping analysis of this paper. The relative noise is somewhat larger for the averaged data, due to the shorter time record, and the low level of infrasound from these machines.
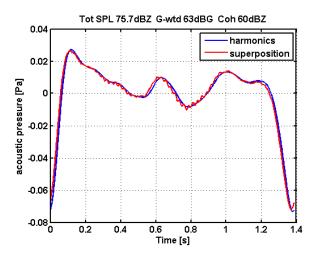
These Siemens units have a large spacing between blades and pylon, and this may be a factor in keeping with the ideas in [2]. Since these data were taken over only 30 minutes, it appears noisier than our other measurements. We should point out that without the time-warping order analysis, these data would not yield an infrasonic pulse at all.

A measurement made on a moderately windy day at the Enbridge Ontario Wind Farm with 110 Vestas V82 1.65MW turbines, located near Underwood, Ontario showed yet another pattern as demonstrated in the irregularity shown in Fig.16. The data were downsampled to 3200Hz from the original 12.8kHz file.
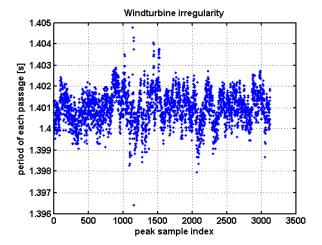
Figure 16. Period irregularity of a turbine that seems to be jumping between 2 rotational speeds.

This turbine may be changing its gear ratio to optimize the power gleaned from a fluctuating wind speed. Again it gives a very clear infrasound pulse as shown in Fig.17. The total SPL from the farm was over 16dB greater than the coherent portion. Note that the overall shape of this pulse differs from several of the others.
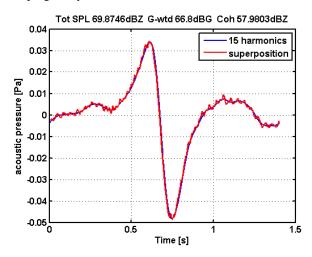

Figure 17. Coherent infrasonic pulse of the turbine whose period character is shown in Fig.16.

The irregularity of a turbine in the Mohawk Point Wind Farm of six 1.65MW turbines, located near Lowbanks, Ontario, is shown in Fig.18. These units have a very consistent period near 1.401 seconds. The relative variation is so low that we used the full 12.8kHz data to reduce quantization errors in determining the number of samples in each period.

Figure 18. Period variation of a turbine that is quite regular, with a speed variation of only 2%.

The extracted infrasound pulse is shown in Fig.19. The reduction of signals from the other turbines was still very good, even though these units are nearly but not quite synchronous. It is necessary to take at least an hour of data to achieve such separation, and the warping analysis is vital.

Figure 19. The infrasonic pulse extracted from the fairly regular turbine whose period character is shown in Figure 18. Even here the order analysis is crucial to obtain such clear data.

The data from our 5 chosen wind farms display a wide range of behaviour. Yet each instance benefits greatly from time-warping order analysis.

Multiple Simultaneous Measurements

If there are several sources to be measured, each having an optical or other reference track, but only a

single microphone or a common data track, can we separate out each of the quasi-periodic processes? The answer is that it may be possible, but only if certain conditions exist. We will not go into an exhaustive analysis, but indicate some avenues to explore.

The first thing that comes to mind is that using each reference track, we can analyze the common data to retrieve mainly the data component related to each reference. We can resample the data using each reference, and consider what crosstalk might not be avoidable. If the several processes are quite irregular and their irregularity is unrelated, we would expect a large DFT to isolate each process to a good degree, treating the others as noise, much as our earlier examples. If the irregularity is similar for each component, then it will be difficult to separate them out.

If the processes are only slightly different in average period, we may be able to choose a data length that captures a different integer number of periods of each. In principle then, each process should be virtually orthogonal to the others. A DFT will then show different lines for the fundamental and harmonics for each process. For more than 2 processes of close to the same frequency, it will be very difficult to find a data size that will treat all of them as close to an integer number of periods.

A method that we are currently assessing is to record multiple turbines in a wind farm with a video camera. With appropriate software, we can post-process the video stream to produce a number of blade passage records. The clock oscillators in modern video and data acquisition equipment are sufficiently accurate that there will be negligible drift during an hour or so of measurement.

Conclusion

Our paper shows how the coherent part of the infrasound from a single WT can be extracted from a microphone signal by using a blade passage reference track from the turbine under study.

Our analysis reveals a characteristic infrasonic pulse. We conjecture that the pulse from a single WT is caused by the interaction of the blades against the pylon, while the rather more complex background

signal relates to the radiation of the Tyler-Sofrin spinning modes.

The random component of the infrasonic signal exceeds the coherent part, and this random component is related to wind noise, which appears to be similar whether one is near or far from a wind farm.

Our paper avoids the issue of health effects from WT infrasound. Information on both sides of the controversy abounds in the literature.

References

[1]   J. M. Tyler and T. G. Sofrin, "Axial Flow Compressor Noise Studies," SAE Technical Paper 620532, 1962.

[2]   Harvey H. Hubbard and Kevin P. Shepherd, "Aeroacoustics of large wind turbines", J. Acoust. Soc. Am. vol. 89 (6), p.2496-2508, June 1991

[3]   Kevin A. Dooley and Andy Metelka, "Acoustic interaction as a primary cause of infrasonic spinning mode generation and propagation from wind turbines", presented at 166th Meeting of the Acoustical Society of America, San Francisco, California, 2 - 6 December 2013. Session 3aNS: Noise. Available on POMA website.

[4]   http://en.wikipedia.org/wiki/ Order_tracking_(signal_processing)

[5]   G-weighting: ISO 7196:1995, Acoustics – Frequency weighting characteristic for infrasound measurements.

[6]   A.N. Salt and T.E. Hullar, "Responses of the ear to low frequency sounds, infrasound and wind turbines", Hearing Research 2010, Sep 1; 268(1-2):p.12-21.