Wind Turbine Amplitude Modulation: Research to Improve Understanding as to its Cause & Effect

Work Package F (WPF) - Collation of Work Package Reports and Final Reporting

UNIVERSITY OF Southampton
Institute of Sound and Vibration Research

HOARE LEA

ROBERT DAVIS ASSOCIATES

University of Salford MANCHESTER

renewableUK
The voice of wind & marine energy

Andrew Bullmore, Matthew Cand
Hoare Lea Acoustics
Tel: +44 (0) 1454 201 020
Fax: +44 (0) 1454 201 704

140 Aztec West Business Park
Almondsbury
Bristol BS32 4TX

Audit Sheet

Revision	Description	Date	Issued by	Reviewed by
1	Draft for comment	23 February 2012	AB	MC
2	Revised draft	25 March 2012	AB	MC
3	Revised draft	03 April 2012	AB	MC
4	Revised draft	11 April 2012	AB	MC
5	Finalised draft	18 April 2012	MC	

Table of Contents

EXECUTIVE SUMMARY

Background

This report describes the results of theoretical and experimental investigations into the causes of amplitude modulation of noise from wind turbines, the subjective response to this phenomenon in terms of annoyance, and potential mitigation methods should it occur.

The work has been funded by RenewableUK and carried out by a consortium of UK companies and universities in partnership with the lead contractor, Hoare Lea Acoustics. Additional complementary research was commissioned from a specialist at the National Aerospace Laboratory in the Netherlands. The members of the research team were selected on the basis of their technical expertise and experience in the fields of aerodynamics and wind turbine acoustics. The research was overseen by a technical steering group and will be peer-reviewed by other specialists working in the field. The experience of the research team and steering group represent a balanced perspective both from those generally perceived to be involved with the wind energy industry and those who are likewise generally perceived to be wholly independent of the wind energy industry.

Amplitude Modulation – general

All wind turbines generate noise. The main noise source for modern turbines is 'aerodynamic noise' - the noise generated by the interaction of flow turbulence with the surfaces of the rotor blades. The noise is said to be amplitude modulated when its level (loudness) exhibits periodic fluctuations at a rate corresponding to the frequency at which a rotor blade passes a fixed point (the 'blade-passing frequency'[1]).

Amplitude modulation (AM) is always detected close to a rotating wind turbine, and is commonly described as 'swish'. The principal source of noise from the blades is trailing-edge noise, caused by the interaction of turbulence in the boundary layer (the slower-moving air close to the blade surface) with the trailing edges (the rear, thinner edges) of the rotor blades. Because this noise source has particular directional characteristics, an observer close to the wind turbine will hear the noise from each blade separately, particularly in the crosswind direction from the rotor as the blade comes down. These trailing edge directivity effects will be present at locations near a turbine even in completely uniform flow. This characteristic swish has been explained theoretically and demonstrated by measurements carried out prior to the current research.

For the purpose of this research, AM resulting from this trailing edge noise directivity effect is termed 'Normal AM' (NAM). Based on theoretical models and experience to date, NAM would not be expected to be apparent (except at insignificant levels) downwind or upwind in the far-field of a wind turbine.

However, on some wind farm sites, AM has been detected in the far-field down-wind from wind turbines. In some cases, the magnitude of the variations in noise levels is higher than predicted and the noise is described as being more impulsive in character, better described as a 'whoosh' or 'thump' rather than a 'swish'. These occurrences cannot be accounted for by the established trailing edge noise mechanism and it is therefore concluded that other source mechanisms or propagation effects must be responsible.

For the purposes of this report, AM phenomena with characteristics that fall outside those expected of NAM are termed 'Other AM' (OAM).

Reported incidences of OAM are relatively limited. However, wind turbine OAM is a recognised phenomenon and has been the subject of several publications and presentations at international

[1] For a modern multi-megawatt 3-bladed wind turbine a typical upper rotational rate is 16 rotations per minute, for which the blade-passing frequency would be 0.8 Hz (slightly less than once a second).

conferences in recent years. Since the causal mechanisms have not been understood to date, no specific information has been available to guide operators towards the likelihood of occurrence of OAM or remedial actions which may be required. Furthermore, where OAM is known to occur, there has been no universally accepted means of measuring its magnitude or determining whether complaints from neighbours are justified.

Scope of Research - Objectives

This project was directed towards:

- identifying the causes of OAM and therefore also potential methods of controlling it should it occur;
- defining a robust methodology for measuring amplitude-modulated wind turbine noise and an associated metric that represents the degree of modulation;
- determining a dose-response relationship between AM, as rated using the adopted metric, and annoyance, in a way that takes account of the time-varying characteristics of amplitude-modulated noise in addition to its 'average' level;
- dissemination of results to the research community and all involved in wind farm planning and operation.

The discussions in the present work are directed principally towards variable-speed pitch-regulated wind turbines of upwind configuration, the most common type of large turbine, although many of the principles will apply to other turbine types.

Organisation of Work Packages

The work programme was divided into 8 work packages (WPs) as follows. Principal contributors to each WP are stated, although the WPs are essentially inter-linked and inter-dependent and all contributors were involved (to a greater or lesser extent) in all WPs through frequent communications and regular project meetings.

WP	Description	LEAD
A1	Source generation effects modelling	NLR
A2	Fundamental Research into Possible Causes of Amplitude Modulation	ISVR
B1	Development of an Objective AM Measurement Methodology	ISVR
B2	Development of an AM Dose-Response Relationship	ARC
C	Collation and Analysis of Existing Acoustic Recordings	HLA
D	Measurement and Analysis of New Acoustic Recordings	HLA
E	Wider Dissemination of Results	ALL
F	Collation of Work Package Reports and Final Reporting	HLA

WPs A-D have been completed and the outcomes presented in the individual WP Reports, and collated and summarised in the present overarching report which has been produced under WPF. WPE will be achieved through the publication and dissemination of the results.

Existing data was obtained from a number of sources, both in response to general requests at the Fourth International Meeting on Wind Turbine Noise, Rome, April 2011 and also through personal contacts with others working in this field. This data was provided at no cost and thanks are due to the donors. This was supplemented by WPD, which required detailed and innovative measurements to be carried out at three wind farm sites to study the characteristics of OAM in more detail. The measurements were related, wherever possible, to turbine operational parameters, such as wind speeds and directions at each turbine and other relevant meteorological variables.

Causes of OAM

Of the potential OAM 'source' effects, the prime candidate is transient separation of airflow from each blade ('stall'). The turbine blades operate at an 'angle of attack' (determined by a combination of the incoming air velocity and the velocity from rotation). Above a given angle of attack (mainly determined by the air velocity and the blade profile), the air flow over the upper (suction) surface of the blade may detach, resulting in the generation of a region of turbulent air on a region of the blade surface (stall) and a loss of lift.

The noise generated by the interaction of the turbulent air in the stalled region with the blade surface will result in increased noise (compared with the un-stalled, attached-flow case). In consequence, stall occurring over a small area of each turbine blade in one part of the blade's rotation only (for example as it passes over the top of its path) will result in cyclic increases in noise level (and therefore OAM). Stall noise also has a lower characteristic frequency than noise from an un-stalled blade and, importantly, it will also exhibit different directivity. Based on a model developed as part of this work, this change in directivity in particular is predicted to result in significant modulation levels and downwind directions, which is consistent with observations of OAM made in WPD. Downwind directions are those in which the highest overall levels of turbine noise are generally experienced in the far-field of the turbines. This results from a combination of source directivity and propagation effects and would explain the different characteristics and impact of OAM when compared to NAM. Although the same model predicts radiation of OAM in the upwind direction, it does not account for propagation effects which would generally significantly reduce far-field noise levels in this direction, consistently with observations made as part of WPD. It should be noted that, if stall occurs round the whole of the blades' rotation, OAM would not occur, though the characteristics of the noise will change (with increased overall noise generation rather than modulation).

The effect of changing blade pitch to initiate blade stall was investigated on a test turbine (in WPD). On other sites, the observed directivity of the OAM in the far-field was found to be highest downwind and more limited cross-wind. In each case, the observations were consistent with the 'transient stall' model.

Non-uniform air flow into the rotor means that the angle of attack of a blade varies as the blade rotates even when the pitch remains the same. This means that the blade may stall over a small part of the rotation. A number of factors which could lead to non-uniform inlet flow are identified in WPA. These factors include non-uniform wind profiles, for example due to a vertical or lateral variation in wind speed, or a spatial variation of the angle of the wind onto the rotor, where high wind shear or local wind gusts could provide the conditions for this to happen.

However, some of these factors are almost always present, for example variations in wind speed with height above the ground (vertical wind shear). Theoretical analysis made as part of WPA suggests that the effect of vertical wind shear cannot, in itself, account for the observed characteristics of OAM, and that partial stall was required in addition. Furthermore, OAM has been observed (on one site studied in WPD) when wind shear has been low but absent when wind shear was high (for the same wind speed) indicating that high vertical wind shear is not an essential factor.

It is also conceivable that particular blade designs could experience partial blade stall in certain conditions without the influence of non-uniform inflow conditions to trigger it. Turbine blades are also subject to aero-elastic effects (they bend and twist under load) and the resulting changes in angle of attack may be a factor (although not assessed in this report) in the occurrence of transient stall and therefore OAM.

Non-uniform turbulence distribution entering the rotor disk, due to upwind obstructions or meteorological conditions, could in theory cause time-varying levels of inflow turbulence noise as a result of the blades passing into and then out of the regions of higher turbulence, although this is not strongly consistent with the evidence (as determined in WPD in particular).

Whether or not a wind turbine on a particular site will exhibit OAM is therefore dependent on a large number of complex factors, including the local atmospheric conditions (particularly variation in wind speed and direction over the area of the rotor disk), local topography (which may influence rotor inlet flows in different wind directions) the design of the turbine blades and the way they are controlled. Several of the potential causal factors which have been suggested in the past have been shown, through the review of WPC and measurements of WPD, to have little or no association with the occurrence of OAM. However some of them may represent potential contributory factors. It is not therefore possible to be prescriptive as to whether any particular site or wind farm design is more or less likely to give rise to OAM being generated. This is considered likely to be due to a combination of site and installation-specific factors, including meteorology.

Where a wind installation exhibits OAM, it is then natural to consider how it can be assessed in terms of annoyance, and, in the event that the assessment shows that OAM requires to be mitigated, how this can be achieved.

Measurement of OAM

To devise an assessment method for AM (NAM or OAM), it is essential first to define the magnitude of this modulation. The ‘modulation depth’ or difference between ‘peak’ and ‘trough’ noise levels has often been used to date. What is important here is that any such method should be representative of the subjective response to amplitude-modulated noise and should be able to be robustly and objectively applied to real measured data. This is not as straightforward as may first appear, as different analytical methods applied to the same signal can produce a wide range of values (metrics), each of which may be valid in rating the severity of AM in terms of subjective response.

WPB1 has considered the philosophy of devising metrics for rating AM. The key feature of OAM that assists in its detection and analysis is the fact that the noise has a periodic character. Using this, Fourier-transform-based signal analysis techniques were used to objectively identify the modulation frequency in a noise signal. The magnitude of the variation in level of the signal at that frequency can then be rated in an objective manner. This is often not possible from the manual review of measurement results in realistic conditions due primarily to the compounding effects of spurious, non-wind turbine related, noise such as bird song that may equally vary with time and affect the measured levels of individual peaks and troughs of the OAM noise.

In parallel with defining an appropriate metric, these methods were also shown to be effective for detecting wind turbine AM ‘automatically’ in a measured noise signal by post-analysis of continuous measurements. This is an essential tool since OAM, where it occurs, is infrequent and its onset cannot generally be predicted, although in some cases experience may indicate that it is more likely to occur in some ranges of wind speed and wind direction. Most sources of extraneous noise can be excluded by applying signal filtering prior to the Fourier analysis, therefore focusing the analysis on the audio frequency bands in which the modulation is most significant and thereby providing a more robust procedure.

The results of analysis of field recordings of OAM, undertaken in WPB1, WPC, WPD and WPF, show that such methods perform well, even in challenging conditions present in rural environments.

How People Respond to Amplitude-Modulated Wind Turbine Noise

An extensive series of listening tests was carried out as part of WPB2 in a specialist facility at the University of Salford. This was done to establish if and how noise with a modulating character can be more annoying than steady noise of the same measured level, in order to supplement existing published information on this subject.

Simulated recordings, based on an analysis of actual field recordings, and with a wide range of input parameters, were played back to a range of up to 20 subjects of different ages and sensitivity but of normal hearing.

The frequency spectra and levels of sounds were intended to represent the varying characteristics of wind turbine AM as it might be perceived in a rural garden. Subjects were asked to rate the noise in two ways: on an absolute annoyance rating, and with a rating relative to un-modulated noise (by the subject adjusting the levels of modulated and un-modulated noise to achieve the same annoyance rating), with the presence in some cases of background noise with a spectrum and character representative of a rural garden.

Responses were not significantly affected by the frequency content of the modulated noise (either dominated by medium or lower frequencies), the modulation waveform or the presence of limited amounts of wind-disturbed vegetation noise.

The annoyance ratings were, however, significantly related to the frequency (rate) of the modulation, the overall A-weighted level (or loudness) of the test sound and the modulation depth. It was noted that the term 'modulation depth' has no accepted definition and the exact value depends on the protocol adopted for analysing the modulated signal. This factor highlighted the central importance of relating any specific response to measured AM levels in a consistent way.

Annoyance ratings were correlated with mean noise level and a range of metrics defining the degree of modulation. This showed that annoyance increases slightly with modulation depth. However, the observed effect is continuous with there being no evidence of a clear onset of increased annoyance at a particular modulation depth, particularly when considering the large spread of ratings. In contrast, the mean overall noise levels were shown to dominate the annoyance rating.

The tests for which an un-modulated wind turbine noise was adjusted for comparable annoyance with the AM stimuli resulted in levels which were relatively constant from modulation depths of approximately 3 dB(A) upwards. The adjustments were on average 1.7 dB(A) for a 40 dB(A) stimulus and 3.5 dB(A) at 30 dB(A). The use of the L_{A90} (the A-weighted noise level exceeded for 90% of the time over a certain period) as a measure of noise levels produced comparable results for moderate modulation depths.

The study necessarily relied on tests carried out under controlled laboratory conditions. Rating annoyance is subject to contextual and attitudinal issues and these factors are thought to be responsible for the wide error bands in the plotted data. However, the results of the WPB2 listening tests are generally consistent with those of existing research into subjective response to amplitude-modulated noise.

The effects of additional factors, such as the frequency of occurrence of OAM 'events' and their duration, were not addressed in the current research. The view of the steering group and the research team is that the significance of these factors would have to be assessed using professional judgment and experience and that it would not be practicable to assess them via further subjective testing.

Can OAM be Effectively Mitigated?

There is nothing at the planning stage that can presently be used to indicate a positive likelihood of OAM occurring at any given proposed wind farm site, based either on the site's general characteristics or on the known characteristics of the wind turbines to be installed.

In the immediate term, the only guaranteed solution to mitigate OAM if it occurs in practice on a particular site is the cessation of operation of offending turbines during those conditions under which OAM is found to occur. The conditions leading to OAM, and the characteristics of that OAM when it occurs, appear to be very site-specific and would therefore need to be established specifically for each operational site considered.

Notwithstanding the above, it has also been concluded that, as the existence of OAM in the far-field requires some effect to have occurred at source, even though the effect may be exacerbated by propagation effects, the control of the source effect will remove the OAM experienced in the far field.

Given the characteristics of the partial stall mechanism identified, the effective mitigation of OAM in practice will require the future involvement and close cooperation of wind turbine manufacturers, and possibly involve detailed measurements that focus on better understanding the surface pressure distributions on the turbine blades themselves, particularly as the stall point is approached. Simple analysis methods have been developed to assist in identifying the most likely relevant conditions. It is believed that with such cooperation, methods will be capable of being developed for avoiding local stall conditions.

Such methods may involve software ‘fixes’ that seek to modify the logic of the control system algorithms, perhaps even through the application of more advanced cyclical pitch control. More fundamental, physical design changes may also prove worthwhile, such as innovative blade designs or the addition of blade vortex generators which may delay the onset of stall. Such methods would be likely to only have a limited or negligible impact on the generating capacity of the turbines.

1.0 INTRODUCTION & PROJECT OVERVIEW

This report presents the summary findings of a research project awarded by RenewableUK (ReUK) in March 2011. The research project is entitled 'Wind Turbine Amplitude Modulation: Research to Improve Understanding as to its Cause & Effect'. The project comprises a total of six separate work packages, the outcome results of each of the work packages have separately resulted in their own dedicated final reports. This is the final report of the seventh work package (WPF) which consists in an overarching final report in which the key findings across the separate work packages have been collated and discussed.

Prior to embarking on any detailed discussion of the project, it is useful first to define what constitutes 'Wind Turbine Amplitude Modulation', as referred to in the project title, and also to clarify the definitions of some key terms that are subsequently used throughout the reporting of the project. This is considered important because the subject matter can, in parts, rely on quite detailed theoretical knowledge of several aspects of acoustics, including aero-acoustics, noise propagation, signal processing and subjective response to noise. In this respect, the reader is also referred to the glossary of terminology which is included at Appendix A to this report.

What is 'Wind Turbine Amplitude Modulation'?

When considering noise, the term 'amplitude modulation' refers to any noise whose amplitude (perceived loudness) modulates (i.e. goes up and down) over time. Over short periods of time these amplitude modulations may repeat themselves with an almost constant period, thus resulting in what is termed 'quasi-periodic' noise. There are many common sources of noise whose amplitude modulates in a quasi-periodic manner over time, examples being emergency vehicle sirens, the passage of cars in a steady stream of traffic, the regular chirping of birds, etc. However, to this list can be added wind turbine noise, for the reasons set out below.

As the blades on a wind turbine pass through the air they produce aerodynamic noise. It will subsequently be discussed how this aerodynamic noise arises from a number of separate source generation mechanisms but, for the present, it is merely noted that the intrinsic nature of this aerodynamic noise, as it is generated at the blade surface, is generally a continuous type sound. However, this continuous 'whoosh' is not what is experienced by anyone standing on the ground close to a turbine. Due to the rotation of the rotor, any such listener will experience the tips of the blades successively moving towards them and then receding away from them. As the dominant source of aerodynamic noise is radiated from the trailing edge of the blades (the thin end of a cross section through a blade) in the region lying towards the blade tips, and as the noise level at any given receiver location varies as a function of distance from the source, the movement of the each blade tip towards and then away from the listener will result in the noise level experienced by that listener going up in level as a blade tip approaches them and then down as a blade tip recedes from them.

This up and down variation in noise level will repeat with the passage of every blade. As a consequence, the perceived noise will be amplitude modulated (i.e. it will go up and down) at a rate equal to what is termed the 'blade passing frequency' of the turbine, or BPF for short. For a modern, large scale wind turbine having three blades and an upper rotation rate of 20 revolutions per minute, the blade passing frequency (and hence also the modulation rate) typically equates to 3 blades x 20 revolutions per minute, which equals a modulation rate of around 60 per minute. This is more commonly expressed as a modulation rate of 1 per second, which is alternatively expressed as a modulation frequency of 1 Hz[2]. The fact that most modern, large scale wind turbines are now variable speed machines[3] means that the frequency of this amplitude modulated aerodynamic noise may vary between anything from around 0.5 Hz to around 1 Hz, the actual frequency depending on the actual rotational speed, and hence blade passing frequency, of the turbine under consideration. Smaller turbines, such as were prevalent in the 1990s or still exist nowadays in the smaller end of the market, tend to rotate faster. The result of this faster rotational rate is that the time between each blade passage is reduced, with a consequent increased rate of modulation of typically double the values discussed above.

[2] Modulation frequency should not be confused with audio frequencies, as discussed below. See also the Glossary at the end of this report.
[3] Variable speed machines will adapt the speed at which their rotor turns depending on wind conditions.

Thus the 'amplitude modulation' of wind turbine noise, which is now commonly termed 'AM' for short, is nothing more than the quasi-periodic variation in the perceived level of aerodynamic noise as experienced by a listener standing relatively close to a turbine. This feature of wind turbine noise has long been accepted as a natural consequence of wind turbine operation, certainly ever since the introduction of the first commercial wind energy installations in the UK, which occurred some two decades ago now. The question must therefore reasonably be asked as to the need for the present project given that AM is such an accepted feature of wind turbine noise. The answer to this question lies in the specific character and incidence of the AM being considered. The project was commissioned on the basis that AM noise has received an increasing amount of attention more recently, mainly following the emergence of a number of reports of the existence of a certain type of AM whose character and/or incidence was not expected.

This is why it was felt necessary, at the introductory stage of the present project, to make clear the distinction between the long-accepted form of what is herein termed 'normal' AM (NAM), which is otherwise commonly referred to as 'blade swish', and any form of 'other' AM (OAM) which lies outside the range of what would be considered to be 'normal' AM.

It is the issue of OAM, as opposed to NAM, that the present project specifically seeks to address. This is because, as discussed in the relevant sections of this report, it is this OAM which has formed the focus of attention in recent years, with the presence of such OAM being cited in some cases as being a specific cause for complaint.

The 'problem' with the forgoing definitions of NAM and OAM is that the distinction between the two is not always so clear-cut as may be desired in the context of defining the scope of the present project. In this respect, two separate means of defining NAM and OAM may be considered: one in terms of source generation mechanisms and the other in terms of subjective character.

Defining NAM and OAM in terms of source generation mechanisms

Previous research into the issue of blade swish has resulted in a clear understanding of the fundamental source generation mechanisms involved in the generation of such noise and its resultant character and spatial distribution in the relative near-field of wind turbines (i.e. at distances of up to around 3 rotor diameters). Theory, backed up by detailed measurements, has concluded that blade swish noise results from noise generated at the trailing edge of the blades whose level and character varies quasi-periodically over time at any given listener location. This variation has been shown to be due to a combination of the specific directivity of the radiation of trailing edge noise relative to the blade geometry, coupled with the fact that the turbine blades are moving relative to the listener. This relative movement between source and listener results in variable amplitude effects due to a combination of the relative proximity of blade to the listener at any given moment in time and convective amplification, plus variable frequencies due to the Doppler effect[4].

Thus a potentially useful definition of NAM is that element of AM noise which can be fully explained by way of existing models of trailing edge noise that have been proven to be capable of successfully modelling 'blade swish' noise from wind turbines. The definition of OAM therefore becomes, by default, any form of AM noise whose physical characteristics (e.g. spectral content, amplitude, variability, directivity, etc.) and incidence (particularly at a distance from the turbines) cannot be explained by accepted source generation models for NAM. In making any such distinction between NAM and OAM on purely physical grounds it is important that no a priori assumptions are made concerning the relative subjective impact, or indeed acceptability, of the two different forms of AM.

Defining NAM and OAM in terms of perceived acoustic character

As an alternative to the foregoing distinction between NAM and OAM on the purely physical grounds of source generation mechanisms, the difference in the perceived acoustic character of the different forms of AM may also be considered. Such a distinction may be driven by noise complaints arising from specific characteristics of AM noise, whereby those affected have reported the noise to alter in character from the more usually encountered 'swishing' sound to that of a 'thumping' sound, with the latter sound being the cause for complaint. Some recent studies have also shown measurements, albeit limited, of the incidence of what could be characterised as OAM at a distance from some wind farms.

[4] The Doppler effect relates to a change in perceived pitch due to the effect of the movement of the source towards or away from a receiver, for instance the effect which is commonly experienced as noise from an emergency vehicle siren changes in pitch as it comes towards, and then moves away from, an observer.

Whilst any such difference in acoustic character may outwardly be an attractive means of differentiating NAM from OAM, the approach does raise its own issues. The main problem is that it relies on subjective judgement as to what constitutes NAM and what constitutes OAM. This subjective judgment can either be made by those involved in the research or on the basis of reported complaints. In the former scenario there must necessarily be some presumption on the part of those involved as to what may or may not be considered 'normal' (and possibly therefore also acceptable). In the latter case the successful discrimination between NAM and OAM relies on the reliable and consistent reporting of adverse noise effects by those living in the vicinity of wind farms. In this respect, any review is complicated by the difficulty in establishing the relevance of reported experiences and disturbances of wind farm neighbours from wind turbine noise, and in particular the relevance of the widely varying descriptions used by those reporting it and their potentially widely varying sensitivity to noise, especially when considered together with other potentially confounding, noise unrelated, issues. Even for some of the clearer reports (shown in different studies) of objective measurements of noise which could be characterised as OAM, the evidence was limited and often contradictory.

Choosing a definition of NAM and OAM for the present project

There is a further potential issue when considering either of the foregoing definitions as a means of differentiating between NAM and OAM. This issue relates to any change in character that may have occurred as a natural consequence of wind turbine development, and in particular the increase in physical size of turbines as they have increased in generating capacity from a norm of less than 500 kW two decades ago to typically greater than 2000 kW today. This general increase in physical size has included an increase in the diameter, chord and thickness of the rotor blades, coupled with the introduction of variable speed machines with lower rotational rates. The increased physical size of the blades would be consistent with a general reduction of the dominant frequencies radiated as a consequence of trailing edge noise generation mechanisms (as previously identified as being the main cause of NAM) from between around 500 Hz to 1000 Hz for earlier turbines to around 300 Hz to 800 Hz for modern, large scale turbines, although there is some conflicting data on this decrease[5]. The reduction in the rotational rate of the blades has resulted in a reduction in the periodicity of the AM from a maximum of around 1.5 modulations per second (i.e. a modulation frequency of 1.5 Hz) for the earlier machines, to less than 1 modulation per second for current machines (i.e. a modulation frequency of less than 1 Hz, with the maximum typically being closer to 0.5 to 0.8 Hz). Taken together, these two factors result in NAM for current turbines occurring at approximately half the rate of earlier turbines, with a dominant audible frequency content also potentially approximately half that of the earlier turbines. This being the case then, defining NAM based on source generation mechanisms, it is the case that what constituted 'normal' AM for earlier turbines may be different to what constitutes 'normal' AM for current turbines. Whether or not this change is sufficient enough to result in a subjectively significant change in response (and possibly also acceptability) is a question that the project seeks to resolve.

Notwithstanding the complicating factors just discussed, the project requires some consistency of definitions if discussions are not to become confused. Therefore, the definitions of NAM and OAM adopted are those based on the physical source generation mechanisms involved, with NAM being defined as that capable of being fully described in terms of 'standard' models of trailing edge noise and OAM being any form of AM lying outside this definition of NAM. It is stressed here, however, that in line with the foregoing discussion concerning this definition, no a priori assumption is being made as to the relative acceptability or otherwise of NAM as opposed to OAM: it is for the project to deliver conclusions in this respect.

In specific situations, tonal noise emissions have been found to vary in time (i.e. modulate at the blade passing frequency) as the mechanical source producing this tone varies in time, this can be readily assessed using existing methods to evaluate (non-stationary) tones [3]. This was therefore excluded from the scope of the current project.

Modulation frequency and audible frequency

The foregoing discussion has highlighted a potential area of confusion when interpreting the outcome findings of the project, which is the difference between audible frequencies and modulation frequencies. Audible frequencies are those frequencies which are capable of detection and interpretation via the human hearing mechanism. These are the frequencies at which the air pressure actually fluctuates. Audible frequencies are typically stated to range from around 20 Hz to around 20 kHz. As an example,

[5] See WPC data and Annex B to this report.

within this range the dominant speech frequencies typically lie between around 400 Hz and 4 kHz, with hearing sensitivity peaking at around 1 kHz such that the peak in hearing sensitivity matches the general peak in speech frequency. Aerodynamic noise generation from wind turbine blades produces a 'broad band' of frequencies across the human audible frequency range. Whilst this noise is all aerodynamic in origin, it arises from a number of quite different interaction mechanisms between the blade and the air, with some source mechanisms (and therefore also their associated noise frequency ranges) becoming more prominent than others depending on operational conditions.

The various noises arising from the different aerodynamic source mechanisms combine to produce the overall audible aerodynamic noise that may be heard when standing relatively close to a wind turbine. However, as previously discussed, one of the features of wind turbine aerodynamic noise is that it may increase and decrease in level with time. These increases and decreases in level occur at a rate equal to the blade passing frequency of the blades which, for modern large scale turbines, has been identified to be around 1 Hz or less. These 1 Hz variations in level can be subjectively quite discernible when present. However, this is not because the listener is hearing very low frequency sound of 1 Hz. Rather, the listener is detecting the regular variation in level, about once every second, of aerodynamic noise which is itself in the audible frequency range. Thus modulation frequency relates to the rate of the subjectively discernible low frequency modulation of noise which lies within in the normal audible frequency range: it does not relate to audible noise having a frequency of 1 Hz.

Bearing the above in mind, it is important when interpreting the outcome findings of the present project, that audible frequencies are not confused with modulation frequencies.

Project specification overview

The project specification was divided into a number of separate work packages. One of the key aims of project's execution was that, whilst each work package should produce its own reported outcomes, the various work packages should be interlinked in terms of information flow such that the ongoing findings of each package could inform the development of the other. With this aim in mind, a key element of the project was regular and open communications between the teams undertaking the various work packages.

The various work packages that went to make up the totality of the project are listed in the following Table 1.1.

An overview of the specification for each work package is presented in the following section 3 of this report.

The detailed outcome results of each of the work packages is presented in separate final reports, with a dedicated final report for each of work packages A1, A2, B1, B2, C and D. In terms of work package E, the wider dissemination of results, this has to date been delivered through presentations at different conferences[6], as well as through the open publication and peer-review of all project outcome reports.

WP	Description	LEAD
A1	Source generation effects modeling	NLR
A2	Fundamental Research into Possible Causes of Amplitude Modulation	ISVR
B1	Development of an Objective AM Measurement Methodology	ISVR
B2	Development of an AM Dose-Response Relationship	ARC
C	Collation and Analysis of Existing Acoustic Recordings	HLA
D	Measurement and Analysis of New Acoustic Recordings	HLA
E	Wider Dissemination of Results	ALL
F	Collation of Work Package Reports and Final Reporting	HLA

Table 1.1 – Work packages included in the present ReUK research project

[6] This included: Fourth International Meeting on Wind Turbine Noise, Rome, April 2011; Institute of Acoustics Conference, Cardiff, Jan 2012; Acoustics 2012, Nantes (France), April 2012;

The project commenced in March 2011, with an initially targeted duration of 8 months culminating in the issue of all associated final reports in November 2011. However, due to reasons largely associated with restricted access to the data necessary to undertake certain key areas of the work, the project programme slipped by 4 months, with a revised completion date for the major part of the project in March 2011.

Aim of the present overview report

The aim of the present report is to bring together the various outcome findings of the individual work packages, as set out in each of the final reports associated with each work package, into a single summary document. Clearly this document will be best read alongside the final reports from the different work packages. In particular those final reports provide much extended references and more detailed technical discussions. However, the present document has nevertheless been written to provide a stand alone summary of the project as a whole, as well as providing some additional considerations that may arise through the combination of outcome findings from the various separate work packages.

Project team

In order to best fulfil the project objectives, which are set subsequently in this report, the delivery of the project brought together a project team with strong combined expertise in both the aero-acoustic mechanisms of wind turbine noise and the environmental impact assessment and planning issues associated with wind farm noise. The experience of the project team was also selected to provide a balanced perspective both from those generally perceived to be involved with the wind energy industry and those who are likewise generally perceived to be wholly independent of the wind energy industry. The project team comprised members from the following organisations:

- Hoare Lea Acoustics (HLA) – Andrew Bullmore, Matthew Cand and others.
- National Aerospace Laboratory (NLR) - Stefan Oerlemans.
- Institute of Sound and Vibration Research, University of Southampton (ISVR) – Paul White, Malcom Smith.
- Acoustics Research Centre, University of Salford (ARC) – Sabine Von Hünerbein.
- Robert Davis Associates (RDA) – Robert Davis.

with the technical monitoring of the Project being undertaken on behalf of ReUK by:

- Dr Jeremy Bass, RES
- Mr Dick Bowdler, Independent Consultant

The perceived benefit of the assembled project team was that it included both those having direct experience of current AM issues and, equally importantly, those without such direct experience of AM but with a long established track record in more general acoustics research, including aero-acoustics. It was considered that, through this combination, the research could at the same time focus on the immediate issues of AM facing the wind energy industry in the UK, whilst also remaining open to ideas and scrutiny from a wider acoustics perspective.

One of the project team members generally lead the technical input in each project area. Hoare Lea Acoustics (HLA) additionally acted in a general technical and project management role. The various work packages, and the lead project team organisation responsible for their delivery, have already been presented in the final column of the preceding table of work packages.

The reports for the different work packages will be referenced throughout as WPX, where X is taken from Table 1.1. For example, WPC corresponds to Work package C.

It should be noted that WPA1 was commissioned by ReUK to NLR on a separate basis and has therefore formed the subject of a separately issued final report. The results of the WPA1 package of work have, however, been extensively referenced to inform the work undertaken as part of the various work packages reported herein.

Acknowledgments

The financial support of RenewableUK for this research project, as well as the input from the wind turbine noise research community on the subject, are both gratefully acknowledged.

2.0 BACKGROUND TO AM ISSUES

The issue of AM arising from the operation of wind turbines has recently been receiving an increasing focus of attention. Whilst the acceptability of audible noise from wind turbines continues to be the subject of considerable debate, the specific issue of AM has come to the fore following the publication of a number of studies claiming that the existence of such noise may result in an enhanced possibility of adverse impacts, in terms of increased annoyance.

The issue of AM is not a new one, having been the subject of a previous study undertaken by the University of Salford in 2007 [1]. That study was initiated following complaints of what was believed to be problematic levels of low frequency noise arising from a limited number of operational wind farms.

Following the publication of a number of press articles on the matter early in the 2000's, one of the key wind farm noise issues that began to attract attention was that of infrasound. During this time, reports of the potentially adverse impact of 'infrasound' from wind turbines were discussed in the press. In order to investigate the magnitude of this infrasound 'problem', the DTI commissioned a study into the subject which included measurement campaigns at three UK wind farms where infrasound and/or low frequency noise had specifically been raised by wind farm neighbours as being a problem. However, when the output of this study was reported [2], it concluded that infrasound from wind farms occurs at such low levels that it lies well below the thresholds of human perception. The overriding conclusion was, therefore, that infrasound poses no threat to health, nor could its presence at such low levels contribute to potential disturbance.

Notwithstanding the foregoing conclusion, the investigations did identify that the level of wind farm noise within the normal audible range was, at times, measured to increase above the threshold of audibility inside the premises at which noise measurements were undertaken. Where causal noise issues were positively identified, these tended to result not from low frequency noise itself, but rather from audible broad band aerodynamic noise that was being modulated in amplitude at a frequency related to the rotational rate of the wind turbine blades. Analysis of noise recordings made at the complainants' properties identified that, during the affected periods, the degree of 'swish' arising from the rotation of the blades was higher than that normally encountered, and perhaps also higher than the level originally assumed when setting wind farm noise limits in accordance with the document entitled 'The Assessment and Rating of Wind Farm Noise', ETSU-R-97 [3]. This document, which has generally become referred to simply as 'ETSU-R-97', currently provides the commonly accepted best practice for the setting of noise limits for wind farms across the UK, and has done so since its original publication in 1996.

At roughly the same time as the DTI study into low frequency noise, claims were emerging from other researchers that, particularly under certain atmospheric conditions more prevalent at night, wind turbines were capable of generating noise having characteristics outside of that expected of them [4]. The characteristic being referred to was an enhanced level of amplitude modulated aerodynamic noise. This type of AM resulted in the blade swish becoming more impulsive in character, such that those exposed to it tended to describe it as more of a 'whoomp' or a 'thump' than a 'swish'. This was similar to the character of the noise identified in the DTI study [2]. Appendix A to the final report of WPC presents a more detailed review of such information.

As a consequence of the above, adverse effects from wind farm noise that had hitherto perhaps erroneously been blamed on low frequency sound or infrasound itself, tended to refocus on AM.

A note of caution is advised here in that residual confusion is still often encountered as to what aspect of wind farm noise is actually being complained of when reference is made to 'low frequency sound'. This is because the term 'low frequency sound' is often erroneously used to refer to amplitude modulated aerodynamic sound. This is where the need for a very clear discrimination between low frequency sound itself, as opposed to the low frequency modulation of audible (higher frequency) aerodynamic sound, is required. With specific regard to low frequency sound, what is often meant in the context of wind farm noise is infrasonic sound in the generally sub-audible frequency range of less than 20 Hz. As far as such infrasound is concerned, repeated studies into noise and vibration from wind turbines, including the DTI study referred to above, have confirmed the lack of sufficient energy in these very low frequency bands to result in either direct adverse health effects or subjectively perceptible effects. In contrast, the same or similar studies confirmed the potential presence of audible wind turbine aerodynamic sound that is amplitude modulated at a low frequency of typically around 1 Hz.

It was the confirmed presence of this type of AM that provided a possible causal link between a physically measurable and subjectively perceptible acoustic effect and reported adverse responses. A consequent resurgence in interest in AM has therefore arisen during the past five years or so[7]. An issue that is of particular concern to potential neighbours to wind farms, and the wind energy industry alike, is the question as to how commonplace the occurrence of 'problematic' levels of AM is. In 2006, the DTI study [2] presented the results of noise measurements undertaken at 3 wind farms. This was 3 out of a total of 5 wind farms where disturbance due to 'low frequency noise' had at the time been formally reported. The alleged problem sites therefore represented less than 4% of the total number of operational wind farms in the UK at that time. At one of the tested sites the problem had already been alleviated by the introduction of a wind speed and sector noise management reduction system, which had to be switched off in order to acquire the data required for the report. At each of the other two sites, noise monitoring equipment had to be installed for a minimum of between 4 and 8 weeks for occurrences of the problem situations to be self-reported by the occupants of the dwellings concerned. Based on these facts alone it appeared, at least at that time, that the problem could not be regarded as commonplace across UK wind farms, either geographically or temporally.

The finding of the DTI study was further corroborated by the subsequent study undertaken by the University of Salford [1]. The evidence available at the time of the Salford Study concluded that the issue of AM was of limited extent across UK wind farms. Anecdotal evidence relating to wind farms in other countries reinforced this initial impression of limited impact. Nonetheless, AM was formally identified through the University of Salford study to occur at a limited number of wind farms.

Since the time of the Salford Study [1], there have been claims of increased occurrences of 'problematic' AM, with one suggestion being that the adoption of physically larger turbines invalidates the assumptions made in ETSU-R-97 concerning the 'acceptable' AM that may be expected to be produced through the normal operation of wind turbines. In this respect the 'norm' at the time of ETSU-R-97 being developed was based largely on turbines having typical hub heights and rotor diameters both of less than 40m. Current large scale wind turbines are now based around hub heights and rotor diameters of typically double this size at 80m or more. At the same time typical generating capacities have increased from less than around 500kW to 2,000kW or more.

Whilst recent research [5,6] undertaken prior to the present project has demonstrated that some degree of AM is inherent to the operation of all wind turbines, whether large or small in size, the features of the AM identified as a result of that research cannot explain the characteristics that have been reported at some sites (see, for example, the final report for WPC). Such characteristics include situations in which AM is distinctly audible in the far-field of turbines, predominantly in the downwind direction, is more impulsive in nature, and is observed to occur intermittently.

Whatever the root cause of this 'other' form of AM, the consequence of its occurrence on some wind farms has resulted in concerns increasingly being raised both by local authorities and the public alike as to the likelihood of it occurring around newly proposed wind farms. The argument has increasingly been presented that, if the current state of knowledge precludes the accurate prediction as to the form with which AM will occur at a proposed wind farm, then a precautionary condition should be set in place. The problem with imposing an AM based planning condition on such a precautionary basis is that the robustness and effectiveness of any such conditions would, within the current state of knowledge, be totally untried and untested. It would also necessarily be based on an incomplete knowledge of how any objective measure of the AM may relate to subjective response.

It is with the foregoing issues in mind that ReUK, on behalf of the UK onshore wind energy industry, commissioned the present research project. The main aims of the project are to address two key issues:

- First, the causal mechanisms for any AM that falls outside that which is known to result from a wind turbine's normal operation need to be better understood. Only in this manner can its occurrence be predicted (and therefore precluded) at the design stage of a wind farm development. Achieving this goal of understanding the causal mechanisms would mean that certainty could be placed on proposed wind farm developments with regard to the occurrence of AM issues. This being the case, then the imposition of a condition against such an occurrence would become unnecessary.

[7] It is noted that, even before this, the possible existence of 'enhanced' or 'excess' levels of AM (what constitutes 'normal' and 'enhanced' or 'excess' levels of AM remains the subject of some debate, and is an issue that the present project considers in some detail), the issue of AM had not been ignored. Indeed, even in the very early days of wind turbine development, the adoption of downwind configuration horizontal axis wind turbines, particularly in the USA, had resulted in reports of high levels of 'blade thump' noise even at large distances from the turbines

- Second, for those situations where it cannot be guaranteed that AM issues won't occur (or in the event that the outcome of the foregoing investigation concludes that such an assurance is simply not possible), then a robust objective metric for the rating of AM effects is required to be developed for inclusion in any AM based condition. This metric should necessarily relate directly to the subjective impact of AM where it occurs, such that any AM condition based around the metric protects those exposed to it, whilst equally being wholly objective and repeatable in its derivation such that it provides certainty from the perspective of the wind farm operator.

3.0 OUTLINE DESCRIPTION OF WORK PACKAGES

The following descriptions of the various key work packages are based on the scope of coverage of each. In the main, these descriptions reflect the proposed content of the separate work packages at the contract definition stage. However, in some cases the detail of the approaches has varied from that originally envisaged as a consequence of the findings of one work package informing the others as the project progressed.

It is acknowledged that there may be some repetition of basic considerations in this section, as some of these have necessarily already been covered by way of general introduction in the preceding sections. However, it is felt useful to also incorporate those considerations herein in order to aid the understanding of the key issues facing the teams undertaking each work package.

Work Package A – Fundamental Research into the Causes of OAM

At the time of project definition, it was speculated that the primary cause of OAM at blade-passing frequency was non-uniform air flow into the rotor plane resulting in cyclic changes in blade-loading and resulting in variable noise generation. The non-uniformity may result from wind shear, inflow turbulence (from topographic forcing or the wake flow from other turbines), or yaw error (since the turbine control system can only control the rotor orientation by reference to the mean hub-height wind direction). AM may also be enhanced in propagation as a result of wind and temperature gradients. The relative importance of these contributory factors would then vary between sites and (probably) between turbine designs.

It was also recognised that there exists a large body of knowledge on aerodynamic noise from of all types of rotors, including wind turbines, helicopter blades, turbo-machinery in aero-engines, etc, and the catalogue of noise sources mechanisms are well known from the classical literature (see WPA2), including noise due to flow past the trailing edge of the aerofoil, the tip vortex, flow separations and inflow turbulence. Each source has particular frequency and spatial characteristics. Most relevant current research is led by the needs of the aircraft industry, involving work on high-speed fans and turbines and, increasingly, on noise generated by airframes (appendages such as flaps and undercarriages), which is now a significant noise source for commercial aircraft on approach.

The key to delivering an improved level of understanding of the key drivers for OAM was identified as being the ability to establish a clear correlation between the characteristics of noise (i.e. spectral, temporal and spatial features) on a few typical installations compared with the expected characteristics of the potential mechanisms. Oerlemans [5,6] has shown that some degree of NAM is inevitable, even under ideal environmental conditions, because of the inherent directivity of noise sources on a moving aerofoil. However, this does not explain the range of the feature reported in practice.

Therefore, in a first instance, the previously developed model of NAM [5,6] was developed and extended to cover more general inflow conditions and source generation mechanisms, in an attempt to replicate the observations of OAM: this was done by S. Oerlemans in WPA1.

WPA2 then continued this theoretical work from a more general point of view. Initially the research under WPA2 was based on an analysis of the existing data which is available immediately under WPC, thus limiting lead times for this element of the work. However, an important output from WPA2 was to inform the gathering of additional corroborative data of a specific and targeted nature to be carried out under WPD.

The scope of Work Package A2 comprised:

1. carry out a review of published literature relating to AM aerodynamic noise effects in rotating machinery;
2. examine existing wind turbine noise measurement data available from WPC;
3. produce a number of hypotheses for the origins of AM derived from the literature and current measurement data review and from 'brain-storming' meetings involving key staff from the project

team and representatives from ReUK, as well as contributions from other researchers in the field and from wind turbine manufacturers;

4. determine what the likely spectral, temporal and spatial features of each of the candidate mechanisms are likely to be, and compare this with existing data and use it to inform the data gathering undertaken as part of WPD;
5. use all available data to identify the key drivers for AM and hence establish potential causal mechanisms;
6. based on the identification of the key AM drivers, produce a list of AM ‘risk factors’ to assist developers in predicting the likelihood of varying AM occurring at any particular site with a particular turbine type and configuration (due to the necessarily limited scope of WPA2 it was recognised at the project definition stage that this list was highly unlikely to provide a means of calculating AM risk with any degree of precision, rather it should provide significantly better informed guidance than is currently available).

Work Package B – Development of Objective Amplitude Modulation Measurement Methodology & Development of a Dose-Response Investigation

There are three interlinked issues that this element of the project was required to address:

- the desire to develop an objective methodology which would allow the automated discrimination of wind farm amplitude modulated noise, in much the same way as there is an accepted methodology for the automated discrimination of tonal noise;
- the desire to enable the discrimination method to produce an objective rating of the identified amplitude modulation which could be related to a subjective ‘dose response’ relationship;
- the need to establish a ‘dose response’ relationship based on subjective testing.

It was clear that there would need to be some cross-fertilisation of ideas between the various elements of the work package, as each of the three foregoing issues relies to some degree on the other two.

Fundamental to developing a dose response relationship is the requirement to develop an objective metric which represents the characteristics of the stimulus (the amplitude-modulated noise) and weights these characteristics to generate (ideally) a single number value that can be shown to correlate with subjective response. However, the potential complexity of achieving the foregoing aim was recognised form the outset. For instance, signal-to-noise ratio was acknowledged to be a potential factor. In the presence of background noise levels that lie close to the mean level of a modulated superimposed sound, the superimposed sound can be partially masked and the modulation depth (peak-trough) reduced, depending on the frequency content of both components. Additionally the acoustic parameters will be different indoors and outdoors and whether windows are open or closed. In each case the background noise and spectrum will be different, with mid frequencies potentially being relatively more important outside and low frequencies being more important inside a double-glazed house, for example.

Another potentially complicating factor is that annoyance is also affected by non-acoustic factors. Such factors include an individual’s attitude to wind turbines in particular or to renewable energy in general. Annoyance will also be related to disturbance of concentration if an individual is undertaking various tasks, such as working at home, reading, listening to music or radio or television, trying to rest or sleep, working in the garden, or relaxing on a patio. Once a task has been disrupted, annoyance may increase rapidly. As such there is likely to be a relatively sharp threshold level between not annoying and very annoying. This threshold will vary from individual to individual, and from occasion to occasion for each individual according to the tasks being undertaken degree of concentration, number of times the noise is noticed, *etc*. The threshold may also depend on previous experience, as some individuals may become sensitised to a particular noise once they notice and recognise it, whereas others may become habituated. Thus, in designing the experiment, the questions posed to the subjects needed to be carefully considered, and the limitations and value of tests undertaken in a laboratory environment needed to be recognised.

In practice, the work undertaken under work package B was perhaps the element of the project which deviated most from the initially planned approach. In particular, the development of objective discrimination tools and the validation of the use of these various methods proved to be significantly more involved than had originally been envisaged. This was not least due to the fundamental issue of defining what actually constitutes a reasonable measure of the degree of modulation present in a signal, the variable ranges of objectively quantified values that can result from seemingly small changes in assumptions concerning the underlying models, and how these values may then be related to subjective response. In order to deliver the above, the work package was divided into two separate components: B1 considered the development of an automated AM identification and rating tool based on signal processing techniques, whilst B2 considered the subjective response to AM.

B1 - Development of an AM Descriptor ('metric')

The requirement was essentially to develop a signal processing method capable of robustly identifying and quantifying the presence of the modulation of the sound envelope of the noise signal. The basic starting premise for identifying potential identification methods was that modulation occurs in a generally periodic manner at regular intervals related to the blade passing frequency of the turbine. However, it was recognised that, in practice, the signals would be 'quasi-periodic' (i.e. the modulation would not be truly regularly spaced in time) due to the variable speed nature of turbines. Multiple modulation rates may also feature in any noise signal due to turbines across a wind farm site having different rotational speeds at any given moment in time. Extraneous sources could also cause the sound to also vary in level, quite independent from the presence of any wind turbine noise. Quite apart from these physical considerations, it was also recognised that any objective measurement methodology derived from the study may need to be modified to take account of the findings of the subjective testing being undertaken as part of WPB2, whose aim was to determine what parameters are important in controlling the dose-response relationship. However, in the first instance, the metric was expected to take account of such factors as:

- level and frequency content of the modulated sound;
- depth of modulation;
- waveform of modulation (modulation rate, sinusoidal, strongly impulsive etc);
- temporal variation in depth and waveform of modulation.

Notwithstanding the above, the final form of metric could not be developed in isolation: a key criterion for the appropriate metric was that it must produce repeatable results and show good correlation with subjective response. The necessary approach was therefore required to:

- develop a candidate metric (or metrics) which would ideally account for all variables that are likely to be subjectively detectable;
- design subjective tests in such a way that the range of stimuli presented to subjects represents all variables included in the metric(s), and the range of stimuli presented should correspond with the range observed on wind farm sites (based on supporting input from WPC and WPD);
- refine and develop the metric if needed to optimise correlation with subjective response.

It was also considered important to evaluate how these metrics could be applied to realistic field signals.

B2 - Listening Tests

It is known that level alone is only one piece of the equation that goes to make up an individual's total annoyance response to noise. Other factors (some non-acoustic) often lead to the exposed individual becoming increasingly sensitised to the noise. Alternatively, through prolonged exposure, individuals may become habituated to a given noise, or they may develop coping strategies (either consciously or subconsciously) which result in a lessening of the impact. Two of the key factors other than level that affect an individual's response to noise include spectral content and frequency of occurrence (when, how often, and for how long). It was therefore clear that laboratory based tests under controlled conditions could result only in relative judgements as to the perceived differences between different noises. Such short term tests could not produce an absolute measure of annoyance. The latter could only be established via a large scale social survey of individuals in their home environments, but then with the associated uncertainties concerning the actual exposure of the test subjects to the specific noise of interest.

It is plausible that OAM noise may be more disturbing than both steady and NAM wind farm noise. This possibility is driving an ongoing debate as to whether current guidelines for the assessment of wind turbine noise, such as ETSU-R-97, are sufficiently taking into account the possible occurrence of OAM. As a consequence, wind farm developers, planners and policy makers are interested in finding out how OAM noise is generally perceived and how listeners respond to it.

Unfortunately, not much is known about the occurrence of and response to OAM. The work undertaken under this work package was aimed at developing a scientifically based procedure for the rating of OAM effects. Previous work has suggested that the following general behaviour might be observed:

- a threshold of onset of annoyance will occur when the fluctuations become sufficiently pronounced that the modulating sound becomes more annoying than the steady sound of the same sound pressure level;
- for AM values above the threshold, there will exist a fixed relation between AM characteristic parameters and a mean annoyance score, with previous research suggesting that annoyance might be observed to systematically increase with the increased prevalence of certain characteristic parameters.

It was therefore recognised that, if one or both of the above hypotheses could be validated for OAM noise, then their quantification could aid the development of guidance for any OAM based planning condition. For example, a relation between mean annoyance scores and fluctuation strength (as defined in [7]) could potentially be used to define an assessment scheme for OAM by matching the mean annoyance scores of OAM noise with continuous noise. This would provide the basis for rating the OAM characteristic of wind turbine noise, if considered necessary.

The main objectives of the listening tests were therefore two-fold:

1. to validate different AM metrics as a measure of AM which correlates with subjective response;
2. to investigate the relationship between AM value and mean annoyance score, with the following aims:
    - to establish (in terms of AM 'value') the threshold of onset of annoyance (i.e. when modulated sound becomes more annoying than steady sound of the same mean level);
    - to establish the relationship between AM value and mean annoyance score at AM values above the threshold by matching the mean annoyance scores of AM noise with continuous noise.

The key stages of the development and implementation of the listening tests were as set out below.

1. key decisions had to be made as to the scenarios to be tested (for example indoors or outdoors, fixed AM or variable AM, range of AM depths, range of overall levels, key variable parameters, etc.);
2. test signals were generated using simulated data to represent the range of characteristics of AM noise generated by wind turbines based on input from WPC and WPD (the use of simulated data as opposed to recorded wind turbine noise was required to permit the necessary control over the various parameters whose effects were required to be tested);
3. test signals (including both background and wind turbine sounds) were selected to be representative of those experienced at typical nearest residential neighbours to wind farms in rural situations outdoors (i.e. below 45 dB(A));
4. pilot tests were carried out to validate and develop the experimental method (this involved presenting test sounds in a quiet listening room and asking subjects to record an annoyance score on the basis that they were relaxing outside their home, plus the alternative approach of requesting subjects to adjust the levels of the test samples on the basis of establishing 'equal annoyance' with certain standard signals);
5. samples were presented for durations of about 20-30 seconds in a random sequence to include background noise and to test 'wind turbine' noise at different levels and with varying AM characteristics, as achieved through varying the test parameters;
6. the pilot tests were undertaken involving a limited number of subjects, the results of which were to narrow down the range of parameters to be varied and tested for the final listening tests;

7. final listening tests were carried out with an increased number of subjects in more controlled conditions;
8. the results of the final listening tests were analysed using the various objective metrics developed under WPB(1) to establish a dose response relationship.

Work Package C – Collation and Analysis of Existing Acoustic Recordings

Third party sources that may hold relevant acoustic data relating to AM were identified. These parties were then formally approached and access to their data requested on the basis of its confidential use. All parties thus identified and approached were added to a database developed as part of this Work Package C. Where available, stored data included relevant supporting information including, amongst other factors, site topography, turbine details (including operational data where available), reports of AM at the site, details of the noise measurements, etc.

It was highlighted as a potential risk to the project that the exercise of seeking data from third parties could result in no (or very limited) positive data returns, and that this would throw into some doubt the ability of WPC to deliver even a 'broad' estimate of the frequency and severity of the AM problem across UK wind farms. In an attempt to circumvent the foregoing issue, it was also requested that the identified data owners could alternatively release summary details of their own analysis of the audio data they hold, even if they are not prepared to release the audio data itself, on a site anonymous basis. It was also identified that confidentiality issues may also limit the returns available from even this reduced request. In reality, requests made under WPC were more successful than had originally been anticipated at the project tender stage, although requests were understandably received from a number of wind farm operators to keep the information confidential. This requirement of anonymity has not, however, taken away from the utility of the data in informing the project team and in contributing to the final outcomes of the project as a whole.

One of the complicating factors of the present study is that wind turbine noise generally only becomes an issue where it adversely affects residential neighbours to wind farms. The issue as to whether or not the operation of a particular wind farm may cause problematic noise therefore generally relies on the reliable and consistent reporting of its effects by those living in the vicinity of the wind farm. The review and interpretation of available information is complicated by the difficulty in establishing the relevance of reported experiences of wind farm neighbours, and in particular the relevance of the widely varying descriptions used by those reporting being affected by wind farm noise. The potential issue of relating subjective descriptions to actual effects applies particularly to the specific subject of Amplitude Modulation noise from wind farms, especially when it is appreciated that this potential feature of the noise may vary and/or take different forms, each of which may be self reported using different descriptors across different subjects. All the foregoing factors were recognised in terms of risk to the outcome conclusions resulting from this work package which relied on existing data allied with reports of AM.

In terms of formal deliverables for Work Package C, these comprised a compilation of audio recordings suitable for the testing of candidate AM assessment methodologies under WPB. It was also initially planned that this would allow an informed estimate (based on the data returns) of the extent of the AM problem across UK wind farms, at least as far as is practicable within the available budget, timescale and confidentiality constraints of the available information, but this did not in fact materialise in practice. In producing the foregoing outcome results from the available data, possible factors positively contributing to the occurrence of increased levels of AM were also to be considered.

Work Package D – Measurement and Analysis of New Acoustic Recordings

At the project outset it was initially conceived that additional measurements would be undertaken as part of Work Package D at up to seven separate sites, the original aim being to collect additional recordings to those of work package C, but using an essentially similar approach which would focus on recordings at the locations of neighbouring dwellings, or proxies thereof. On this basis, the proposed work programme for Work Package D comprised the following stages.

1. Based largely on the outcome of Work Package C, sites were to be identified at which AM is known or is alleged to occur to such a degree that it gives rise to adverse comments or formal complaints.

2. A project programme was to be prepared for undertaking the proposed noise measurements, along with details as to how supporting meteorological and operational data would be acquired. This programme would contain details of the proposed sites at which measurements would be undertaken, the locations at which it was proposed to undertake the measurements and the measurements to be undertaken (e.g. attended short term / unattended long term, internal / façade / external measurement locations, monaural / binaural recordings, parabolic / array microphone systems, mast based / remote sensing meteorological parameter recordings, wind farm operational data requirements, etc.).
3. Measurement sites were to be secured and noise monitoring / met monitoring systems were to be deployed at these locations.
4. The operator's of the selected wind farm sites were to be approached in order to acquire relevant wind farm operational data for the duration of the noise monitoring exercise. A running analysis of any data acquired during the course of the measurement exercise at each site was to be maintained in order to inform the utility of further measurements and to amend the data acquisition programme, where deemed necessary, to maximise the chances of a successful project outcome.
5. A detailed objective AM analysis of the data acquired under Work Package D, where possible using the analysis tools developed under WPB1, was to be undertaken in order to correlate the measurement results against potential controlling factors.
6. The outcome results of the preceding analysis were to be used to refine, where found necessary, the analysis techniques used for the objective assessment of AM and to feed the results into WPB1 and WPB2, and also to provide recommendations as to the preferred measurement technique for the identification of AM effects.
7. The database developed under Work Package C was to be supplemented with all relevant information and audio records resulting from Work Package D.

Prior to implementing the initially envisaged measurement campaign under WPD, a review of the current knowledge and experience of AM was undertaken to provide a starting point for the investigation. The review sought to include relevant reports and potential data sources from projects from both within and outside the UK. Based on the information gathered, coupled with the experience of the project team, it was recognised that, even at those wind farm sites where AM has been reported to be an issue, its occurrence may be relatively infrequent. Also, for the reasons set out above under work package C, reliance on reports of AM could be open to interpretation.

It was therefore acknowledged that the capture of time periods when subjectively significant AM occurs may involve elapsed periods of several weeks or even months, even at those sites where it was accurately reported as occurring. As a consequence, the scope and methodology of WPD evolved significantly with time. It was determined that, rather than simply acquire more examples of AM audio recordings in the far-field, more detailed supporting information, such as better defined meteorological information and turbine operational information was crucial to obtaining an improved understanding of the potential causes of OAM. The approach adopted was therefore to focus on a more limited number of sites but to undertake more extensive and detailed measurements than originally envisaged at each of these sites using a variety of strategies, the requirements for data capture being guided by the findings of other work packages.

One of the benefits of the originally envisaged approach to the measurements to be undertaken under WPD was that recordings could be taken at any accessible location without necessarily the need for approval from any involved parties. However, the downside to this approach was quickly recognised to be the subsequent inability to access relevant supporting information from wind farm operators who were understandably cautious about allowing their wind farms to be used as part of the project.

The revised approach therefore required the direct involvement of wind farm operators if the necessary supporting information was to be made available. The benefits of the revised approach were determined to far outweigh any drawbacks, although it had to be accepted that one of the drawbacks was the significantly extended time required to arrange site access and to adequately address confidentiality considerations. Also, many logistical and technical challenges were associated with the more detailed and innovative nature of some of the measurements proposed, some of which required working in close collaboration with wind turbine manufacturers in addition to the wind farm operators. This led to some delays to the overall project programme, but the value of the additional measurements which were undertaken were deemed to more than outweigh any issues associated with these delays.

It was also therefore not expected that the measurements undertaken as part of this work package would offer more evidence into the frequency of occurrence across the country of the phenomenon, but it was considered preferable to focus on acquiring further insight on the characteristics of 'other AM' and its potential causes.

Ultimately, a staged approach was taken in which different approaches were employed at three different sites in order to achieve separate aims, as summarised below.

Site A involved measurements that were most similar to wind turbine noise immission (far-field) measurements, as generally undertaken following the guidance of ETSU-R-97. However, in accordance with the expected requirements for the measurement of any modulation, high-resolution data including audio recordings were also collected as part of these works. Recordings were made at two residential properties neighbouring a wind farm site for which residents at the measured properties had been complaining about the noise from the wind farm, with investigations currently underway (by others) to assess the complaints further. Access was granted to the two properties which enabled detailed measurements to be made at suitable external locations over a period of several weeks.

Site B comprised a location where the existence of OAM had been positively identified. Efforts were therefore focused on undertaking a detailed series of measurements, allowing a study of conditions in which varying levels of AM were experienced and at different locations relative to the turbines. This also involved capturing detailed anemometry measurements and turbine operational data at a high resolution. The schedule and requirements of the measurement campaign were based on the theoretical considerations and requirements outlined in the WPA2 report, in as much as practical and budgetary conditions reasonably allowed. These measurements were significantly more detailed than those generally undertaken for immission measurements of wind turbine noise but were nevertheless still limited to some degree through practical constraints.

Site C involved targeted measurements on an operational wind farm site using novel techniques to investigate the influence of turbine operational parameters on the character of the AM noise produced, as informed by the results of the other work packages in this project. In particular, the hypothesis raised in WPA1 and WPA2 that partial blade stall may lead to increased levels of modulation at large distance downwind warranted particular investigation. Considering this, the influence of a turbine's pitch regulation system was considered a crucial element to be tested. Designing a measurement campaign capable of addressing the foregoing issues required detailed cooperation with turbine manufacturers, especially as a detailed knowledge of the blade geometry and control system operation of commercial wind turbines is generally not freely available due to confidentiality considerations. It was agreed with the turbine manufacturer to undertake tests in which the pitch of the turbine was controlled directly in order to attempt to trigger detached (stalled) flow on the blade, and to assess the relative impact on far-field noise, and in particular the potential for varying pitch angles to induce or control amplitude modulation. The project team is not aware of such investigations having been undertaken previously.

4.0 SUMMARY OF FINDINGS

As a consequence of how knowledge concerning AM has developed since the present project was conceived, it is useful to review what the important outstanding issues were then, and are now. To this end a summary comparison of the issues, and how they have developed since the start of the present project, is presented in Appendix A.

The present section seeks to summarise answers to the key questions concerning AM out of the findings of the various work packages. In doing so the aim is to highlight those areas where firm conclusions may be drawn, and those areas where information is still lacking. In this respect, whilst the outcome findings of the present project are undoubtedly positive in that they have moved the understanding of AM forward in several major areas, it will become apparent on reading the totality of the project reports that some major issues still await clearer explanation. Furthermore, the explanations for some issues that were previously thought would be relatively easily established have, in some cases, become less certain as a consequence of observations made under this project. These are areas where further work is recommended, as outlined in the conclusion.

At the time of the commencement of this project, a number of theories were being forwarded by various acoustics practitioners active in the field of wind turbine noise as to the likely causes of OAM. These included, amongst other factors:

- blade passing tower
- angle of attack changes
- high wind shear conditions
- stable atmospheric conditions (also leading to high wind shears)
- low-level jets
- high turbulence
- yaw error
- rotor/wake interaction effects
- interaction between turbines, particularly when arranged in a linear array;
- synchronisation or relative phasing between turbine rotations;
- propagation effects;
- 'stubby' towers[8].

It was one of the starting aims of the project to look at each of these theories and assess their merits. However, it was equally recognised at the outset of the project that to focus solely on such theories, often based purely on anecdotal evidence, could lead the investigation to miss other, truly significant, factors. The starting point of the project was, therefore, a diverse project team meeting. This project team inception meeting included invited contributions from other experts outside the immediate project team (including Professor Richard Sandberg, an aerodynamicist based at the University of Southampton) to look at all the evidence available to the project team at that time. This included listening to samples of AM which included both NAM and OAM (see the discussion in the introductory Section 1 for a definition of these terms) and consideration of circumstances under which each had been reported, as far as such information was available.

It was quickly established that the available supporting information was often contradictory. One of the outcome conclusions of the very first meetings of the project team was not that more examples of AM audio recordings were necessarily needed per se. Rather, what was required were AM audio recordings with more detailed supporting information, such as better defined meteorological information and turbine operational information. As the project progressed, efforts were made to collect such supported audio recordings from third parties under WPC, with some success. However, as a consequence of the identification of the various potential causal mechanisms of AM in its different forms, it was recognised by the project team that the validity of a number of these mechanisms could not be either proved or disproved on the basis of the available data. It was additionally recognised that it was highly unlikely the theories could be proved or disproved by the initial remit for WPD of the project which anticipated the

[8] Turbines with a large rotor diameter relative to heir hub height, although the ratio is not described precisely.

collection of recordings of wind farm noise at typical residential dwelling locations around a number of different UK wind farms, but again without any detailed supporting information.

It was conceived, however, that some relatively simple measurements could be devised to assist in identifying the validity of at least some of the postulated theories. The problem facing the project team was that these measurements, although relatively straightforward in concept, would require the close cooperation of a wind farm operator and wind turbine manufacturer, as they would necessitate access both to the full range of turbine control system parameters as well as the ability to manually vary the control of a turbine outside its normal operational envelope. Such tests were ultimately arranged and undertaken under Work Package D once wind farm operators and wind turbine suppliers had been identified and had offered to work in full cooperation with the project team. However, through no fault of the various parties involved, whose proactive contributions to the project are gratefully acknowledged, getting to this stage was a lengthy process and one which contributed significantly to the delay to the project.

The particular piece of information identified as being ideally required was that relating to the precise flow conditions on the surface of the turbine blades themselves. It was therefore unfortunate that the collection of such data was outside the scope of the present project and it was not possible to obtain such information from any third parties. However, interesting discussions were held as part of the present project with third party research organisations in Denmark. As a consequence of those discussions it remains possible that, in terms of future work, access could be made available to test turbines fitted with the necessary instrumentation to obtain this data on the back of unrelated research that is being carried out by those organisations.

Due to the manner in which the research undertaken under the project's various work packages was interrelated, it is considered unhelpful here to summaries the findings of the project on a work package by work package basis. Instead, the present section is divided into a number of 'headline' issues that each call on the results of more than one work package to fully explain the project's outcomes concerning each issue. These issues comprise:

- What is AM?
- What causes AM?
- Can AM be objectively identified and quantified?
- How does the theory for AM identification apply to real world data?
- What is the subjective impact of AM?
- What are the characteristics of OAM?
- Can OAM be predicted?
- Can OAM be mitigated?

Prior to attempting to answer the foregoing questions, it is first useful to set the scene as to the available anecdotal evidence that has been collated during the course of the project, and to see how this information may immediately impact on some of the previously postulated hypotheses for OAM. It should be remembered here that the definition of OAM adopted for the purpose of this project relates to any occurrences of AM observed to lie outside the characteristics of that which may generally be classified as 'normal' blade swish.

- the modulation depth (the difference between the levels of adjacent peaks and troughs) can be significantly greater than that of normal blade swish, with changes in A-weighted levels of up to 10 dB having been measured;
- the dominant frequency characteristics are lower than for normal blade swish, with a shift in the dominant frequency range to typically around the 300-400 Hz region;
- the effect has only been reported on a limited number of wind farm installations;
- the effect is not restricted to an individual turbine type;
- the effect has been reported both for individual turbines and for wind farm arrays;

- a particular turbine type that exhibits the effect on one wind farm site will not necessarily exhibit the effect on another site;
- even on those sites where the effect has been positively identified to occur, it is intermittent;
- the effect may occur for just a few rotations of the blade, or it may persist for periods of several minutes or hours;
- the effect has been reported to be most common during evening and night time periods, although it has been observed at times during day-time periods;
- the effect is more dominant in the far field (typically 10 rotor diameters or more from the turbine) and may not be simultaneously discernible in the near field (typically less than 3 rotor diameters) of the turbines;
- the effect is generally strongest in the downwind direction and can also be present in the upwind direction of turbines, but it has not been clearly recorded in the cross wind direction in which normal blade swish is most prevalent;
- the effect has been reported on wind turbine installations on both flat and hilly terrain;
- the effect has frequently been associated with conditions when wind shear may be expected to be high, but has also been reported to occur on some hilly sites or during damp conditions of light or even heavy rain when wind shear may be expected to be low.

The foregoing observations, whilst only anecdotal, immediately lead to the conclusion that if any one factor or combination of factors were responsible then the effect would occur at all sites featuring those factors and it would occur frequently at those sites. However, neither of these situations has been observed to be the case in practice. As an example, the effect has been observed to occur on a single isolated turbine with a high tower, so interaction effects between linear arrays and too closely spaced turbines cannot be a decisive factor, neither can a large rotor on a short tower. As another example, the effect has been reported for one make and physical configuration of turbine on one site, whereas it has not been reported for that same make and physical configuration of turbine on all installations. Such contradictory observations serve to illustrate the problems that have been facing the wind energy industry, and reinforce the reasons why the resolution of the AM issue has been proving such a non-trivial task. It is the aim of the present project, as summarised over the remainder of this report, to draw together the various observations and theories in an attempt to provide answers to some of the questions raised.

The starting point of the discussion should naturally be, as proposed previously, to better define what is meant when we refer to ‘AM’.

4.1 What is AM?

Descriptions and definitions of AM, NAM and OAM have already been presented in some detail in the introduction to this report (Section 1). This was deemed necessary as these definitions are so important to placing all subsequent discussions into context and to avoid subsequent confusion over terminology. In order to avoid unnecessary repetition, it is assumed that the reader of the present section will have read the introduction, such that the basic terminology and issues under consideration are appreciated. The present section of the reports therefore expands on the discussion contained in the introduction section only where it is considered helpful to do so.

Given that the definitions of AM and also (at least for the purpose of the present project) NAM and OAM have already been provided as being various incarnations of the regular variation in level of wind turbine ‘aerodynamic noise’, it is perhaps useful to present a brief overview of wind turbine aerodynamic noise in general.

As discussed in greater detail in WPA1 and WPA2, all wind turbines generate aerodynamic noise when their blades rotate. This aerodynamic noise is caused by the interaction of the blades with turbulence in the air flow. Some turbulence is present in the inflowing wind itself. This turbulence causes so called ‘inflow turbulence’ noise. However, turbulence is also generated by the boundary layer of the flow over the blades, and this is the origin of a number of ‘self’ noise mechanisms. Self noise would be produced by a turbine even in a uniform and non-turbulent flow in the absence of any inflow turbulence as it is produced by the displacement of the air as the blade moves through it.

There exist five key mechanisms for the generation of aerodynamic self noise from blades in general:

- boundary-layer turbulence passing the trailing edge - this is the dominant aerodynamic noise source on wind turbines under normal operating conditions;
- separated-boundary layer / stalled-aerofoil flow - this is a potentially major source in particular conditions, but it is not expected to be significant for wind turbines during normal operating conditions;
- vortex shedding due to laminar-boundary-layer instabilities - this is unlikely to contribute to wind turbine noise as the flow regime does not apply;
- vortex shedding from the blunt trailing edge of the blade - this is a known feature of wind turbines, but generally occurs at high frequencies;
- the turbulent vortex flow existing near the tips of lifting blades - this is normally a relatively high frequency problem and has also been largely controlled by careful tip design on modern blades.

The dominant self noise mechanism when dealing with the A-weighted spectrum of wind turbine noise is generally considered to be trailing edge noise. As indicated in the first of the preceding bullet points, trailing edge noise occurs when the turbulent boundary layer of the flow over the blade is convected past the sharp trailing edge. Due to the relatively higher speed of the tip of a blade as it passes through the air; trailing edge noise is more prevalent towards the outermost part of the blade lying closer to the tip than the root.

In the absence of any other factors, this trailing edge noise would be perceived subjectively as a constant 'whoosh' sound. The spectral content of this 'whoosh' sound would be controlled largely by blade geometry considerations. So, as previously indicated in the introduction, the physical size of turbines over the past two decades has seen an increase in the diameter, chord and thickness of the rotor blades. This increase in physical size of the blades has resulted in a general shift of the dominant frequencies radiated as a consequence of trailing edge noise generation mechanisms from between around 500 Hz to 1000 Hz for earlier turbines to around 300 Hz to 800 Hz for modern, large scale turbines, although there is some conflicting data on this decrease due to the overlap of the spectral regions between different turbine types of a similar size[9].

However, it is commonly observed that the sound heard by a listener located on the ground in the vicinity of a wind turbine is not constant, but instead it goes up and down in loudness in a regular manner. This effect is termed the 'Amplitude Modulation' of the aerodynamic noise, or AM for short. In subjective terms, the presence of AM causes the otherwise constant level 'whoosh' sound of the trailing edge noise to be perceived as a 'swish' sound, commonly referred to as 'blade swish', that is regularly varying in level. The regularity of the occurrence of peaks and troughs in the 'swish' sound is related to the rate at which the blades pass by the listener. Due to the directional radiation characteristics of the trailing edge noise, the difference between the peaks and troughs in noise level is greatest in the plane of rotation of the turbine, and in particular when the listener is located on the downward stroke side of the turbine. The same directional radiation characteristics mean that little temporal variation in the loudness of the trailing edge noise is experienced when the listener is located upwind or downwind of a turbine, particularly as the observer's distance from the turbine increases beyond a few rotor diameters.

This AM, which manifests itself as what is commonly referred to as 'blade swish' has been fully characterised through both theory and measurement. The theory has confirmed blade swish to be an inherent feature of wind turbine noise [6], although it has long been recognised as such. For example, blade swish is discussed in the ETSU-R-97 guidance document [3] published in 1996 (see, for example, pages 12 and 68 of that document).

[9] See WPC data and Annex B to this report.

Given the apparent acceptance of blade swish as an inherent feature of wind turbine noise, the question must therefore be asked as to why, after approximately two decades of wind farm development in the UK, the issue of AM has come to the fore? The answer to this question can not be definitively formulated, but it is believed to result from reports of a change in the characteristics of the AM being heard at some residential dwellings neighbouring wind farms. These changes include a general shift to lower frequencies of the dominant noise spectrum, even on an A-weighted basis, and an increase in modulation depth, in particular with higher depths of amplitude modulation occurring in the far-field downwind (and sometimes even upwind) of the wind turbine. These characteristics cannot be explained by current models of AM based on trailing edge noise. It is for this reason, for the purpose of the present report and as set out in the introduction, that all occurrences of AM falling outside that which can be explained through the trailing edge noise generation mechanism (i.e. 'normal' AM, or NAM) are referred to herein as 'other' AM, or OAM. Where the term 'blade swish' is used, this historically relates to occurrences of NAM.

At the time of the formulation of ETSU-R-97 [3], blade swish was considered in some detail:

> ETSU-R-97 page 12 - *'an amplitude modulation of noise in the frequency range which is associated with trailing edge noise radiated form the outer portion of the turbine blade and discrete frequencies associated with trailing edge thickness. This rhythmic swishing sound, dependent upon tip speed and blade profile, is normally centred around the 800-1000Hz region of the frequency band for trailing edge noise and at higher frequencies for trailing edge discrete frequencies depending on edge thickness. ….. modulation of the A-weighted noise level is of the order of 2-3 dB(A) for typical wind turbine configurations. …. this level of amplitude modulation may be greater if analysis is performed using third octave or narrow band analysis of the radiated noise from a wind turbine. This modulation may be caused by directivity effects associated with the generation of noise at the blade and is most apparent when standing close to a wind turbine, less than 50 m from the base of a supporting tower.*
>
> *As observer distance increases from the turbine, the rhythmic swishing becomes less pronounced. This may be due to a number of single effects or a combination. As distance increases, the modulation caused by the directivity of the radiated sound wave emitted by a turbine blade will become less significant. Therefore, it would be expected that any directivity effects which may be audible close to the turbine wil1 be reduced in audibility. Atmospheric attenuation will cause a reduction of high frequency blade noise relative to lower frequency blade noise. This removes the high frequency "swish" spectral content which increases its distinguishing character. As the observer distance increases, the level of sound from the turbine incident at the observer position will decrease. However, in exposed locations, it should be expected that the background noise level will remain, in general, the same. Therefore, increased masking by the background noise will reduce the subjective impact of the turbine noise. This rhythmic swishing has been noted to vary between turbine types and between sites where similar turbines have been installed.*
>
> ETSU-R-97 page 68 - *The modulation or rhythmic swish emitted by wind turbines has been considered by some to have a characteristic that is irregular enough to attract attention. The level and depth of modulation of the blade noise is, to a degree, turbine-dependent and is dependent upon the position of the observer. Some wind turbines emit a greater level of modulation of the blade noise than others, Therefore, although some wind turbines might be considered to have a character that may attract one's attention, others have noise characteristics which are considerably less intrusive and unlikely to attract one's attention and be subject to any penalty*
>
> *This modulation of blade noise may result in a variation of the overall A-weighted noise level by as much as 3dB(A) (peak to trough) when measured dose 10 a wind turbine, As distance from the wind turbine/wind farm Increases, this depth of modulation wouId be expected to decrease as atmospheric absorption attenuates the high frequency energy radiated by the blade. However, it has been found that positions close to reflective surfaces may result In an increase in the modulation depth perceived at a receiver position remote from a site, lf there are more than two hard, reflective surfaces, then the increase in modulation depth may be as much as ± 6dB(A) (peak to trough)[10].*

[10] The issue of increases in modulation depth close to reflective surfaces has subsequently been recognised as affecting both the peaks and the troughs of the AM noise to an equal extent, such that the effect in itself will not affect the difference in level between

The selection of the measurement position can also result in particular frequencies exhibiting a greater depth of modulation due to standing wave effects from reflected waves off the surrounding structures, These effects are very specific to the positions at which measurements are undertaken and are more the result of building layouts at the receiver position than a change in the character of the emitted wind turbine noise

It is the opinion of the Noise Working Group that there is insufficient data available at this lime to formulate an accurate measurement methodology for blade swish where it occurs. It is envisaged that further research will be required to enable proper measurement and assessment to be devised, if in the future this is felt to be necessary. Work is already under way aimed at establishing the causes of blades wish, the frequency and magnitude of its occurrence and developing an appropriate metric for its measurement.'

Section 2 of this report, together with Annex A of WPC, has considered in greater detail the current state of knowledge concerning AM and the reasons why its consideration has risen to the fore over the past five or so years. This was largely driven in the UK by the publication of a UK Government sponsored report, undertaken by the Hayes McKenzie Partnership, into the study of low frequency noise from wind farms [2]. The outcome results of that report prompted a second study, undertaken by the University of Salford, specifically addressing the issue of the prevalence of 'problem' AM across UK wind farms [1]. The issue was further highlighted by the publication of various research papers (see WPC Annex A).

The desktop based study undertaken by the University of Salford [1] was based on responses to a questionnaire issued to all local authorities with wind farms within their areas of control, the study concluded that, out of the then operational 133 wind farms across the UK, the issue of 'other' AM[11] could be identified at 4 wind farms. The study also raised the possibility that AM may additionally have contributed to complaints at a further 8 wind farms. Subsequent analyses of the responses to the University of Salford survey questionnaire [8] separately concluded that many more of the responses may have indicated the presence of AM as a contributory factor to the complaints. As pointed out by F. van den Berg, in all but one of the cases where a description of the sound was available (13 out of 14 cases) the sound was described with one or more words indicating a regular variation in level.

The foregoing perceived differences in positive responses to the questionnaire issued by the University of Salford serve to illustrate the difficulties facing the acoustics practitioner when interpreting responses to questionnaires of this nature. As so often proves to be the case, the devil is in the detail. In this case the detail concerns whether or not the complaints are related to 'other' AM, or merely to the fact that the 'normal' AM (blade swish) arising from the operation of a wind farm is audible and this, either in itself or in combination with the absolute level of turbine noise experienced, has caused those exposed to the noise to complain. This issue illustrates well the potential pitfalls associated with personal interpretations of subjective descriptors.

This brings us to the first question that needs to be answered in the context of the present project: exactly what is AM? As a consequence of collating available information during the course of the present project, what can be confidently stated is that the AM reported to occur from wind turbines covers the whole range of effects from mildly audible 'swishing' to much more pronounced 'thumping', that the effects may vary with distance and direction from the source and that, for any given location, the effects may vary with time over periods of a few seconds or they may present in a relatively constant manner over periods of several hours. Such observed variations may, however, sometimes be as much to do with variable masking noise effects as with variable AM effects.

Whilst the subjectively different audible effects of AM have previously been self-reported by those allegedly exposed to them using a whole raft of different descriptors, the problem with formally defining different types of AM is that there isn't necessarily a clear-cut dividing line between them, either in terms of physical characteristics or in terms of people's subjective response. Rather the observed effects exist along a continuum that may, or may not, see an abrupt change between states.

The issue of AM, and the differentiation between 'normal AM' (swish) and some form of 'other AM' has been usefully summarised by Bowdler and van den Berg [9] as follows:

peaks to troughs. Where increases in modulation depth have been observed close to reflective surfaces, they have generally resulted from the increases in the level of the peaks rising further above the general background noise.

[11] Although the term 'other AM' is not used in [1], the definition of AM used was '*A sharper attack and a more clearly defined character than normal blade swish […] like a distant train or distant piling operations*', which is consistent with the term as used in the present work.

> *'The presence of 'swish' near a wind turbine is a well known phenomenon. We know that swish is the type of AM which is ubiquitous next to a turbine. […] [It] falls off rapidly up and down wind, but transmits over longer distance at 90 degrees to the wind direction. However, in the latter direction the absolute level is significantly lower, so overall it seems that swish does not generally transmit over distances of a few hundred meters at levels that are likely to be a significant problem to neighbours.*
>
> *So what is the AM that troubles residents and why is AM a problem at some sites? […] Apparently residents are not bothered by this 'normal' swish but another form of AM which has been called 'thump'.'*

It is for all the foregoing considerations that the present project has chosen, as already set out in some detail in the introduction, to differentiate between 'normal AM' (NAM) and 'other AM' (OAM) in terms of:

- NAM is all AM that can be explained through current aerodynamic trailing edge noise source theory;
- OAM is any AM that falls outside the definition of NAM.

The convenience of this definition is that it means the causal mechanisms of NAM are well understood. What the following sections therefore present is consideration of the observed characteristics of OAM, and how these characteristics impact on the potential understanding of how it is generated, how it may be measured and objectively quantified, and how it may impact on subjective perception of wind turbine noise.

4.2 What causes OAM?

WPA1 and WPA2 have presented overviews of the potential different noise generation mechanisms in relation to the modulation of wind turbine noise. It is useful to view these findings in light of the observations made around actual wind farm sites, as reported in WPC.

The analysis of WPA1 is based on detailed operational and aero-acoustic parameters derived from the results of a specialised turbine design and analysis package, which was then used to model the noise radiation characteristics of a typical large, modern turbine design. The model's output has identified differences between AM in the near-field and the far-field. These differences represent a key consideration which has been relatively little studied to date. However, this represents a key issue at the heart of the question of amplitude modulation of wind turbine noise which may enable a better identification of the different potential AM effects, particularly through measurements in the far-field from wind turbines, as now discussed.

The expected decrease of NAM from wind turbines in the far-field may be due to several possible factors (see WPA2), as summarised below and addressed in further detail in the following section:

- air absorption dissipating the frequencies dominating the modulation of the turbines;
- the specific directivity of the modulating part of the noise produced;
- the reduction of directivity effects with increased separation distance;
- propagation effects such as refraction due to wind speed gradients;
- the effects of the noise from different turbines adding up together to 'smooth out' the modulation.

The situation is complicated by the fact that observations made in the far-field are undertaken in complex sound fields with low signal strength and therefore with the inherent potential for significant contamination from background noise sources. This is particularly problematic when considering measurements of modulation, as an increase in the noise floor will tend to reduce the measured depth of modulation (see WPB1). Therefore, studying the influence of periods of high wind shear, which has been cited by many observers as being a key contributor to higher levels of AM, is complicated by the fact that these periods often (in rural, relatively flat sites) correspond to increased levels of sound clarity because high wind shear often corresponds to reduced background noise levels. Thus it is not always possible to separate out the effects of the high wind shear at the turbine as opposed to its effects at the receiver (*i.e.* is the high wind shear actually causing higher levels of AM to be generated at source, or is the AM that was present in any event simply more audible under conditions of high wind shear due to the lower background noise at the receiver?).

One potential clue to the answer to this question can be found in the various observations WPC, WPD whereby high levels of AM are discernible at more distant receiver locations while there is no apparent increase in the levels of AM measurable closer to the source. Whilst these observations do not provide a definitive answer to the question, they do at least provide an indication that the existence of higher levels of AM at more distant receiver locations (for whatever reason) may not be as a direct result of changes to the level of AM experienced 'at source'.

Previous work

Initially, the work of Oerlemans [6] assumed stationary and uniform flow conditions which were considered 'typical' of standard operating conditions. This work established that, in close proximity to the turbine (i.e. less than 2 Rotor Diameters (RD)), substantial swish tended to occur to a similar degree in all directions around a turbine (as can be observed in practice in the field), whereas at increased distances (of 3 RD or more) significant AM only occurs in cross-wind directions, in which overall levels of noise are lower. Furthermore, as will be discussed below, we can also consider additional effects which will mean that the noise will tend not to propagate in these directions. As discussed in WPC, WPA1 the model was validated, to a high degree of success, based on field measurements of the type of turbine modelled for a range of measurements undertaken at a distance of approximately 3 RD from the turbine. Furthermore, the model was extended to far-field distances of up to 10 RD (WPA1).

It could be argued (in a perhaps naïve analysis) that the modulation will decrease with distance because the effect of the changing directivity of the trailing edge noise will decrease with distance due to purely geometric considerations. It may be thought that the relative angle of view to the blade may not change significantly at a greater distance away from the turbine. However, the results of the WPA1 analysis show that this is not the case in practice and that, in cross-wind directions, the predicted effects of directivity alone would remain significant out to large distances of at least 10 RD. A similar investigation was reported in [10] which was based on a flat plate model for the rotating blades. This paper reached similar conclusions in terms of directivity patterns as WPA1, specifically noting that the *'variation range of the directivity angles […] is invariant with the distances'*.

In both these models it was observed that the '45 degree' direction (i.e. between the downwind and cross-wind directions) was the direction in which the combination of both overall noise levels and modulation reaches a maximum (see also WPC]). It was therefore suggested (for example in [10]) that this would correspond to the maximum likelihood of audibility of the modulation. This does not, however, correspond to the observations of OAM in which high modulation levels were observed WPC in directly downwind directions rather than specifically in the 45 degree sector. This feature is also noted in WPA1, as well in [9] or in papers by di Napoli [11][12].

It should be appreciated, however, that it was in practice difficult to differentiate precisely between the instances of NAM and OAM and their respective directivity effects as a consequence of the non-optimised measurement data generally available at the commencement of this project. This is an area where the dedicated measurements undertaken under WPD add significantly greater certainty to the provisional observations based on the existing (generally ad-hoc with limited supporting information) measurements of AM reported under WPC. Notwithstanding this potential weakness of the existing available data, the work by di Napoli [12] further notes that the corresponding upwind OAM was less strong than the downwind OAM, and that it was weaker still in the crosswind direction. The observations of di Napoli also indicate that the AM was weakest close to the turbines but more prominent at larger distances (of up to 2 km) from the turbines [12].

One of the important conclusions of WPA2 was that a key theoretical condition necessary for high levels of OAM to be observed in the far-field downwind of a wind turbine is that the flow into the rotor is non-uniform. This would require either:

- that the wind profile is non-uniform, for example due to a vertical or lateral variation in wind speed or a spatial variation of the angle of the wind onto the rotor, where high wind shear or local wind gusts could provide the conditions for this to happen;
- that the turbulence entering the rotor disk is non-uniform due to upwind obstructions or meteorological conditions, thus causing time-varying levels of inflow turbulence noise as each blade enters the region of high turbulence.

Wind shear as a model of non-uniform flow

As part of the current research project, S. Oerlemans extended in WPA1 his existing model of trailing edge noise to account for non-uniform flow: using the specific and relatively simple case of an atmosphere with a certain degree of wind shear.

This analysis is based on a theoretical, engineering model, which does not necessarily aim to model every single aspect of what is admittedly a complex situation, but which provides an extension of an existing, well-validated model.

The approach undertaken in WPA1 also addresses concerns which have been raised in the past [13] that the wind profile due to high levels of wind shear would necessarily lead to increased modulation of the noise from the turbines. Such claims have centred on the possibility that the increase of wind with height, which can increase in magnitude in period of atmospheric stability, would directly lead to increased modulation as the noise emissions reach a maximum at the top of the rotation (TDC) where the wind speed will be higher, and that the modulation rate and depth would therefore be directly related to the rate of wind shear.

However the results of WPA1 demonstrate that even a relatively strong rate of wind shear (a shear exponent of m=0.6) across the turbine rotor does not lead directly to a significant change in the predicted modulation pattern or rate, but only to a relative small bias in the directivity. This finding is consistent with the results of a model by Boorsma and Schepers [14] which showed that the predicted effect of wind shear and the influence from the downwind support tower were minimal compared to the effect of the directivity of the trailing edge. These results may perhaps not be directly intuitive, but result from the reality of this situation and, in particular, the fact that the acoustic contributions of all 3 blades of the turbine must be added together as they rotate so, whereas the variation in sound emission levels for the outer part of any one turbine blade may vary strongly across its rotation (see Figure 1 of WPA1), the contribution from the other 2 blades will ‘fill in’ the troughs between the maxima which is experienced during the downward stroke of the blade.

Inflow turbulence as a model of non-uniform flow

The potential effects of inflow turbulence are reviewed in WPA2. The characteristics of this potential source of noise are similar in some aspects to stall noise (which will be considered subsequently) in that:

- its directivity is thought to be very similar, i.e. it represents a dipole radiating outwards from the blade (see [14]);
- it results in an increased dominance in the spectra in the lower frequency region (below 200Hz).

However, for modern turbine designs, it is often assumed that inflow turbulence will not represent a significant contribution to the overall A-weighted levels from the turbine, which are expected to be dominated by the trailing edge noise mechanism (see [14]). There is, however, some uncertainty as to how valid these assumptions would be, particularly when considering discrepancies between models and predictions at the lower end of the frequency spectra (below 200Hz). The emission levels of this source will also, of course, be strongly related to the amount of atmospheric turbulence present at different times, which may be very variable depending on site-specific factors.

As noted in WPA2, in order to represent a significant mechanism of modulation of wind turbine noise, the distribution of turbulence would need to be non-uniform across the rotor blade. The scale of variation would need to be significant: for a 7 dB AM depth, WPA2 estimates a 10-fold increase in turbulence intensity would be required. Furthermore, to lead to sustained OAM would require an equally sustained period of non-uniform turbulence distribution. Finally the observation of OAM during periods of high wind shear during stable atmospheric conditions at some sites does not suggest this mechanism to be the root cause, at least in those cases, as atmospherically stable periods correspond to reduced rather than increased levels of turbulence.

Detached flow

WPA1 notes that, for particular flow conditions, there may be two possible situations that may occur for certain sections of the turbine blade: the flow may be attached or detached (this is the aforementioned ‘stalled flow’), as shown diagrammatically in Figure 1.

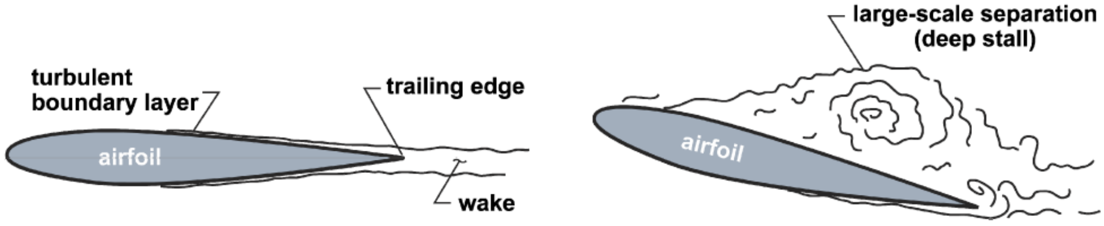

Figure 1: attached or detached flow on an aerofoil

Assumptions are made in WPA1 in terms of the exact level of noise emission from the stalled region. This was a necessary approach given the relative lack of current knowledge on this aspect of the noise generation. Nonetheless, it stands to reason that the basic assumptions are quite reasonable that noise emissions would increase significantly as a consequence of stall, and dominant frequencies would reduce (by about half) given the increased length scales, coherence and levels of turbulence associated with a significant stall region.

Importantly, the model in WPA1 predicts that the effect of partial stall on part of the blade fundamentally changes the nature and directivity of the modulation. When wind shear causes the turbine blade to experience detached flow for only part of the turbine rotation, the level of modulation in the far-field becomes significant upwind or downwind (rather than in cross-wind directions in the standard model), whereas the directivity pattern of the overall levels does not change significantly. This is particularly evident when comparing instantaneous noise footprints between the two cases: see Figure 2.

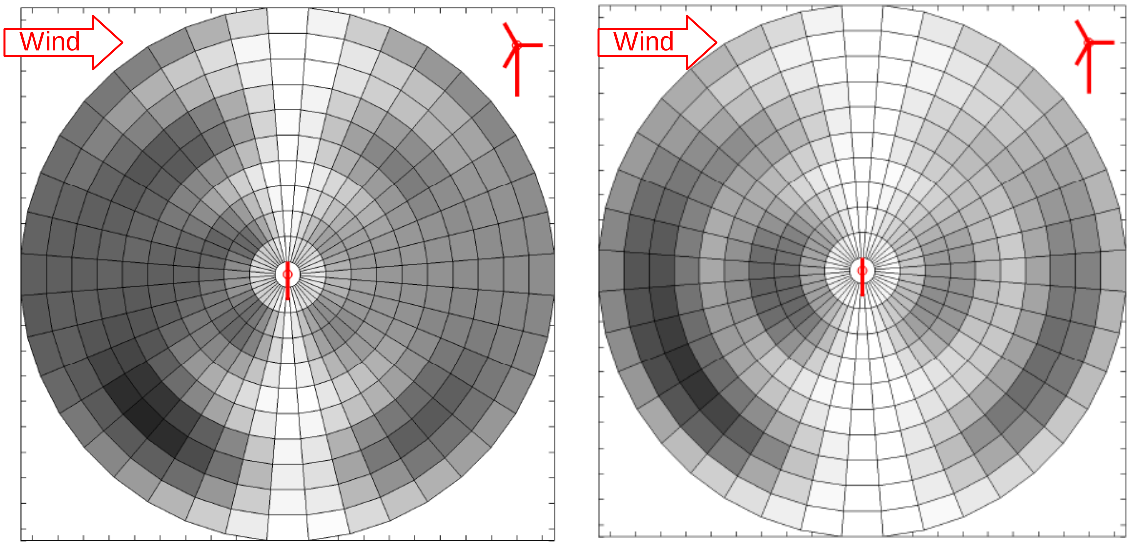

(a) Attached flow **(b) detached flow**

Figure 2 – Sample instantaneous noise footprints calculated for moderate wind shear, from WPA1 – white corresponds to low noise levels, and black to louder levels, and variations illustrate the presence of modulation, which is mainly present cross-wind in a (attached flow) and up/down-wind for b (detached flow).

When the modelled wind shear increases further, this increases the zone in which blade stall is predicted to occur but reduces the modulation amplitude: this is because the greatest variations in noise levels occur if the flow separation only occurs for a limited period of time.

WPA1 stresses that wind shear is only used as a simple model of non-uniform flow, for ease of computation but that, as detailed in WPA2, some other sources of flow non-uniformity may lead to similar localised stall effects, either temporarily or for more prolonged periods. Possible sources of non-uniformity are identified in WPA1 as wind veer, topography, large scale turbulence and the wake of other turbines.

The model presented in WPA1 doesn't initially show directly how modulation amplitudes of more than 6dB could be experienced in the far-field. However, this is shown by the author to depend strongly on assumptions regarding the features of the stall noise source and by how much stall would increase the local noise radiation from the stalled region of the blade. WPA1 has initially made an assumption of 10 dB which results in a modulation depth in the far field of up to 6 dB. Were this instead 13 dB, then the maximum far field modulation depth would increase to 9 dB. As this type of aerofoil noise source has received relatively little study it is easy to see how the assumptions made can affect the outcome conclusions. It is also possible, in line with the discussions of WPA2, that propagation effects may also enhance the modulation depths observed in the far field over and above the effects depths observed 'at source' or in the near field of a turbine, as discussed below.

The model presented in WPA1 does not directly explain either why this type of OAM would be more likely to be experienced in the far-field of the turbines as opposed to the NAM model. We can, however, consider, in the next section the assumptions and limitations of the model of WPA1 and how additional effects may explain the observations made in some cases. The potential influences of such effects are considered further in WPA2.

Limitations of the model in WPA1– propagation effects

The model presented in WPA1 does not take into account the longer-distance propagation effects, described in WPA2 (see also Chapter 3 of [9]), which are known to affect the long-range propagation of noise over larger distances. The model was effectively validated at distances of 3 RD, but not at further distances which are relevant to the far-field region in which wind farm neighbours can generally be found. Over these distances, the potentially most significant propagation effects are:

- spherical spreading of sound (which will reduce overall noise levels equally in all directions and is not considered further);
- refraction effects due to wind speed gradients in upwind/downwind;
- atmospheric absorption effects at high frequencies.

The potential influence on the predictions of WPA1 of the propagation effects to the far-field identified in WPA2 can be considered: this is developed in further detail in Annex B. It was noted that the effects of the foregoing factors were accounted for to some extent in a similar model developed in [14], but, as flow separation (stall) noise was not modelled by these authors, this could not be evaluated on this basis.

In summary of the considerations of Annex B, the following can be noted.

- **Atmospheric absorption effects:** are not considered to significantly affect the modulation predicted by the model, given the dominating frequencies characteristic of the dimensions of modern turbines. For OAM, it is predicted in the model that the peaks and troughs of the modulation will have different spectral characteristics (dominated by stall and trailing edge noise respectively), but because of the above conclusion, the effect of atmospheric absorption is not predicted to further enhance the difference in A-weighted levels in the observed modulation.
- **In down-wind conditions**, in addition to the inherent directivity of the turbine noise, the atmospheric refraction effects mean that far-field turbine noise levels are significantly higher in the downwind direction than in other directions. It therefore seems significant that, in the case of detached flow, significant modulation is predicted in the downwind direction. This key change in directivity is due to the inherent characteristics of stall noise compared to trailing edge noise, as shown in Fig 12 of WPA1 report. This is particularly evident when comparing instantaneous noise footprints between the two cases, as shown in Figure 2 above.
- **In upwind conditions**, overall levels tend to be lower generally, but the effect of the changing source height (as the blades rotate) may conceivably lead to this attenuation varying in some frequency bands as the blades rotate, as discussed in WPA2. This may combine with or affect the changed directivity effects of OAM illustrated in Figure 2. The influence of turbulence then provides a further complicating factor and can alter these predictions slightly.

In addition, the above consideration of propagation effects may further enhance the predicted effects of partial blade stall. As noted above, the predicted (under neutral conditions) peaks and troughs of the modulation in the downwind direction (in the far-field) correspond to different source mechanisms with different directivities. It is therefore conceivable that the refraction effects highlighted may in reality affect the resultant noise footprints in increasingly different ways when compared to the effectively neutral conditions modelled in WPA1. This effect could usefully be subject to additional, more detailed modelling to better understand the potential magnitude of such combined source and propagation effects.

Finally, the fact that propagation effects are being considered here as a possible contributor to OAM, as opposed to just source effects, is noteworthy, especially in the light of the (admittedly arbitrary) definitions of NAM and OAM proposed at the start of this section: namely that NAM is AM that may be, and OAM is AM that may not be, described by current aerodynamic noise *source* theory. Clearly, if propagation effects may also affect NAM, then this may become a relevant consideration.

In summary of the foregoing discussion, which has mainly focussed on WPA1/2 (with supporting observations from WPC) to review the potential combination of source generation mechanisms and propagation effects in exacerbating the presence of AM, there are some clear indicators that appear to be supported by observations in the field. However, as previously suggested, the problem with the available observations is that they have generally been derived from noise data acquired on an opportunistic basis, without detailed supporting meteorological and turbine operational status data also being available. For this reason it is not possible to draw firm conclusions. This is not to say that the observations discussed above are not useful in their own right, particularly in providing much-needed circumstantial evidence for feeding into the experimental plan for the dedicated measurement campaign undertaken in fulfilment of WPD.

4.3 Can AM be objectively measured and quantified?

The preceding section has defined NAM and OAM in the context of different source generation mechanisms. This is a potentially convenient differentiator. However, the possible contribution of noise propagation effects to the AM experienced at far field receptor locations represents a complicating factor in this analysis. In particular, it has been observed that OAM may be discernible at typical residential locations in the far field of wind turbines whilst not being easily determined in the near field of those turbines: although this is consistent with the theory of WPA1, this must be borne in mind when considering an attractive, but perhaps over-simplistic, definition based on source effects only.

A second possible means of differentiating between NAM and OAM is via their resultant acoustical characteristics, as opposed to their origins. These acoustical characteristics may be based on subjective descriptors, or they may be based on the outcome of objective analysis. What is ideally required is an objective analysis leading to clearly defined and measurable parameters that can then be related to subjective descriptors and subjective response.

The present project has therefore included extensive work on investigating various methodologies for the automated discrimination, analysis and quantification of OAM. The key objectives of this element of the work were twofold:

- first, a method was sought for the automatic detection of OAM in the presence of extraneous noise, with the selected method needing to be robust and repeatable both in terms of its ability to effectively discriminate OAM from other sources of noise (i.e. it should be resilient to returning false positives or negatives) and also in its subsequent objective quantification of the identified OAM element of the noise;
- second, the objective quantification of the OAM was required to be capable of providing a rating level that will enable the correlation of that level against subjective response to the noise.

As a useful starting point, the concept of the degree to which is a signal is modulated is a natural one to consider. It is for this reason that, hitherto, general attempts to relate a measurable characteristic of AM to subjective response have tended to focus on determining the variation in level of the overall A-weighted equivalent noise level, as measured on a short time basis of typically 100 ms or 125 ms: see [15] or [16]. In some cases the signal has been filtered in the frequency domain (typically by applying a low pass filter with a cut-off of around 500 Hz) prior to establishing the variation in the overall noise level, and in other cases the variation in individual octave or third octave frequency bands has been considered: see WPB1 or [17].

There are two key issues that arise from the foregoing. The first concerns the objectivity of the measurement itself, whilst the second concerns the appropriate measure of modulation depth in terms of subjective significance. The latter is particularly problematic, being circular in nature. In order to correlate subjective response with acoustic characteristics one must have an objective descriptor for the relevant characteristics, but the development of an objective descriptor requires some knowledge of the characteristics of the signal that are of interest subjectively and are therefore required to be quantified. However, the problem is farther reaching than this, particularly given the potential confusion that can arise from the meaning of the term 'modulation depth'. WPB1 provides a useful insight into this issue, from which it hopefully becomes clear that this apparently simple definition is anything but clear cut.

This particular element of the project, concerning the objective detection and quantification of AM, has been considered in WPB, with WPB1 focussing on the development of robust objective metrics (on the basis of the perceived dominant characteristics of OAM) and with WPB2 focussing on the subjective testing of the metrics derived under WPB1.

Developing a robust objective metric for AM

WPB1 presents an overview of various techniques allowing the quantification of modulated noise in general terms, and more specifically when applied to noise from wind turbines.

The first key result from this work is that there is no simple, single definition of the magnitude of modulation in an acoustic signal. When discussing the effects of the modulation of signals it is therefore crucial to be consistent in the definition of the metric and the normalisation that is adopted. This point cannot be emphasised strongly enough. A similar issue was identified in the research undertaken on behalf of ETSU for the assessment of tones [3], where the importance of applying a consistent set of analysis parameters to determine tonal levels was found to be central to the success of correlating objectively measured tone levels against subjective responses.

As a starting point for the quantification of modulation depth, the difference in peak to trough levels has understandably often been used (albeit, in retrospect, on a perhaps somewhat naïve basis) from readings of short term (typically 100 ms to 125 ms) equivalent continuous overall A-weighted sound pressure levels. This approach is not without merit. Indeed, WPB1 notes that the advantage of using this measure is that it will be invariant regardless of whether pressure signals are squared or not. However, whilst there is no ambiguity as to the definition of such a value for well-defined modulated signals, as is the case for the 'artificial' signals prepared for the present project to illustrate general principles, it is not so easy to apply this peak to trough measure in a consistent and repeatable manner for realistic field recordings. In particular, WPB1 notes that, primarily due to the compounding effects of extraneous noise, both the maximum and minimum (effectively instantaneous) values are particularly difficult to measure or evaluate when compared to the longer time averaged values (typically of at least 1 minute, and most often longer, duration) which are more generally used in the assessment of environmental sound fields.

Typical time-averaged metrics, such as L_{Aeq} or L_{A90}, are often analysed over periods of several minutes to evaluate different noise sources. Whilst this averages out short-term variations and can assist their evaluation, conversely it also erases information relating to variations over timescales of less than a few seconds. However, it is precisely these shorter term-variations in sound pressure level that are of prime relevance to the types of noise being studied. It is therefore necessary to capture the variations of the signal over timescales equivalent to less than a few Hertz, with a sampling frequency of around 100 ms being preferred.

However, considering the variations in terms of maxima and minima of typical short-term acoustic metrics can be particularly prone to difficulty. See, for example, Figure 3 below which considers the evolution of a short (20 second long) sample of noise for a variety of averaging periods, both with and without A-weighting, short-term energy-averages (L_{eq}) or for a time-weighted metric like the 'Fast' noise levels[12]. In this example, the short-term variations in the signal are mainly caused by bird noise. Given the significant variability in the real signal, it is not immediately obvious which 'troughs' and which 'peaks' should be considered when evaluating such differences, even using when trying to consider subjective restrictions placed on the analysis, such as the need to consider 'consecutive' or 'adjacent' samples. Apparent variations of between 5 and 15 dB can be noted depending on the metric and/or the timescales chosen.

[12] This time-weighting, using an exponential function with a constant of 125 ms = 1/8th of a second, is often considered representative of the response of the human hearing, although this may not be necessarily directly relevant to the present study in this sense.

The L_{eq} metric in particular will be sensitive to very short-term energetically intense events, which may be averaged out using other metrics.

Further significant variation could still be obtained on a similar analysis if the signal was further filtered for particular frequency bands, depending on the composition of the signal, as also discussed in WPB1.

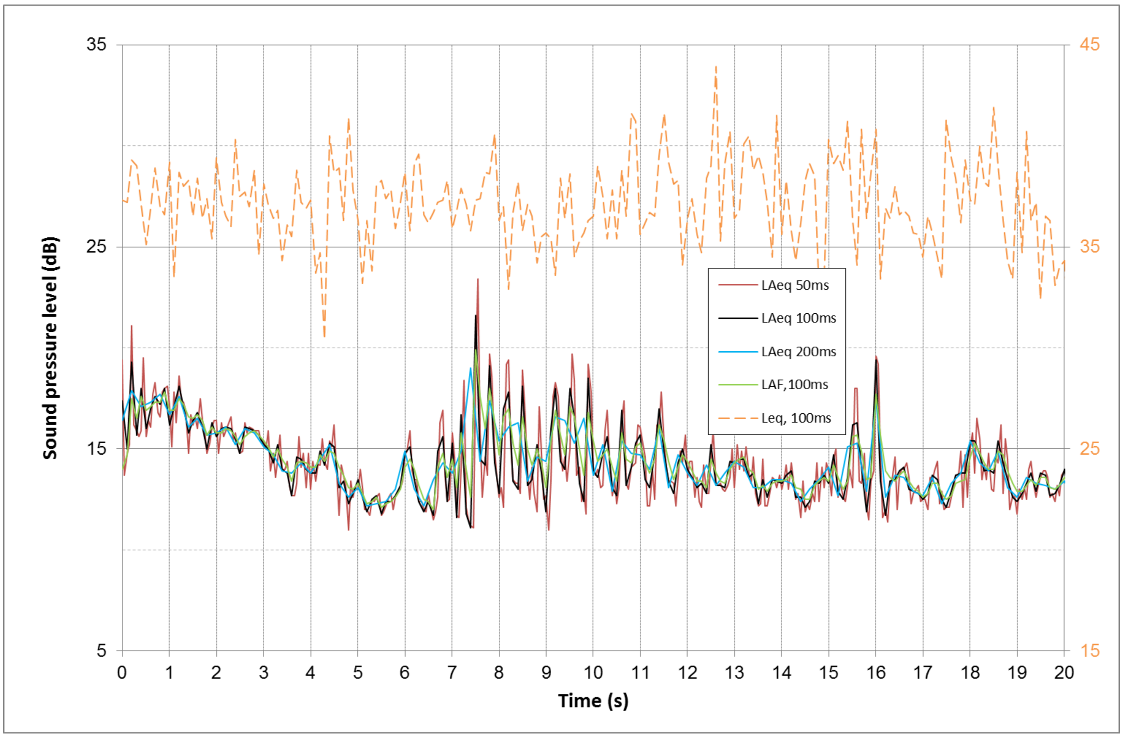

Figure 3 – Analysis of 20s of recorded audio data measured in a rural location in presence of wind turbine and bird noise (relative, non-calibrated levels) using a variety of short-term noise metrics

Short period measurement results can, in some cases, be usefully plotted out against time to graphically illustrate levels of AM (when present), although in practice only when there is a very clear and dominant contribution of the AM noise. With experience, it is possible in some cases to subjectively discern some patterns in the data which may be representative of different sources. This has led some [16] to suggest using such data to assess directly the character of the signal. However, as highlighted in the example of Figure 3, the process of trying to directly rate the modulation amplitude in this manner is too arbitrary, as picking peaks and troughs from all of those present is open to a degree of interpretation which is potentially too large to afford the method the required degree of objectivity and/or repeatability. Figure 3 is also an example of the potential for 'false positives', as other sources of environmental noise may lead to short-term variations in the measured noise levels which may be wrongly identified using this criteria, as shown in [15]. Perhaps most significantly, this type of method does not fully exploit the inherent characteristics of modulating wind turbine noise, as discussed below.

Other techniques which have been used in practice in similar contexts include:

- Comparing statistical metrics such as the difference between L_{A10} and L_{A90} when considering low-frequency tones [18]. Whilst such methods based on standard indices are easy to implement, and possibly effective in simple situations, the L_{A10} index in particular is likely to become strongly influenced by extraneous noise sources.
- Comparing the difference between L_{A90} and L_{Aeq}. A typical difference between L_{A90} and L_{Aeq} is described in ETSU-R-97 as being between 1.5 and 2.5dB for wind turbine noise (after excluding the influence of other sources). More recent experience has shown that this remains the case for modern turbines. As shown from the results of WPD, or even in the analysis of artificial stimuli derived as input to WPB(2) (see Table 22.1 of the WPB2 report), and perhaps contrary to informed intuition, this difference between L_{Aeq} and L_{A90} remains the case even for large modulation depths of up to 6dB(A) peak to trough.

Defining the problem

Whilst the general question of analysing signal modulation is difficult, we know that the modulation of the aerodynamic noise produced by a wind turbine noise occurs at a set frequency which is determined by the blade passing frequency (BPF), which in turn is determined by the rotational rate of the rotor. For variable speed machines, the exact modulation frequency in question will vary to some extent depending on wind conditions, but generally not to a significant extent[13]. Due to variations over time of the wind speed at any given turbine, or due to variations in wind speeds across a wind farm leading to different turbines simultaneously exhibiting different speeds, this variation may not be truly periodic but more likely 'quasi-periodic'. However, this quasi-periodic character, with the basic frequency being set by the basic BPF of the turbines, means that the AM of WTN can usually be set apart from other sources of changing noise present in rural or urban environments but which don't exhibit the same periodicity. This periodicity also allows the possibility for a model to be fitted to the data in order to rate the AM aspect of the noise more precisely.

Objective techniques

On the foregoing basis, WPB1 presents more refined signal analysis techniques which can determine more precisely the variations of particular features of a signal at a particular frequency. Importantly, WPB1 shows that the techniques presented represent the optimal way (in a precisely defined sense) of determining the modulation parameters in a measured acoustic signal.

- **The first method** assumes the basic signal being modulated is a white noise.
    - First a short-term envelope of the signal is obtained by squaring the signal and deriving energy averages over a short timescale (a Hilbert transform envelope technique produces similar results).
    - The Fourier transform of the envelope of the signal then provides a modulation spectrum
- **The second method** relaxes the assumption of white noise, allowing signals of more defined spectra (as will be the case in reality).
    - First, the signals are filtered in narrower frequency bands: this can be done through Fourier transform (narrowband analysis as in WPB1) or using 1/3 octave band filtering for example;
    - Another Fourier transform is then applied for each band in the time domain;
    - Finally the modulation spectra for results for all the different frequency band are summed up with a suitable weighting (such as A-weighting for example).
- **The third method** studied is more general in scope, in that it could detect a wider range of periodic signals. However, it was found in WPB1 to require a relatively high signal to noise ratio and is not discussed further.

Importantly, for improved analysis methods such as those described in WPB1, the analysis methods are much less sensitive to the resolution in time of the envelope data in terms of the returned levels, which was one of the issues noted for the naïve techniques described above based solely on consideration of variations in short term signal levels. The resolution of the more advanced methods would, however, affect the range of modulation frequencies which can be detected: to detect a modulation frequency of f_m requires data to be captured at a sampling frequency of at least 2 f_m (in accordance with sampling theory). For example, using a 100 ms time resolution (or 10Hz sampling rate) would allow modulation of up to 5 Hz to be detectable, whilst using 200 ms would decrease this to 2.5 Hz. However in both cases, the calculated modulation amplitude at 1Hz would be similar whereas, for example, the peak-trough variations in $L_{Aeq\ 200ms}$ and $L_{Aeq\ 100ms}$ shown in Figure 3 differ strongly.

[13] A typical variation between 10 and 16 RPM for a 3-bladed turbine will correspond to a variation from 0.5 to 0.8 Hz in Blade Passing Frequency, or a period range of 1.25 to 2 s in periods between each blade passing a set point.

As noted in WPB1, whilst the criteria for determining the presence of modulation can be defined precisely, there are potentially several metrics that can be defined to represent the magnitude of this modulation, and these can furthermore be normalised in different ways.

Potential confusion

When considering this type of modulation analysis, it is important not to confuse audible frequencies and modulation frequencies. By performing a modulation analysis, the envelope of the signal (as opposed to the signal itself) is transformed from the time to the frequency domain, as in the example below. The resultant frequency spectrum reveals the modulation frequencies. For the example considered in Figure 4 the modulation frequencies is approximately 0.8 Hz, with harmonics at approximately 1.6 Hz and 2.4 Hz. Had the noise signal itself (as opposed to the envelope of this signal) have been transformed from the time to the frequency domain, the dominant audible frequencies would have been found to exist over a range more typically extending from 300 Hz to 1000 Hz (i.e. a factor of typically 100 or more times the modulation frequencies).

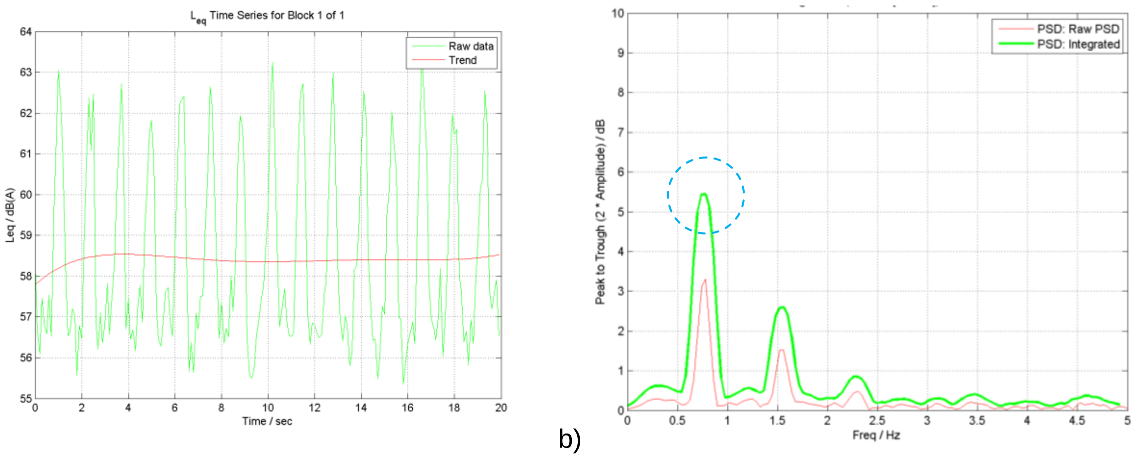

a) b)

Figure 4 – a) Time history ($L_{Aeq100ms}$, arbitrary scale, as a function of *time*) and b) resulting modulation spectra (as a function of *modulation frequency*), for an example of generated AM stimuli. Peak at modulation frequency highlighted in blue.

Normalisation and parameters

As discussed above, before considering subjective relationship and response to modulation at a certain level, it is crucial to consider how this level has been determined (*method*) and how the rating has been defined and scaled (*normalisation*), and also what inputs were considered (*parameters*). Despite the apparent theoretical simplicity of the subject for simple signals, and even some consistency in the overall methods, there is some variation in the latter two aspects in the literature.

For example, [19] uses the ratio of the stationary to the modulated components of the signal envelope as provided by a Fourier analysis (which is akin to the second, filter bank method referred to previously) to estimate a modulation factor based on the decibel peak-to-trough variation in the acoustic signal envelope.

In WPB1, the depth of modulation is initially defined as 10 $\log_{10}$ of the amplitude of the peak of the power spectral density obtained. WPB1 then describes a further normalisation to make the measure scale-invariant in which the data is scaled by a level corresponding to the mean of the quietest 25% of the signal as an approximation of the stationary part of the signal. The resulting amplitudes are, however, not intuitively intelligible.

Given the common use and consistency (albeit only when applied to signals dominated by AM) of the simpler peak to trough metric defined above, it may be considered more appropriate to design the AM detection algorithm so that it is normalised to provide identical values for well-defined, artificial signals for which the peak to trough values can be determined with precision.

It is important to consider the parameters for which the analysis is undertaken, in particular the duration of the data period blocks over which the modulation spectrum analysis is undertaken as, in effect, the modulation depth represents an average over the time period of analysis. WPB1 proposes 10s as practical example, but states that this will depend significantly on the stationarity of the modulation process, which is an indicator of how the amount of AM changes over time. A stimuli sample length of 20s was used for the stimuli of WPB2, which had a relatively constant modulation. Typical observed instances of intermittent modulation, shown for example in WPC, show variable levels of modulation over a scale a few seconds, and in some cases certainly less than 10 seconds. Intervals of less than 10 seconds would not provide enough data for a meaningful analysis.

Additional filtering

As explained in WPB1, given the nature of the signals encountered in the far-field of wind turbines, it can be useful to apply a low pass filter to the measured data prior to further envelope analysis. This will exclude most common bird call signals, as they are usually high-pitched sounds, and are most likely by their repetitive nature to corrupt the analysis of modulated signals in rural environments.

It is thus apparent that there exist a number of different facets of signal analysis that may come into play when determining the 'optimal' signal processing algorithm for the detection and objective quantification of AM. Whilst the various types of algorithm may be tested using both artificially generated AM signals, or real world signals containing high levels of AM, the performance of each algorithm is best judged against more typically encountered real world signals in which AM may well be immersed in a general ambient noise of a similar level to, or higher than, the AM itself. It is this application to real world data that is considered in the following section.

4.4 How does the theory for AM identification apply to real world data?

Whilst WPB1 provides an in-depth discussion of possible analysis tools for use in the investigation of AM, and in doing so highlights some of the pitfalls of adopting differing definitions of seemingly similar metrics, the test of any method designed for the automatic identification and objective quantification of AM must rely on its robustness when applied to real world data. Bearing this in mind, the techniques identified in WPB1 have been implemented as computer codes and tested on various data sets: this is detailed in Annex C of this report. The metric method used for this analysis will now be described.

A computer code implementing a signal envelope Fourier analysis technique in the MATLAB software was kindly provided by RES as part of their input into this project. This algorithm effectively represents an implementation of the first type of the methods described in WPB1, as it undertakes a frequency analysis of the acoustic signal envelope, using a Fast Fourier Technique (FFT), to produce a modulation spectrum. But if applied to the envelope of a narrower-band signal (such as 1/3 octave band data), it effectively becomes closer to the second type of method described in WPB1.

In this implementation:

- the signal is A-weighted;
- short term (1/8th or 1/10th of a second as in WPB1) L_{Aeq} energy averages are calculated from audio signals or directly provided as input;
- the envelope is separated into 'blocks' of a set length (typically 10s to 1 minute);
- the stationary component of the signal is removed using a de-trending technique;
- a power spectral density (PSD) of the envelope is calculated using standard windowing and fast-Fourier transform techniques;
- the PSD is normalised as $2\sqrt{2\,PSD}$;
- the resulting spectrum is then integrated in the frequency domain using a moving average of width equal to 10% of the estimated likely modulating frequency;
- a local maximum or 'peak' in the integrated modulation spectrum can then be determined close to the modulation frequency (with a defined tolerance, typically +/- 0.4Hz);

The method therefore requires a likely estimate of the modulation frequency (BPF) to be provided as an input, in addition to the parameters discussed above. The output provided for each data block comprises:

- a modulation spectrum;
- the amplitude and frequency of the main peak in this spectrum (near BPF).

Annex C to this report presents detailed results of the analysis of the implementation of this routine to signals of increasing complexity, starting with simple signals, and progressing to the artificial stimuli used for WPB2 as well as actual recordings of modulated wind turbine noise collected as part of WPC.

The normalisation and integration process undertaken means, when applied to simple and well-defined artificial signals with a trivial modulation pattern, the amplitude provided in the spectrum at the modulation frequency matches the peak-to-trough modulation depth of the A-weighted signal envelope.

The analysis of Annex C also shows that, using this AM metric routine for more complex signals, the calculated amplitude of the peak[14] of the modulation spectrum identified at the modulation frequency (as highlighted in blue on Figure 4) generally represents a useful and representative measure of the amount of modulation present in typical signals, which provides values which are consistent with subjective response and are similar to those used to date in the literature[15]. WPA2 has shown that the annoyance was not significantly related to the shape or duration of the signal, which suggests the effect of the higher harmonics[16] was not significant in the response and the main peak amplitude is sufficiently representative.

Annex C or WPB2 (Figure 9.6) show that, for the test stimuli used in the listening tests, the resulting AM metric values are approximately 1dB lower than those obtained from estimated peak-to-trough variations in A-weighted levels. The use of frequency filtering allows spurious sources to be excluded in most cases, and results in marginally more elevated values (close to the estimated peak-to-trough values).

This code will therefore be designated as the *main 'AM metric routine'* throughout the rest of this report for clarity, although it must be stressed that based on the review of the literature undertaken, there is currently no standard or generally accepted implementation of such metrics.

Modulation analysis - practical considerations

It is important, in the context of this project, to consider current practice in the environmental field measurement of wind farm noise immissions.

Planning requirements for most wind turbines in the UK follow the guidance of the ETSU-R-97 report [3], and measurements of noise levels are made at noise-sensitive properties on the basis of overall statistical indices over 10minute periods (L_{A90} in particular). The measurement of the shorter-term variations in noise levels is therefore not directly required, and would in general not be stored by many integrating sound levels meters which will only record these time-averaged metrics.

As the criteria on overall noise levels relate to A-weighted values, more detailed frequency information (such as octave-band or third-octave band data) is generally not collected or stored.

Audio records are also not directly needed for the analysis of overall noise levels or even to exclude spurious periods which could have been affected by anomalous sources of noise, as the analysis of the time history of noise levels or their relation to wind speeds is generally sufficient. However, the requirement to undertake tonal analysis will require audio samples at regular intervals, in order to collect representative recordings in a range of operating conditions. It is then general practice to collect 2 minute

[14] Local maximum in the modulation spectra.

[15] Normalisation of the output results is not an essential step. Rather, it is undertaken to produce output results for which the value at the modulation frequency is similar in level to the peak to trough difference seen in the original signal when that envelope is created using a time average of approximately 100 ms. In this sense the output results are easier to interpret on a physical basis. It should be noted, however, that for real world (time variant) signals, the output results will never match up exactly to the peak to trough level difference in each sample, even with such normalisation, as the peak to trough difference will vary with each modulation. Instead the method will provide an averaged measure of the peak to trough level differences across the sample. What is important is that, whatever normalisation process is adopted, it must be consistently carried through to all aspects of the analysis including any attempts to relate the metric to subjective response. The comparison of results obtained either using different analysis methods, or using different normalisation processes for the same analysis method, would not be meaningful.

[16] Peaks in the modulation spectrum which are found at integer multiples of the modulating frequency: see Fig 4b).

samples, as required by the method of ETSU-R-97, which can be done at regular 10-minute or hourly intervals.

Storage and power requirements to record detailed frequency or audio information are considerable, particularly when considering measurements made over periods of several weeks typical of wind turbine immission measurements (due to the need to capture a range of wind conditions). For example, continuous recording of A-weighted and 1/3 octave band at a 10Hz sampling rate, and period recording (recorded nearly continuously, i.e. for 80% of the time) of uncompressed audio at a 25.6kHz sampling frequency (therefore effectively capturing the frequency range of 0-10kHz), may require more than 800 Mb storage for 3 hours of recording. 99% of this size was taken by the audio data. This represents considerable practical and logistical difficulties, in addition to the significant costs involved.

The latest developments in the capabilities of noise measurement equipment have made this more accessible, as power and processing capabilities continue to improve. This type of complete and detailed measurement is nevertheless considered challenging and may not be easily undertaken by most consultants or local authorities.

Recording audio data therefore involves considerable difficulties, but the lack of this data limits the range of post-processing that can be undertaken on the data, such as filtering out spurious noises (which is not possible if only A-weighted levels are recorded). For the study of wind turbine noise however, relatively low sampling frequencies may be sufficient as, for example, a 4 kHz sampling rate would capture with reasonable accuracy audio frequencies of up to 2 kHz.

WPB1 describes filter bank methods, which are based on an analysis over a range of narrowly filtered bands of data. The application of the first type of method in WPB1 to a single 1/3 octave-band represents a simple example of such a filter bank method. But when considering the analysis of different WPC samples which is undertaken in Appendix B, it can be observed that the analysis of the short-term evolution of the variations in certain 1/3 octave bands can be sufficient to characterise the modulation in many circumstances. This approach is also shown to be implemented with a good success in WPD to large sets of data. It is then seen that this provides an efficient way of eliminating several types of spurious types of sources such as bird or wind noise, without the need to record further detailed audio information.

Compared to acquiring audio data, octave band data is more difficult to interpret subjectively when trying to evaluate the character of the noise environment being analysed. To this end, it may be useful to also acquire shorter term samples audio data on a periodic basis throughout any measurement campaign (as is done for the narrow-band, tonal analysis in which 2 minutes of audio data is required to be recorded for every 10 minute period).

It may also be useful to investigate in further detail the application of such filter techniques to multiple 1/3 octave band simultaneously, which would allow this full frequency information to be used meaningfully, along the lines discussed in WPB1. However, it is thought that this may provide relatively limited advantages compared to the consideration of a single, well-chosen band which is representative of dominant modulation frequencies, provided the latter can be determined (as shown was possible through the results presented in WPD).

The expected modulation frequency would in practice be relatively well known, as it will be determined by the blade passing frequency. This can therefore represent an input to the analysis of modulation. For a typical three-bladed turbine rotating at R rotation per minute (RPM), $f_m = 3R/60$. For variable speed machines, as is the case for most modern turbines at present, the main shaft rotational speed of the turbine will usually be recorded by the turbine SCADA control system, and this could therefore inform the variation in f_m with time. The range of variation may not however be very significant in practical terms. Foe example, for a typical large turbine with a rotation rate varying between 10 and 17 rpm over its typical operating range, this would represent a variation between 0.5 and 0.8Hz. Knowing the expected range of modulation frequencies may be sufficient to identify modulation likely to be associated with the turbine. Again, WPD shows an example of how this can be applied in practice to provide more certainty to the analysis.

WPB1 did not consider in detail metrics such as fluctuation strength [7], which are based on psychoacoustic models, but these were considered in WPB2. However the practical implementation of such methods is subject to considerable difficulties, particularly in the presence of complex realistic signals (rather than test stimuli), and would in practice rely on specific software implementation of such methods which often depend on specific models of loudness and simplifying assumptions to some extent.

Such metrics are not scale independent and would therefore rely on careful control of calibrated sound levels, in ways that are difficult to control in practice. Finally, such methods require audio signals and are computationally intensive.

4.5 What is the subjective impact of AM?

WPB2 identifies the difficulties arising with characterising AM, and the lack of sufficient knowledge in the literature about listener response to its characteristic physical properties. The scope and also the potential limitations of the exercise undertaken are summarised in the following table. Bearing in mind that practical considerations will always result in some limitations of any such exercise, the important consideration here is that any conclusions drawn from the results do not stray outside the bounds of what can reasonably be concluded within the identified limitations.

FEATURE	COMMENTS
Use of synthesised AM samples	Necessary as control required over signal parameters, and removal of variability of real world signals and extraneous noise
Use of short term samples each with fixed (i.e. non-variable) AM effects	Although the level of AM is known to be highly variable and intermittent, it was not practicable to assess these effects in a laboratory setting
Considering external noise environments only	Based on previous experience, simulating internal levels cannot be done conclusively and involves significant difficulties regarding the necessary assumptions concerning outside to inside sound transmission and the resultant low signal levels. There would also be limited relevance to the current practice of assessing free-field levels from wind turbines.
Use of headphones in the sensitivity measurements	Done to assess relative impact of different parameters to restrict test size as opposed to directly rating AM noise
Seemingly arbitrary choice of objective metric against which to assess the subjective test results	A simplistic metric was chosen as a starting point to help design the stimuli, and then all results were subsequently evaluated over a wide range of objective metrics
High scatter in test results	Seen as inevitable when testing affective responses such as annoyance, but needs to be taken into account when interpreting significance of results
Limited sample size	It might be useful to extend the sample size, but this has been considered unlikely to fundamentally change the outcome results
Use of 'Direct Annoyance Rating' based on subjects relaxing at home in their garden	Very difficult to ascertain true annoyance unless subjects actually exposed to the noise at their own homes, but equally difficult to consider an alternative in a laboratory environment given overall aims of the study, which is to provide comparative results.
Use of 'Paired Comparison'	Following from above point: this allows more direct measure of a modulated signal compared to a similar level of un-modulated noise, without relying on an arbitrary absolute rating of annoyance in a non-home environment
Comparison between 'Direct Annoyance Rating' and 'Paired Comparison' results	Care needs to be exercised with any such comparison due to the dominant effect of increases in overall noise level over and above the effect of increased modulation depth (which had little statistical significance)

The first objective of Work Package B2 was to test whether the AM metrics developed in Work Package B1 would provide a meaningful measure of AM 'value' that correlates with subjective annoyance ratings. The second objective was to quantitatively investigate the relationship between the AM value and a measure of average annoyance in the form of a dose-response relation.

Key to obtaining representative results was the design of stimuli which were representative of the spectrum and character of actual noise experienced by wind farm neighbours exposed to AM. Extensive work was done to obtain this. Based on input from WPC, test signals were synthesised for a characteristic range of wind turbine sounds, with a wide range of input parameters. This model is described in an annex to WPB1. The model consisted of “pulses” of modulation overwritten on a constant “masking noise” which determined the effective modulation depth. It was the general consensus of the project team that the artificial stimuli obtained were representative of AM experienced in the field, as collated under WPC.

In a first phase, sensitivity tests were undertaken to find the AM parameters that listener response was most sensitive to. A total of 80 test sounds, each of at least 20 second duration, were presented via calibrated headphones in a quiet room and 11 volunteers were asked to score the annoyance on a numeric 11 point scale. The outdoor sounds included test wind turbine noise at typical levels, with varying AM characteristics and some natural background noise.

Using the results of the sensitivity tests, a sub-set of samples was selected to cover those variables having the greatest impact on subjective response. A final set of tests was then undertaken in a quiet listening room with a sound reproduction that mimicked the outdoor directivity of one wind turbine in the distance. A total of 34 test sounds were generated and presented to 20 participants. Two validation tests, containing another 158 and 34 test sounds respectively, were also conducted to clarify results from the headphone based sensitivity tests in the better-controlled listening room. Participants rated annoyance directly as before, but were additionally requested to adjust an un-modulated test sound in level such that they judged it equally annoying as the modulated test sound.

The sensitivity tests showed, in line with previous literature, that annoyance crucially depended on the overall A-weighted level of the test sound, as measured in L_{Aeq}, and to a lesser extent on modulation depth, which is a measure of the modulation strength. It was shown that the response to the noise was not significantly affected by the modulation waveform. Modulation depth was shown to be also best expressed in terms of A-weighting to give consistent results, as initial differences in the response to different types of spectra was found to be explain by increased overall dB(A) levels. The use of L_{A90} as an alternative to L_{Aeq} produced similar results at the low and medium modulation depths most often observed from wind turbines, as the difference between L_{Aeq} and L_{A90} was comparable for most stimuli below modulation depths of about 9dB(A).

In the final test there were three sets of test sounds that were played back with constant L_{Aeq} of 30, 35, and 40 dB(A). For each of these sets the modulation depth was systematically varied from 0 to 12 dB(A) in increasing steps. After taking into account the effect of L_{Aeq}, which was found to always dominate the annoyance rating, the modulation depth was found to increase the annoyance rating slightly, but consistently. However, the effect was not statistically significant because there was a large spread of ratings. See Figure 5 for an example representation of these results.

The above suggests that, given a large enough group of participants, it could possibly be shown that annoyance increases slightly but consistently (monotonically) with modulation depth. In contrast, the L_{Aeq} level of the adjusted un-modulated wind turbine noise remained broadly constant as the modulation depth increased above about 3 dB(A). This answered the question of how much louder would an equivalent un-modulated sound have to be to be equally annoying to a modulated sound. The adjustments were on average 1.7 dB(A) for a 40 dB(A) test sound and 3.5 dB(A) at 30 dB(A). Validation tests at two additional levels of 45 dB(A) and 25 dB(A) confirmed this trend. See Figure 6 for an example representation of these results.

A clear onset of annoyance at a particular modulation depth was not found for either of the two rating methods.

When levels were measured as L_{A90}, results suggest that annoyance ratings were similar for modulation depths of up to 6 dB(A) and generally increased with both modulation depth and L_{A90}. Because results for sets of stimuli with constant L_{A90} and changing modulation depth are not available simple average adjustments cannot be identified and further work would be necessary.

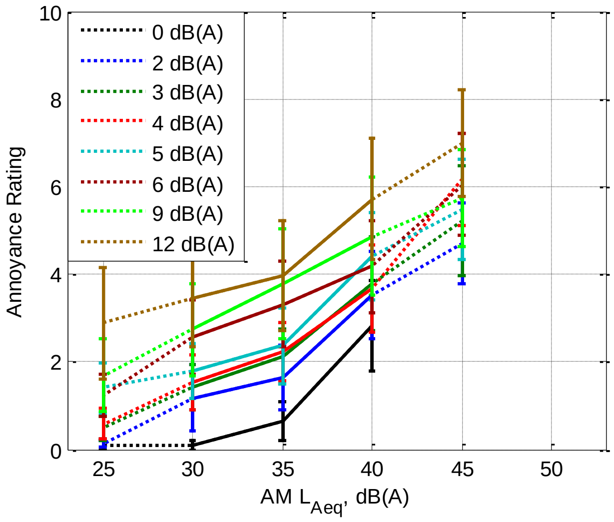

Figure 5 - Mean annoyance rating of AM test stimuli as a function of estimated modulation depth. Solid lines are results from final test, dotted lines from the validation tests with reduced participant numbers.

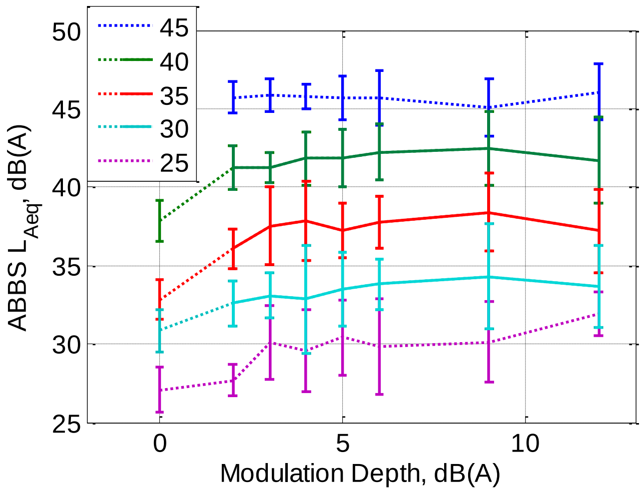

Figure 6 - Levels of adjusted un-modulated noise in comparison to AM stimuli level, as a function of estimated modulation depth – Solid lines are results from final test, dotted lines from the validation tests with reduced participant numbers. The legend specifies the L_{Aeq} of the test stimuli in dB(A).

In a validation test with a subgroup of 11 participants, the spectral characteristics of the test sound were changed to represent Mid-Frequency AM, often described as 'swish', as opposed to Reduced Frequency AM which is sometimes described as a 'swoosh' or 'whoomp'. Also, a limited amount of garden noise[17] was added at a low level to change the character of the sound for both types of AM sounds. For all four groups the results for both absolute annoyance ratings and un-modulated level adjustments appeared very similar. This suggests that the relative effect on annoyance is small as long as the garden noise does not reduce the audibility of the modulated sound.

In a last step the annoyance ratings were compared for 6 different metrics, four of them based on different physical definitions of modulation depth and two using the perceptive measure of fluctuation strength [7]. Given the low sensitivity of the response to modulation depth, the comparison showed that the main effect of the physical metric is to change the range of modulation depths. The same stimuli

[17] The term 'garden noise' is used here to describe a typical ambient noise expected to be experience in a rural outdoor environment due to the natural sound of the wind blowing through the trees. The use of the more usually encountered terminology of 'masking noise' is avoided as, for the purpose of WPB2 in which the feature under investigation was the modulating component of wind turbine noise, 'masking noise' included both 'garden noise' and the steady component of wind turbine aerodynamic noise.

would have a range of 0 – 12 dB(A) modulation depth in one metric but 4 – 32 dB in another metric: this highlights the importance of relating any specific response to measured AM levels in a consistent way. Fluctuation strength results showed an improved correlation with listener response, accounting to some extent for the overall level of the noise, but, as noted above, these psychoacoustic methods are difficult to apply in practice to realistic field data. As noted above, even a perception-based metric can never fully account for contextual and attitudinal aspects of annoyance rating, and the result of survey studies should be considered in parallel.

The effect of factors such as the frequency of occurrence of OAM 'events' and their duration not addressed in the current research. The project team's view is that the significance of these factors would have to be assessed using professional judgment and experience and that it would not be practicable to assess them via further subjective testing.

4.6 What are the characteristics of OAM when measured

In the light of the results of WPD as well as Annex C of this report, it should be borne in mind that: even a carefully filtered and scrutinised analysis of a real measured dataset in a typical rural environment, using the main type of AM metric routine considered in this work (defined above in 4.4), will produce values of the AM rating metric of up to 2 dB (with a limited amount of sporadically higher values), even in the absence of any significant modulation. This is mainly due to residual noise, uncertainties associated with digital processing and the integration process used.

The presence of OAM was able to be detected and rated effectively using the main AM routine described above (see Figure 7 for example). Perhaps counter-intuitively, it is remarkable that, even in periods of marked modulation, the difference between the measured L_{Aeq} and L_{A90} metrics remained typical of that expected from wind turbine noise, at between 1.5 and 2.5 dB.

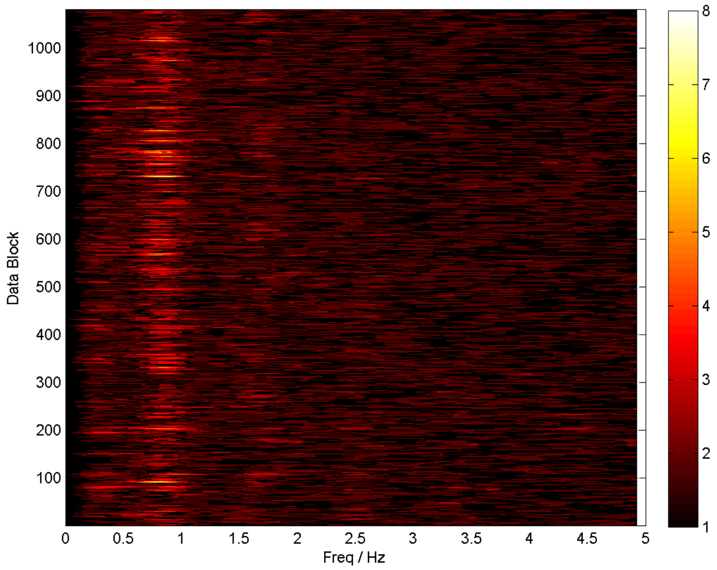

Figure 7 – Evolution over a period of 3 hours of the calculated modulation spectrum based on the 315Hz 1/3 octave band (for contiguous 10s data blocks) at site A, location 1

Based on the different measurement results, and particularly those at Site B, it has been concluded that the general characteristics of OAM noise are consistent with the expected directivity and spectral characteristics of transitory stall noise (as derived in WPA1). Indeed, in the far-field, instances of clear OAM were enhanced in the downwind direction and reduced in the cross-wind direction. Observations in which OAM was discernible in the far field of wind turbines but not simultaneously in the mid- or near-field suggest that AM may be strongly influenced or exacerbated by noise propagation effects (such as ground effects).

At site A, the presence of OAM was clearly established during conditions in which the surveyed locations were downwind of the turbines (see Figure 7 above), whereas for another period in which the location was upwind, no significant levels were detected, which matched subjective reports from one of the residents.

The effects in the near-field were more difficult to discern, although the expected presence of NAM was clearly characterised by higher modulation depths of up to 5 dB in the cross-wind direction, as illustrated in Figure 8 below, which is consistent with the standard AM model [6].

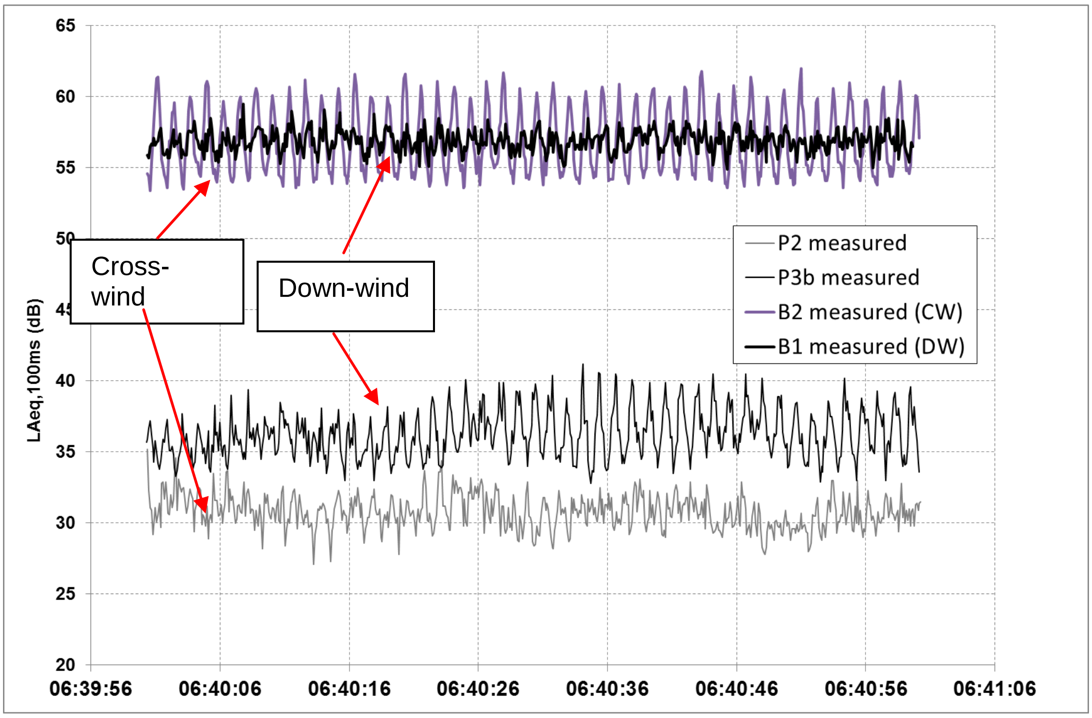

Figure 8 – WPD, site B: time history of measured $L_{Aeq,100ms}$ levels for a sample period showing: two far-field (downwind and cross-wind) and two near-field locations (cross- and down-wind)

The magnitude of OAM rating levels in the far-field were strongly variable and did not seem to be simply associated with the existence of certain meteorological conditions. In terms of the various hypotheses that have historically been as possible causal mechanisms for other AM, whilst the results could not generally rule out any of these as potential contributory factors, they did confirm the ability of OAM to exist in situations where the factors are known not to contribute. In summary, significant OAM was positively identified under conditions of:

- low wind shear (see Figure 9 for example);
- low wind veer;
- uniform turbulence;
- single operational turbines (i.e. no interaction effects);
- on both flat and hilly sites;
- turbines with high tower to rotor diameter ratios.

The only positively identified association between the occurrence of OAM and the operational characteristics of the turbines, in the detailed measurements undertaken at site B of WPD, was that OAM only occurred when active power generation was occurring, and it also appeared to be sometimes exacerbated during periods when changes in the estimated relative angle of attack of the blades also occurred.

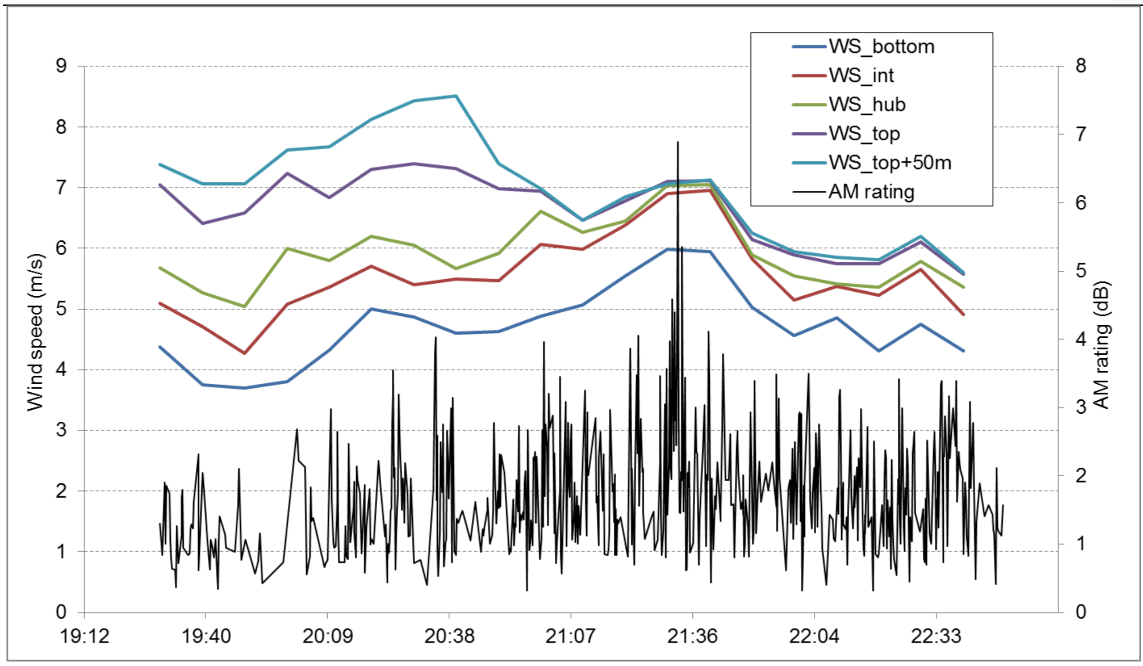

Figure 9 – Time history of the AM rating for a sample period a site B, in parallel the changes in wind profile at different heights.

Other source mechanisms may be at play. It is conceivable that particular blade designs could experience partial blade stall in certain conditions without the influence of non-uniform inflow conditions to trigger it. For example, aero-elasticity effects may vary the blade geometry in such a way that it departs significantly from its design conditions. There is, however, very limited information available to study this in the current state of knowledge.

HLA has also seen some evidence of measurements undertaken at another site (not reported in WPD), in which periods of OAM were associate with night-time periods in which extremely high values of wind shear were regularly experienced, which suggests it could be a contributory factor in some cases, but the above consideration show this is not generally necessary or sufficient and may therefore be site-specific.

It is also the case that, on some sites, the impact of wind shear on effective modulation may be more important at (non-sheltered) residential location surrounded by vegetation. This is because increased wind shear will correspond to relatively reduced wind near the ground, and this will have an effect on the level of masking background noise which may otherwise reduce the effective modulation depth.

In conclusion to this question, the definitive answer as to what causes OAM remains elusive. What has been proven, however, is that many of the causal factors previously claimed to be indicators of an increased likelihood of OAM are not, in their own right, sufficient to trigger the effect. The observed far-field OAM directivity pattern was also shown to be consistent with predictions of the effect of transitory blade stall.

It has additionally been noted that, for the purpose of the present project, NAM and OAM have been defined in the context of different source generation mechanisms. This is a potentially convenient differentiator. However, the possible contribution of noise propagation effects to the AM experienced at far field receptor locations represents a complicating factor in this analysis. In particular, it has been observed that OAM may be discernible at typical residential locations in the far field of wind turbines whilst not being easily determined in the near field of those turbines: although this is consistent with the theory of WPA1, this must be borne in mind when considering an attractive, but perhaps over-simplistic, definition based on source effects only.

4.7 Can OAM be predicted?

As discussed above, a variety of causal factors have been hypothesised by various acoustics practitioners as potential causal factors for OAM. In light of the available evidence, and the theoretical consideration summarised in the appendix to WPA2, as supplemented by the results of WPD, the main factors are considered in the following table 2.

Cause raised	Significance	Evidence and factors
High wind shear: stable atmospheric conditions or presence of ground obstacles such as vegetation	Low-medium	• Large variety of sites with different shear characteristics in which OAM was still experienced; for example: presence of OAM in hilly sites with low shear • May cause transitory stall in some conditions • Weak/negative correlation evidence in WPD • But increase of signal to noise ratio in high wind shear conditions may increase OAM perception if present
Complex flow/turbulence (large-scale)	Low/medium	• Presence of OAM in flat sites with high wind shear conditions (low turbulence) • Large non-uniformity of turbulence required in theory, or large-scale turbulence triggering stall • High variability of AM in time would be consistent with turbulence variations, but sustained periods of OAM sometimes observed are clearly not. • May cause transitory stall in some conditions • Weak/negative correlation evidence in WPD
Interaction between turbines / linear arrangements	Low	• Presence of OAM with single turbine operation as shown in WPD and [11]
Synchronisation between turbines	Low	• Theoretical considerations and presence of OAM with single turbine operation as shown in WPD and [11]
Yaw error or wind veer	Low	• Weak/negative correlation evidence in WPD • See high wind shear above as both are often related
Power regulation (pitch or stall regulated)	Unclear	• Turbines at one of the sites in WPC were stall-regulated (DTI HMP study [2], site 2). • Pitch-regulated turbines present in many of the OAM cases considered, but these are generally more prevalent today in any case
Non-uniform inflow	Unclear	*See above*
Angle of attack	Medium/High	• Likely to be strongly related to incidence of stall • Some correlation shown in WPD with periods of high power and variations of angle of attack
Ground interference, atmospheric attenuation	Low	• Little AM effect in theory, present at all sites • Some evidence in WPD suggested ground interference reduced modulation peaks in some conditions
Background noise masking	Medium	• Will limit the modulation depth at large distances
Temperature gradients	Low	• Likely to be limited except in calm conditions
Aero-elastic effects	Unclear	• May trigger stall, but complex and little evidence

Table 2 – OAM causal factors and review of evidence

Several of the potential causal factors which have been suggested were shown through the present project to have little or no association to the occurrence of OAM. However some of these factors may represent potential contributory factors. It is not therefore possible to be prescriptive as to whether any particular site is more or less likely to give rise to OAM being generated at source. This is considered likely to be due to a combination of site and installation specific factors.

It is only where background noise at noise sensitive receptor locations is high enough under all circumstances to effectively mask any potential OAM noise that it can be positively concluded that the chances of OAM causing problematic noise issues are reduced.

4.8 Can OAM be suitably mitigated?

It would be desirable to develop OAM mitigation measures based solely on the operational characteristics of the affected turbines (i.e. modifications to the turbine control systems, and in particular the blade pitch control systems) as opposed to physical modifications to the turbines themselves or operational curtailment of turbines.

Partial blade stall has been identified as a key potential causal factor in the occurrence of OAM. The effective angle of attack of the flow on the blades in theory represents a key parameter affecting the likelihood of this occurring. The results of WPD have also shown that this variable sometimes demonstrated a correlation with clear instances of OAM, although the overall picture was relatively complex. As noted previously, additional work may usefully identify the exact on-blade conditions in more detail using further investigative measurements which were outside of the scope of the current study.

For modern variable speed, pitch-regulated turbines, the turbine's control system will change its operating parameters to vary the pitch and rotational speed as a function of different observed parameters. Detailed information on this process will often be proprietary and subject to commercial confidentiality considerations. It was speculated during the progress of the project that certain power regulation control approaches for pitch-regulated turbines may result in a greater effective 'stall margin', by which is meant a greater difference between the actual angle of attack and the angle of attack corresponding to full blade stall, as this may conceivably reduce the likelihood of partial blade stall occurring.

Given the current state of knowledge in the field, the Annex of WPD presents indicative analysis techniques which may assist in evaluating the changes in relative angle of attack on the blades, from an analysis of long-term data provided by the turbine SCADA[18] system. Although based on considerable simplifications, and the general need for the analyses to be based on coarse 10-minute resolution SCADA data, the general approach may enable key operational conditions to be identified which are more conducive to blade stall. The identification of any such operational conditions could assist in the development of modifications to the turbine's operational characteristics in collaboration with turbine manufacturers which could positively mitigate against the occurrence of blade stall.

However, the project team are not aware of a situation where this potential method of mitigation has been applied successfully. Evidence has been seen from one site where OAM was being experienced at several locations. At this site, alternative power regulation modes of the pitch-regulated turbine installed were designed to minimise the incidence of potential stall from a turbine, by reducing the relative angle of attack experienced by the turbine in certain conditions identified (based on similar techniques to those of the Annex of WPD). Some initial results suggested that the measures implemented in this regard were not sufficient to fully mitigate the incidence of OAM, and further work in this regard was required.

In some situations, the use of serrated edges on blades has led to a significant reduction in the noise emission from some turbines without the associated reduction of power generation generally obtained with reductions in blade rotational speed. This is because this reduces the radiating efficiency of the trailing edge of the turbine [see Ref. 9, chapter 2]. This type of measure is, however, not considered likely to affect the radiation of stall noise from the blades, which tends to occur at other parts of the blade.

Standard 'reduced noise modes' are commonly used to reduce the levels of overall A-weighted noise produced by many modern turbines. These are often based on reducing the rotational speed of the turbines (and therefore the tip speed). Given the foregoing discussion, there is also no evidence or theoretical grounds suggesting that the use of such modes would affect OAM. In fact, the use of reduced noise modes may even be counter-productive as a reduction in rotational speed could increase the effective angle of attack of the flow.

There is, therefore, currently no clear case history of successful mitigation of OAM noise, except through curtailment of turbine operation in the specific conditions in which it is encountered in the far-field, which can in some case only cover a restricted range of wind speeds and/or wind directions. Given the impact

[18] Supervisory Control And Data Acquisition system, generally installed on wind farms to control and monitor the operation of all turbines.

this can have on power production, it is strongly recommended that additional research is undertaken to identify clear and efficient mitigation strategies that minimise or eliminate the generation of OAM whilst retaining the power generation capacity of the wind farm.

This would require the future involvement and close cooperation of wind turbine manufacturers, and possibly involve detailed measurements that focus on better understanding the surface pressure distributions on the turbine blades themselves, particularly as the stall point is approached. Additional research would then result in methods for avoiding local stall conditions. The Annex of WPD presents some simple calculation methods which have been identified as likely to assist in identifying the most likely relevant conditions likely to lead to stall, based on an analysis of standard turbine operational data.

Such methods may involve software ‘fixes’ that seek to modify the logic of the control system algorithms, perhaps even through the application of more advanced cyclical pitch control, and therefore a limited impact on the operation of the turbine. More fundamental, physical design changes may also prove worthwhile, such as innovative blade designs or the addition of blade vortex generators, for example, as the latter may delay the onset of stall. More advanced numerical modelling, such as for example extending the model of WPA1 but including propagation effects, may also provide further insight.

5.0 CONCLUSIONS

It is an inherent feature of wind turbines that they generate noise when their blades rotate. This is caused by the interaction of the blades with turbulence in the flow and is therefore called 'aerodynamic noise'. The dominant self-noise mechanism when dealing with an A-weighted spectrum of wind turbine noise is so-called trailing-edge noise.

For a listener located on the ground in the vicinity of a wind turbine, the rotation of a turbine's blades results in this aerodynamic trailing edge noise going up and down in loudness in a regular manner. These peaks and troughs in the noise will therefore occur regularly and their occurrence is related to the rate at which the blades pass by the listener.

This regular variation in the trailing edge noise is termed the 'Amplitude Modulation' of aerodynamic noise, or AM for short. This type of AM has been fully characterised through both theory and measurement. It is an inherent feature of wind turbine noise and has long been recognised as such. For this reason this type of AM has been termed in the present work as 'normal' AM, or NAM for short.

Due to the directional radiation characteristics of trailing edge noise, the difference between the peaks and troughs in NAM noise is greatest in the plane of rotation of the turbine, i.e. in a cross-wind direction. The same directional radiation characteristics mean that, in comparison, little temporal variation in the loudness of the trailing edge noise is experienced when the listener is located upwind or downwind of a turbine, particularly as the distance from the turbine increases beyond a few rotor diameters. Based on theoretical models and experience to date, NAM would not be expected to be apparent (except at insignificant levels) downwind/upwind of 'far-field' locations typical of wind farm residential neighbours.

However, in some circumstances the character and spatial distribution of the AM has been observed to alter from that known to result from NAM. Differences include a general shift to lower frequencies of the dominant noise spectrum, even on an A-weighted basis, and an increase in modulation depth, in particular with significant levels of amplitude modulation occurring in the far-field downwind (and sometimes even upwind) of the wind turbine.

The foregoing characteristics cannot be explained by current models of NAM. For this reason, the occurrence of any type of AM that falls outside the known characteristics of NAM is identified in the present work as 'other' AM, or OAM for short. Thus the definitions of NAM and OAM adopted for the project were based on consideration of the physical source generation mechanisms alone: NAM was defined as that element of AM which was capable of being fully described in terms of 'standard' models of trailing edge noise, whilst OAM was defined as being any form of AM lying outside this definition of NAM.

The occurrence of OAM has been cited in a number of complaints concerning noise from wind farms as being the prime cause for complaint. At the time of commissioning of the present project, little was known about OAM. Despite a limited number of reported observations, OAM has formed the subject of an increasing number of papers and publications over the past five or so years. However, whilst adding to the body of evidence relating to the existence of OAM, most such reports were either anecdotal or involved measurements at far-field locations. Little or no supporting information was generally available concerning the noise at the turbine(s), nor the local operating conditions of the turbine(s). Discussions on potential causal factors have been mainly speculative.

Since the causal mechanisms have not been understood to date, no specific information has been available to guide operators towards the likelihood of occurrence of OAM or remedial actions which may be required. Furthermore, where OAM is known to occur, there has been no universally accepted means of measuring its magnitude or determining whether complaints from neighbours are justified. The aim of the present project was, therefore, to investigate AM and in particular OAM further.

How does OAM differ from NAM?

Work undertaken by others prior to the present project has successfully demonstrated, both theoretically and experimentally, that the key acoustic characteristics of NAM (otherwise often referred to as 'blade swish') for modern, large scale wind turbines may be summarised as follows:

- characterised by an audible and regular variation in sound pressure level (i.e. an amplitude modulation of the noise);
- the frequency of the amplitude modulation is equal to the rotational rate of the turbine multiplied by the number of blades, which is otherwise termed the 'blade passing frequency';
- the typical frequency of the amplitude modulation for modern, large scale, variable speed wind turbines lies in the range 0.5 Hz to 1 Hz;
- the maximum 'modulation depth' (i.e. the variation in sound pressure level) exhibited is up to approximately 5 dB when considering the overall A-weighted noise level;
- the maximum modulation depth occurs in the 400 – 1000 Hz range;
- the depth of the amplitude modulation is most pronounced in the near to mid-field of the turbine (i.e. within a few rotor diameters) in cross-wind directions;
- the overall level of aerodynamic noise radiated from a wind turbine is at a minimum in the cross-wind direction (i.e. in the plane of the rotors) meaning that the maximum depth of the amplitude modulation (beyond a few rotor diameters) occur coincident with the lowest overall levels of noise;
- the amplitude modulation reduces significantly with distance, especially in the downwind or upwind directions, and is negligible in the mid- to far-field when the observer is close to the axis of the turbine (i.e. directly upwind or downwind);

Through the definition adopted for the present project OAM is, by default, any amplitude modulation that does not fit with the foregoing description. However, as a starting point for the project investigation, an analysis of available (often anecdotal) information indicated that OAM could be characterised by one or more of the following features:

- the modulation depth (the difference between the levels of adjacent peaks and troughs in the noise signal) can be significantly greater than that of normal blade swish, with differences in level of up to 6 to 10 dB having been measured, and subjective descriptions of "impulsivity";
- the effect is generally strongest in the downwind direction and has also been reported (less frequently) in the upwind direction of turbines, but it has not been recorded in the cross wind direction in which normal blade swish is most prevalent;
- the dominant frequency characteristics are sometimes lower than for normal blade swish, with a shift in the dominant frequency range to typically around 400 Hz;
- the effect is more dominant in the far field (typically 10 rotor diameters or more from the turbine) and may not even be simultaneously discernible in the near field (typically less than 3 rotor diameters) of the turbines;

What causes OAM?

In terms of addressing the possible causes of OAM, in the first instance it was considered helpful to collate available reports of its occurrence from as many different sources as possible (wind farm operators, windfarm neighbours, researchers, internet published information, etc.). The results of this exercise (WPC) lead to the following observations:

- the effect has only been reported on a limited number of wind farm installations;
- even on those sites where the effect has been positively identified to occur, it is intermittent;
- the effect may occur for just a few rotations of the blade, or it may persist for periods of several minutes or hours;
- the effect is not restricted to an individual turbine type;
- the effect has been reported both for individual turbines and for wind farm arrays;
- a particular turbine type that exhibits the effect on one wind farm site will not necessarily exhibit the effect on another site;
- the effect has been reported to be most common during evening and night time periods, although it has been identified during some day-time periods;
- the effect has been reported near wind turbine installations on both flat and hilly terrain;
- the effect has frequently been associated with conditions when wind shear may be expected to be high, but has also been reported to occur on some sites during damp conditions of light or even heavy rain when wind shear may be expected to be low.

If any one factor or combination of factors were responsible, then the effect would occur at all sites featuring those factors and it would occur frequently at those sites. Neither of these situations was observed to be the case in practice. In summary, therefore:

- NAM is an inherent feature of all wind turbines and can be fully explained as a consequence of the rotation of the wind turbine blades;
- OAM is not a common feature of all wind farms and, even for those wind turbines and wind farms where its occurrence has been reported, it is an intermittent and atypical feature.

The project team was therefore unable, based on the available information, to identify a single causal factor giving rise to OAM. It was therefore decided to focus on the identification of possible source mechanisms that could give rise to the observed acoustical characteristics of OAM, where these differ from those of NAM, and most notably in terms of their lower frequency content and different directivity characteristics.

The theoretical study undertaken in WPA2 identified two potential source mechanisms as being either high levels of inflow turbulence or stalled flow over part of the blade. However, the same analysis concluded that a key additional condition being necessary for high levels of OAM to occur in the far-field downwind was that the flow into the rotor would have to be non-uniform in order to result in the observed cyclical variations.

Possible causes of non-uniform flow into the rotor disc include:

- the wind profile being non-uniform, for example due to a vertical or lateral variation in wind speed or a spatial variation of the angle of the wind onto the rotor, where high wind shear or local wind gusts could provide the conditions for this to happen;
- the turbulence entering the rotor disk being non-uniform due to upwind obstructions or meteorological conditions, thus causing time-varying levels of inflow turbulence noise as each blade enters the region of high turbulence.

However, some of these factors are almost always present, for example variations in wind speed with height above the ground (vertical wind shear). Following the line of the first of these two potential causes of OAM, one plausible explanation was that a combination of environmental and turbine operational factors may come together to result in the airflow over small regions of blade becoming ‘detached’ from the blade surface and then quickly 're-attached' such that the blade goes into 'transitory stall' in the affected regions.

A theoretical model was developed in WPA1 for the radiation of transitory stall noise from wind turbines. This model has demonstrated the following key characteristics:

- transitory stall occurring over a restricted blade area can cause a noticeable, but equally transitory, increase in noise radiation from the affected areas of the blade, with the dominant acoustic frequencies of this stall noise being lower than those resulting from NAM;
- this increased low-frequency noise radiation occurs repetitively at a rate related to the blade passing frequency of the rotor, which is the same as for NAM;
- the source directivity characteristics of the stall noise are such that it is preferentially radiated upwind and downwind of the wind turbine and not in the cross wind direction, which is opposite to the source directivity characteristics of trailing edge noise (associated with NAM);
- noise is also preferentially radiated and propagated in the downwind direction of a wind turbine so, unlike NAM, the maximum modulation depth of OAM will occur in the same direction relative to the turbine as the highest overall levels of noise.

In the light of the expected key characteristics of OAM noise, and in particular the notable differences between OAM and NAM, targeted measurements were undertaken across three separate wind farm sites to measure OAM and to confirm its characteristics. This included innovative procedures based on direct modifications of the operational characteristics of a turbine.

The results of these targeted measurements (WPD) have concluded that the general characteristics of the OAM noise observed were consistent with the expected directivity and spectral characteristics of the transitory stall noise: a significantly increased effect was particularly apparent in the downwind direction from the wind turbine, with a frequency composition consistent with this model. Modifying a turbine's pitch when operating, in order to trigger stall, also appeared to generate some increased modulation events, as expected from the theory.

The project has additionally concluded, however, that the total validation of the foregoing conclusion would require further dedicated measurements. These measurements should focus on the surface pressures of the blades themselves in order to directly detect the occurrence of detached (stalled) flow, with the aim of correlating this blade-measured pressure data with measurements of the resultant radiated noise.

The available data and results of the targeted measurements have also been used to test various, previously proposed hypotheses as to possible factors that may give rise to for OAM. Whilst the results of the targeted measurements could not absolutely rule out any of these as potential contributory factors, they did confirm the ability of OAM to exist in situations where the factors were known not to contribute. In summary, OAM was reported or measured (WPC, WPD) under conditions of:

- low wind shear;
- low wind veer (wind direction variation with height);
- uniform turbulence;
- single operational turbines, without wake interaction effects;
- both flat and hilly sites;
- turbines having small or large rotor diameter size relative to their tower.

The only positively identified association between the occurrence of OAM and detailed turbine operational characteristics obtained at one site of WPD was that OAM only occurred when active power generation was occurring, and it also appeared to sometimes be exacerbated during periods when changes in angle of attack of the blades were estimated to occur.

Thus, whilst the project has positively concluded that the source mechanism of OAM is almost certainly transitory stall effects, the results have not allowed the similar positive identification of the factors causing such transitory stall effects to occur. As set out above, several of the potential causal factors which have previously been suggested were shown through the present project to have little or no association to the occurrence of OAM. However some of these effects may represent potential contributory factors, and the evidence further suggests this would be very site-specific.

It is not therefore possible to be prescriptive as to whether any particular site is more or less likely to give rise to OAM being generated. This is considered likely to be due to a combination of site- and installation-specific factors, including meteorology.

This conclusion will, of course, depend on situation-specific circumstances, and most notably the level of ambient masking noise present at the location of the far field observer. If sufficiently high, such masking noise may serve to negate any subjectively perceptible OAM noise. One factor which can affect the relative level of masking noise on some sites is the level of wind shear, in particular for (non wind-sheltered) residential locations that are surrounded by vegetation. This is because increased wind shear will correspond to relatively reduced wind speeds near the ground, with little change to the wind seen by the turbines. So even if the amount of wind shear does not directly affect the generation of modulation, it may have an effect on the level of masking background noise which may otherwise reduce the effective modulation depth.

Other source mechanisms may be at play. It is conceivable that particular blade designs could experience partial blade stall in certain conditions without the influence of non-uniform inflow conditions to trigger it. For example, aero-elasticity effects may vary the blade geometry (as they bend and twist under load) in such a way that it departs significantly from its design conditions. There is, however, very limited information available to study this in the current state of knowledge. It is therefore recommended that further (blade-focussed) measurements be undertaken to establish the potential significance of such effects.

Notwithstanding the conclusion that changes in source effects are likely to be the primary cause of OAM, the project has additionally concluded that OAM may be strongly influenced or exacerbated by noise propagation effects. Observations in which OAM was discernible in the far field of wind turbines but not simultaneously in the near field adds some weight to this possibility.

Environmental noise propagation effects generally result in the upwind radiated sound being rapidly attenuated with distance due to upwards refraction effects, whilst the downwind radiated sound propagates further due to the downwards refraction of sound waves. In addition, turbines do not radiate high levels of noise in the cross-wind direction. So whilst overall levels of turbine noise are normally highest in downwind directions, they normally do not fluctuate significantly. OAM noise resulting from transitory stall on the surface of the wind turbine blade is preferentially radiated in the upwind and downwind directions, as opposed to the cross-wind direction for NAM. When associated with these general directivity effects, this would result in the OAM modulation being more subjectively audible in the far-field than NAM because of the increased overall levels.

A theoretical possibility has additionally been identified whereby propagation effects would result in modulation being perceived or even enhanced in the far field in the upwind direction at certain frequencies. This would be due to the interaction between source-receiver geometry and refraction effects: whereby the noise radiated from the uppermost part of the blade’s trajectory would be able to penetrate into the upwind acoustic shadow zone, whereas the noise radiated from the lower part of the blade’s trajectory would not similarly be able to penetrate this same zone. Thus, for a certain critical range of distances that would depend on local considerations, blade radiated noise would be perceived to increase and decrease at a rate related to the blade passing frequency. However, this has not been commonly observed to be a feature of NAM. As OAM is predicted to radiate in the upwind direction as well, these identified propagation effect may influence this modulation and/or cause the upwind attenuation, which is generally expected, not to materialise in some instances: this may explain some limited observations of OAM upwind made in the field. Nonetheless, the existence of OAM has been found in practice to be prevalent mainly in downwind conditions (see for example site A of WPD), in which case its characteristics have been shown to be consistent with the existence of a source effect.

It has additionally been noted that, for the purpose of the present project, NAM and OAM have been defined in the context of different source generation mechanisms. This is a potentially convenient differentiator. However, it should be noted that OAM has been shown to be discernible in the far field of wind turbines whilst not being simultaneously discernible in the near field to those turbines, which indicates caution is required when regarding this attractive, but perhaps over-simplistic, definition based on source effects only. These observations also highlight the practical difficulties involved with this phenomenon as far-field measurements are always more complex in practice.

Can OAM be objectively identified and quantified?

The present project has included extensive work on investigating various methodologies for the automated discrimination, analysis and quantification of OAM.

It has been discussed how short period measurement results (typically of $L_{Aeq,100ms}$ or $L_{Aeq,125ms}$) can, in some cases, be usefully plotted out against time to graphically illustrate levels of AM. However, it has also been demonstrated how this really only works in practice when there is a very clear and dominant contribution of AM noise. In all other situations, the process of trying to directly rate modulation amplitude in this manner proves to be too arbitrary, as picking peaks and troughs from all of those present is open to a degree of interpretation which is potentially too large, and too readily influenced by spurious noises, such as bird noise, to make the method sufficiently robust in terms of its objectivity and/or repeatability. The resultant derived modulation depths are also dependent on the analysis period used (e.g. 50 ms or 125 ms).

Crucially, WPB1 has highlighted how the analysis of OAM can rely on the fact that the noise is quasi-periodic.

For variable speed machines, the exact frequency of the quasi-periodic variation in level will vary depending on wind conditions, and from turbine to turbine across a wind farm site, but generally not to a significant extent. The fact that the modulation frequency will normally and consistently lie within a fairly restricted range (typically from 0.5 Hz to 1 Hz) means that OAM (when present) can be set apart from other sources of noise present in rural or urban environments, whose periodic variation in level may occur outside the expected frequency range for wind turbine noise, or indeed may not be periodic at all.

On the foregoing basis, the project has presented an assessment of Fourier Transform (FT) based signal analysis techniques with the aim of identifying a particular modulation frequency in a noise signal, and then establishing the variation in level at that frequency. It was shown that these methods are optimal in determining parameters of quasi-periodics signals.

It is important to note here that the transform from the time to the frequency domain using the tested analysis techniques is not undertaken on the signal itself, as this would merely result in the spectrum of audio frequencies within the signal. Rather, the transform is undertaken on the envelope of the signal such that the result is the rate at which the overall signal level modulates (i.e. the modulation frequency). The signal envelope adopted is usefully based on that calculated using, for example, the $L_{Aeq,100ms}$ metric. Thus the analysis techniques developed as part of this research use the recognised merits of the first method discussed above, whereby the $L_{Aeq,100ms}$ or $L_{Aeq,125ms}$ may be plotted out against time to graphically illustrate levels of AM, but adds an additional level of sophistication to isolate those variations in noise level that are specifically and consistently related to the rotational rate of the turbines, therefore lessening the dependence on the signal resolution.

The calculated amplitude of the peak of the modulation spectrum then represents a measure of the amount of modulation present in AM signals which is consistent with subjective response. Results from analysing field recordings of OAM, including those containing higher levels of extraneous noise, suggest that, in general, methods based on the Fourier Transform of the overall energy in a signal can perform well. The downside of such an approach, however, is that their performance tends to reduce as the relative amount of extraneous noise increases.

Applying appropriate frequency-domain acoustic filtering prior to the computation of the energy has been shown to enable the refocusing of the processing on the audio frequency band(s) where the modulation is significant (for example by applying a low pass filter with a cut-off frequency of around 500 Hz or by undertaking the analysis on octave band or third octave band filtered data), thereby reducing the effects of extraneous noise at higher frequencies.

Whilst the identified approach requires the *a priori* identification of a suitable frequency filter band (or bands) in which OAM is prevalent, and it cannot significantly reduce the effects of extraneous noise that may occur in the same filtered frequency range as the modulation, it has nevertheless been demonstrated to be generally robust when applied to actual field-recorded signals. Although there is currently no standard or accepted implementation of such routines, a particular implementation of this method was used principally in this work to analyse new and existing recordings, and was designated as the “main AM metric routine”.

In summary to the above, the project has concluded that an efficient and repeatable objective method for the discrimination and quantification of OAM can be developed. It was found that the resulting metric can be consistent with subjective response to OAM but that input parameters needed to be chosen with care. The key features of the preferred metric comprise:

- it is based on an analysis of noise that is pre-filtered to lie in an audio frequency range in which the effects of OAM are known or expected to be most prevalent;
- it is based on a Fourier Transform from the time to the frequency domain of the envelope of the resultant filtered noise, with the envelope typically being calculated from the short term (~100 ms) $L_{eq,t}$ of the filtered signal, with the choice of the same averaging period being consistently adopted;
- the Fourier Transform analysis is focussed on the range of frequencies for which OAM is known to occur for the turbines being measured (typically 0.5 Hz to 1 Hz for current large scale turbines);
- the resultant energy that is calculated to exist at the identified modulation frequency (the ‘modulation energy’) is normalised using a consistent methodology;
- if the resultant ‘modulation energy’ is used to rate OAM in terms of its subjective impact, it is of central importance that the subjective response tests used to derive a rating scheme were based on analyses undertaken using the same objective analysis techniques throughout.

What is the subjective impact of OAM?

Controlled subjective testing was undertaken in a specialist facility at the University of Salford to supplement the limited amount of available knowledge on the listener response to characteristic physical properties of AM wind turbine noise. Simulated recordings, based on an analysis of actual field recordings, and with a wide range of input parameters, were played back to a range of up to 20 subjects of different ages and sensitivity but normal hearing.

The study necessarily relied on tests carried out under controlled laboratory conditions. Rating annoyance is subject to contextual and attitudinal issues and these factors are thought to be responsible for the wide error bands in the plotted data. However, the results of the WPB2 listening tests are generally consistent with those of existing research into subjective response to amplitude-modulated noise.

The frequency spectra and levels of sounds were intended to represent the characteristics of wind turbine noise as it might be perceived in a rural garden. Subjects were asked to rate the noise in two ways: an absolute annoyance rating, and a rating relative to un-modulated noise (by the subject adjusting the levels of modulated and un-modulated noise to achieve the same annoyance rating). These ratings were correlated with mean noise level and a range of metrics defining the degree of modulation.

The initial sensitivity tests first undertaken, as well as the more detailed and controlled final set of tests, concluded that the subjects' response to the noise was not significantly affected by:

- the frequency spectrum of the modulated noise (either dominated by medium or lower frequencies), once the A-weighted level was taken into account
- the modulation waveform
- the presence of limited amounts of vegetation noise in addition to turbine noise

In contrast, the annoyance ratings were significantly related to:

- the frequency (rate) of the modulation;
- the A-weighted average level of the test sound;
- the modulation depth (a measure(s) of the modulation magnitude).

Regarding the first point: in the range studied, annoyance increases with frequency (i.e. a faster turbine rotational speed is more annoying), which is consistent with psychoacoustic theory. As discussed above and in Annex B, the increase in the size of modern turbines has been associated with a general reduction in the audio frequencies dominating the modulation: other parameters being equal, this does not appear to result in significant differences in the annoyance response. This was accompanied by a decrease in the rotational rate of the turbines, which would correspond to a relative reduction in annoyance. The other two key factors identified were then analysed in further detail.

After controlling for the mean overall level of sound, as measured by the L_{Aeq}, the final tests showed this parameter always dominated the annoyance rating. The modulation depth was found to increase the annoyance rating slightly, but consistently; however it should be noted that the effect was not statistically significant because there was a large spread of ratings. This is consistent with previous results in comparable studies. A clear (significant) onset of annoyance at a particular modulation depth could therefore not be determined.

The tests for which an un-modulated wind turbine noise was adjusted for comparable annoyance with the AM stimuli resulted in levels which were relatively constant from modulation depths of approximately 3 dB(A). The adjustments were on average 1.7 dB(A) for a 40 dB(A) stimuli and 3.5 dB(A) at 30 dB(A).

The use of the L_{A90} as a measure of noise levels resulted in comparable results for moderate modulation depths of less than approximately 9 dB(A), although further work and testing would be required to establish alternative corrections directly. It should be borne in mind that field measurements have established that for real data, the difference between L_{Aeq} and L_{A90} remains close to 2dB(A) even for periods of sustained modulation. L_{Aeq} levels remains strongly susceptible to corruption from other sources of noise in the environment, and that the L_{A90} is therefore used in practice as a proxy for the L_{Aeq} of the wind turbine noise in isolation.

Although psychoacoustic metrics can additionally represent the effect of overall noise level (to some extent), their practical application involves significant difficulties. For the other various objective metrics of modulation depth which were studied, including the main AM metric routine used in other parts of this project, the levels of subjective response were comparable, given the low sensitivity of the response to modulation depth. This comparison mainly highlighted the importance of relating any specific response to measured AM levels in a consistent way: when discussing the effects of the modulation of signals it is therefore crucial to be consistent in the definition of the metric and the normalisation that is adopted. This point cannot be emphasised strongly enough.

The effect of factors such as the frequency of occurrence of OAM 'events' and their duration is not addressed in the current research. The view of the steering group and the research team is that the significance of these factors would have to be assessed using professional judgment and experience and that it would not be practicable to assess them via further subjective testing.

Can OAM be effectively mitigated?

In answering this question, the first issue that must be considered is whether the available evidence points to the fact that OAM is a 'source' issue, or whether it could potentially occur at distances typical of residential neighbours to a wind farm as a consequence of propagation effects in the absence of any source effects?

Although a theoretical propagation effect was identified which could produce a modulation effect in upwind conditions, this did not appear consistent with general observations. The existence of OAM has been found in practice to be prevalent mainly in downwind conditions, in which case its characteristics have been shown to be consistent with the existence of a source effect.

If the existence of OAM in the far field requires some effect to have occurred at source, even though the effect may be exacerbated by propagation effects, then the control of the source effect will remove the occurrence of OAM in the far field. Given that the primary source generation mechanism for OAM has been identified as being local stall on the blades, it is the case that any measures that prevent or reduce the onset of such local stall would result in the mitigation of OAM in the far field.

In the immediate term, the only guaranteed solution to mitigate fully against specific occurrences of OAM is the cessation of operation of offending turbines during those conditions under which problematic OAM is found to occur. These conditions leading to OAM, and the characteristics of that OAM when it occurs, appear to be very site-specific and would therefore need to be established specifically for each operational site considered. The use of the detection and rating methods identified, such as the main AM metric routine described in this report, could provide the basis for analysing these conditions based on the analysis of long-term monitoring.

It has been concluded that the effective mitigation of OAM in practice will require the future involvement and close cooperation of wind turbine manufacturers, and possibly involve detailed measurements that focus on better understanding the surface pressure distributions on the turbine blades themselves, particularly as the stall point is approached. Simple analysis methods have been identified in this project to assist in identifying the most likely relevant conditions. It is recognised that wind turbine design and manufacture are rapidly-advancing fields, the market is highly-competitive, and manufacturers must be protective of technical information relating to the design, manufacture and operation of their turbines. However, it is strongly recommended that wider collaboration is undertaken to devise effective mitigation methods.

It is believed that with such cooperation, methods will be capable of being developed for avoiding local stall conditions. Such methods may involve software 'fixes' that seek to modify the logic of the control system algorithms, perhaps even through the application of more advanced cyclical pitch control. More fundamental, physical design changes may also prove worthwhile, such as innovative blade designs or the addition of blade vortex generators, for example, as the latter may delay the onset of stall. Such mitigation would be likely to only have a limited or negligible impact on the generating capacity of the turbines.

It may also be concluded that there is nothing at the planning stage that can presently be used to indicate a positive likelihood of OAM occurring at any given proposed wind farm site, based either on its general characteristics or on the known characteristics of the wind turbines to be installed.

GLOSSARY OF TERMINOLOGY

TERMINOLOGY	DESCRIPTION
A-weighting	a filter that attenuates low frequency and high frequency sound to better represent the frequency response of the human ear when assessing the likely impact of noise on humans
acoustic character	one or more distinctive features of a sound (e.g. tones, whines, whistles, impulses) that set it apart from the background noise against which it is being judged, possibly leading to a greater subjective impact than the level of the sound alone might suggest
acoustic screening	the presence of a solid barrier (natural landform or manmade) between a source of sound and a receiver that interrupts the direct line of sight between the two, thus reducing the sound level at the receiver compared to that in the absence of the barrier
ambient noise	All-encompassing noise associated with a given environment, usually a composite of sounds from many sources both far and near, often with no particular sound being dominant
annoyance	a feeling of displeasure, in this case evoked by noise
attenuation	the reduction in level of a sound between the source and a receiver due to any combination of effects including: distance, atmospheric absorption, acoustic screening, the presence of a building façade, etc.
audible sound	a sound that can be heard above all other ambient sounds
audio frequency	any frequency of a sound wave that lies within the frequency limits of audibility of a healthy human ear, generally accepted as being from 20 Hz to 20,000 Hz. The frequency represents the actual rate at which the acoustic pressure is oscillating. (see *frequency* below)
background noise	the noise level rarely fallen below in any given location over any given time period, often classed according to daytime, evening or night-time periods (for the majority of the population of the UK the lower limiting noise level is usually controlled by noise emanating from distant road, rail or air traffic)
dB	abbreviation for 'decibel'
dB(A)	abbreviation for the decibel level of a sound that has been A-weighted
decibel	the unit normally employed to measure the magnitude of sound
directivity	the property of a sound source that causes more sound to be radiated in one direction than another
equivalent continuous sound pressure level	the steady sound level which has the same energy as a time varying sound signal when averaged over the same time interval, T, denoted by $L_{Aeq,T}$
external noise level	the noise level, in decibels, measured outside a building
filter	a device for separating components of an acoustic signal on the basis of their frequencies
frequency	the number of fluctuations per second. Note that this can apply either to: • Audio frequency: fluctuations in acoustic pressure about the atmospheric mean pressure (also known as the 'pitch' of a sound) • Modulation frequency: this represents the rate of repetition of the modulation in time when periodic or quasi-periodic • Sampling frequency: the number of times per second a particular signal will be measured. For example, for levels measured every 100 ms this will be 10Hz. Audio recordings are often made at sampling rates of 22 to 44 kHz.
frequency analysis	the analysis of a sound into its frequency components

TERMINOLOGY	DESCRIPTION
ground effects	the modification of sound at a receiver location due to the interaction of the sound wave with the ground along its propagation path from source to receiver
hertz	the unit normally employed to measure the frequency of a sound, equal to cycles per second of acoustic pressure fluctuations about the atmospheric mean pressure
impulsive sound	a sound having all its energy concentrated in a very short time period
instantaneous sound pressure	at a given point in space and at a given instant in time, the difference between the instantaneous pressure and the mean atmospheric pressure
internal noise level	the noise level, in decibels, measured inside a building
L_{Aeq}	the abbreviation of the A-weighted equivalent continuous sound pressure level
L_{A10}	the abbreviation of the 90 percentile noise indicator, often used for the measurement of road traffic noise (exceeded 10% of the time)
L_{A90}	the abbreviation of the 10 percentile noise indicator, often used for the measurement of background noise (exceeded 90% of the time)
level	the general term used to describe a sound once it has been converted into decibels
loudness	the attribute of human auditory response in which sound may be ordered on a subjective scale that typically extends from barely audible to painfully loud
masking	the effect whereby an otherwise audible sound is made inaudible by the presence of other sounds
modulation (amplitude)	for noise this characterises a change in amplitude (or perceived loudness) over time (i.e. going up and down). Over short periods of time these amplitude modulations may repeat themselves with an almost constant period, thus resulting in what is termed 'quasi-periodic' noise.
modulation frequency	The rate of repetition of the above modulation in time when it is periodic or quasi-periodic, in number of periods per second.
noise	physically: a regular and ordered oscillation of air molecules that travels away from the source of vibration and creates fluctuating positive and negative acoustic pressure above and below atmospheric pressure. Subjectively: sound that evokes a feeling of displeasure in the environment in which it is heard, and is therefore unwelcomed by the receiver
noise emission	the noise emitted by a source of sound
noise immission	the noise to which a receiver is exposed
octave band frequency analysis	a frequency analysis using a filter that is an octave wide (the upper limit of the filter's frequency band is exactly twice that of its lower frequency limit)
percentile exceeded sound level	the noise level exceeded for n% of the time over a given time period, T, denoted by $L_{An,T}$
pitch	for a wind turbine, this denotes the orientation of its blades around their length-wise axis
pitch-regulated wind turbine	a wind turbine that sheds power at the highest wind speed to regulate its output by dynamically changing the orientation of its blades around their length-wise axis
receiver	a person or property exposed to the noise being considered
residual noise	the ambient noise that remains in the absence of the specific noise whose impact is being assessed

TERMINOLOGY	DESCRIPTION
sound	physically: a regular and ordered oscillation of air molecules that travels away from the source of vibration and creates fluctuating positive and negative acoustic pressure above and below atmospheric pressure subjectively: the sensation of hearing excited by the acoustic oscillations described above (see also 'noise')
sound level meter	an instrument for measuring sound pressure level
sound pressure amplitude	the root mean square of the amplitude of the acoustic pressure fluctuations in a sound wave around the atmospheric mean pressure, usually measured in Pascals (Pa)
sound pressure level	a measure of the sound pressure at a point, in decibels
sound power level	the total sound power radiated by a source, in decibels
spectrum	a description of the amplitude of a sound as a function of frequency
stall	When the flow around an airfoil (or turbine blade) is interrupted or becomes detached from the surface of the airfoil, resulting in a loss of lift
stall-regulated wind turbine	a wind turbine that sheds power at the highest wind speed to regulate its output by allowing its blade to stall
third-octave band frequency analysis	a frequency analysis using frequency bands one third of an octave wide
threshold of hearing	the lowest amplitude sound capable of evoking the sensation of hearing in the average healthy human ear (0.00002 Pa)
tone	the concentration of acoustic energy into a very narrow frequency range

REFERENCES

Throughout this report, references to [WPX] refer to the report for Work Package X of this project, where X stands for the following:

WP	Description
A1	Source generation effects modeling
A2	Fundamental Research into Possible Causes of Amplitude Modulation
B1	Development of an Objective AM Measurement Methodology
B2	Development of an AM Dose-Response Relationship
C	Collation and Analysis of Existing Acoustic Recordings
D	Measurement and Analysis of New Acoustic Recordings

Additional references are as follows:

[1] Research into aerodynamic modulation of wind turbine noise, Report by University of Salford for Department for Business, Entreprise and Regulatory Reform, URN 07/1235, July 2007, Contract NANR233.

[2] The Measurement of Low Frequency Noise at Three UK Wind Farms, Hayes McKenzie Partnership for the DTI, 01/07/2006.

[3] ETSU-R-97, the Assessment and Rating of Noise from Wind Farms, Final ETSU-R-97 Report for the Department of Trade & Industry. UK Noise Working Group, 1997.

[4] F. Van den Berg, Effects of the Wind Profile at Night on Wind Turbine Sound", J. Sound Vib. 277 (4-5), 955-970 (2004).

[5] S. *Oerlemans 2009. Detection of aeroacoustic sound sources on Aircraft and Wind Turbines.* PhD Thesis, University of Twerte.

[6] Stefan Oerlemans and Gerard Schepers, Prediction of Wind Turbine Noise Directivity and Swish, Third International Meeting on Wind Turbine Noise, Aalborg Denmark 17 – 19 June 2009.

[7] Fastl, H. and Zwicker, E., (2007) Psychoacoustics, Facts and Models. Springer, Berlin, Germany,

[8] Renewable Energy Foundation website, http://www.ref.org.uk, 06 February 2009.

[9] Wind Turbine Noise, D. Bowdler and G. Leventhall (Eds.), Multi-Science Publishing, 2011.

[10] Time Domain Modelling of Aerodynamic Noise from Wind Turbines, Lee et.al., Fourth International Meeting on Wind Turbine Noise, Rome, April 2011.

[11] Carlo Di Napoli, Case Study: Wind Turbine Noise in a small and quiet community in Finland, Third International Meeting on Wind Turbine Noise Aalborg Denmark 17 – 19 June 2009.

[12] Carlo Di Napoli, Long Distance Amplitude Modulation of Wind Turbine Noise, Fourth International Meeting on Wind Turbine Noise, Rome, Italy, April 2011.

[13] Palmer WTN 2009, A New Explanation for Wind Turbine Whoosh – Wind Shear, Third International Meeting on Wind Turbine Noise Aalborg Denmark 17 – 19 June 2009.

[14] K. Boorsma, J.G. Schepers, Enhanced Wind Turbine Noise Prediction Tool SILANT, Fourth International Meeting on Wind Turbine Noise, Rome, Italy, April 2011.

[15] J. Bass, "Investigation of the 'Den Brook' Amplitude Modulation Methodology for Wind Turbine Noise," Institute of Acoustics Bulletin, vol. 36, p. 7, 2011.

[16] The Den Brook Amplitude Modulation Noise Condition, Dr Lee Moroney, Dr John Constable, Renewable Energy Foundation, published www.ref.org.uk

[17] J. Vos and M. M. J. Houben, "Analyse van opnamen van windturbinegeluid," TNO TNO-DV 2010 C014, February 2010 2010.

[18] Moorhouse, A., Waddington, D. and Adams, M., (2005), 'Procedure for the assessment of low frequency noise complaints', DEFRA Contract No NANR45, null, pp.24.

[19] Lee, S., Kim, K., Lee, S., Kim, H., and Lee, S., (2009) An estimation method of the amplitude modulation in wind turbine noise for community response assessment, Third International Meeting on Wind Turbine Noise, Aalborg, DK, June 2009.

ANNEX A - A SUMMARY OF THE STATE OF TECHNICAL KNOWLEDGE

This Annex summarises how the current state of technical knowledge relating to AM has developed, largely as a result of the ReUK funded investigations that have been ongoing over the course of the past year.

The text may assist in making decisions as to how the issue of AM is treated in planning situations, particularly with regards the imposition of AM related noise conditions.

This note is set out on the basis of a direct comparison between the situation that generally existed twelve months or more ago, with a concurrent comparison as to the situation that exists now.

PAST KNOWLEDGE	PRESENT KNOWLEDGE
What is AM?	
• *AM refers to the 'Aerodynamic Modulation' or the 'Amplitude Modulation' of wind turbine blade noise* • *AM can be divided into 'normal AM', or 'blade swish', and 'enhanced AM'*	• *AM more correctly refers to the 'Amplitude Modulation' of wind turbine blade noise which is aerodynamic in its origin* • *AM can be divided into 'normal AM', or 'blade swish', and 'other AM'*
What are the characteristic features of 'Blade Swish' or 'Normal AM'?	
Blade swish, or 'Normal AM': • *an inherent feature of wind turbine noise;* • *typically 3 dB modulation depth close to source;* • *decreases with increasing distance, although;* • *its presence may be exacerbated by local reflections, even at larger distances, with modulation depths of up to 6 dB;* • *dominated by 800 Hz to 1000Hz frequencies;* • *acknowledged and accounted for in ETSU-R-97 in the above terms;* • *since the publication of ETSU-R-97 it has become well understood in terms of its aerodynamic source generation mechanisms.*	*Some conflicting evidence that the frequencies dominating the modulation have decreased as the size of the turbine blades has increased.* *Theoretical models have been developed which are well-validated in the near- to mid-field.*
What are the characteristic features of 'enhanced AM (EAM)' or 'other AM (OAM)'?	
***EAM** typically described as:* • *an atypical feature of wind turbine noise;* • *any AM falling outside the presumptions relating to blade swish in ETSU-R-97 (i.e. greater than 3 dB peak to trough free field, or 6 dB peak to trough in the presence of reflections);* • *no systematic changes to the frequency characteristics set out in ETSU-R-97 for blade swish;*	***OAM** more clearly identified as:* • *an atypical feature of wind turbine noise;* • *audible at distances close to or in excess of 1000 m;* • *even under free-field conditions, modulation depths may exceed 5 dB at these larger distances, with modulation depths of up to 10 dB in overall A-weighted levels being reported;* • *dominated by lower frequencies than 'normal AM', typically around 300-400 Hz ;* • *modulation depths in the dominant lower frequency band generally exceed the modulation depths reported in overall A-weighted levels;* • *subjectively more impulsive in character than 'normal AM';* • *the foregoing features generally result in OAM being described as more of a 'whoomph' than a 'swish' sound;*

<table>
<tr><th>PAST KNOWLEDGE</th><th>PRESENT KNOWLEDGE</th></tr>
<tr><td colspan="2">What is the source generation mechanism of EAM or OAM?</td></tr>
<tr><td>Source generation mechanism of EAM:
• unknown, but various potential generation mechanisms or factors were speculated on.</td><td>Source generation mechanism of OAM:
• still not fully proven, but local transient stall on the rotor blades is most likely to be the predominant source generation mechanism.
The ReUK project has additionally identified that the following (non-source related) mechanisms may add to the effects of OAM:
• propagation effects may increase levels of OAM experienced at larger distances (particularly upwind);
• local conditions at receptor (particularly internal room effects) may enhance noise across common frequency range of OAM;
• masking effects of background noise at receptor location may be a key determining factor in measured modulation depths (this may add to the explanation as to why OAM is sometimes reported as more prevalent under high wind shear conditions).</td></tr>
<tr><td colspan="2">How Common is AM?</td></tr>
<tr><td>'Normal AM', or blade swish:
• an inherent characteristic feature of all wind turbines.</td><td>'Normal AM', or blade swish:
• an inherent characteristic feature of all wind turbines.</td></tr>
<tr><td>EAM:
• reliance placed on HMP LFN study and subsequent Salford study to indicate low prevalence across the UK Wind Farm fleet;
however,
• once identified as a potential feature, increasing reports by objector groups of EAM causing problems at other wind farms;
• high profile of Deeping St Nicholas case raised general awareness still further;
• increasing public availability of reports making a direct link between EAM (amongst other features of wind farm noise) and adverse subjective responses</td><td>OAM, generally still as per 'past knowledge' comments, as no additional systematic data on prevalence has been collected, but additionally:
• seen to be acknowledged by the wind energy industry as a 'key issue through the letting of the ReUK AM Research Project;
• publicly reported as being 'a small problem' but now 'too large to ignore';
• even on those limited sites where it has been reported, its frequency of occurrence appears to be at best infrequent and intermittent.</td></tr>
<tr><td colspan="2">What causes EAM/OAM?</td></tr>
<tr><td>Possible causes variously presented by different authors as:
• blade passing tower
• angle of attack changes
• high shear/stable atmosphere
• high turbulence (or tip vortex)
• yaw error
• rotor/wake effects
• interaction between turbines
• sync between turbines or phasing
• propagation effect
• 'stubby' towers
• etc.</td><td>Based on reported occurrences of OAM, and the focus on the source generation mechanism being localised transient stall, it may be concluded that:
• many of the possible causes listed opposite could be possible contributory factors;
• some of these factors were clearly established as not being associated with incidence or magnitude of OAM;
• if any one factor or combination of factors were solely responsible, then OAM would occur at all sites featuring those factors and it would occur frequently at those sites – neither of which has been observed to be the case in practice;
• as an example, OAM has been observed to occur on a single isolated turbine with a high tower, so interaction effects between linear arrays and too closely spaced turbines cannot</td></tr>
</table>

<table>
<tr><th>PAST KNOWLEDGE</th><th>PRESENT KNOWLEDGE</th></tr>
<tr><td></td><td>be a sufficient factor, neither can a large rotor on a stubby tower;
• as another example, OAM has been reported for one make and physical configuration of turbine on one site, whereas it has not been reported for the same make and physical configuration of turbine on all installations;
• the interaction between the various site/turbine characteristics listed opposite coupled with the operational characteristics of the turbine control system and features of the specific blade design may lead to OAM occurring on some occasions at some sites.</td></tr>
<tr><td colspan="2">Is EAM/OAM likely to be a feature at any particular site?</td></tr>
<tr><td>Based on the foregoing speculative causes it was generally proposed that EAM would be more likely to occur on:
• sites that exhibited high wind shear due to stable atmospheric conditions;
• flat sites on the east coast;
• sites where turbines were spaced too closely;
• sites where turbines were arranged in linear arrays;
• turbines with large rotors on relatively short towers.
• Etc.</td><td>Based on the current understanding:
• it is not possible to be prescriptive as to whether any particular site is more or less likely to give rise to OAM being generated at source;
• Several of the potential causal factors which have been suggested to date were shown to have little or no association to the occurrence of OAM: for example interaction between closely spaced turbines in linear arrays;
• However some of these features may represent potential contributory factors;
• it is only where background noise at noise sensitive receptor locations is high enough under all circumstances to effectively mask any potential OAM noise that it can be positively concluded that the chances of OAM causing noise issues are reduced.</td></tr>
<tr><td colspan="2">Is there an accepted objective metric with associated dose-response relationship for AM?</td></tr>
<tr><td>No, but:
• a seemingly objective metric has been proposed by MAS Environmental (as implemented at Den Brook Wind Farm);
• the proposed metric is essentially based on observing whether the peak to trough modulation depth in the overall A-weighted level exceeds 3 dB under free field conditions (believed to be selected as this because this what is stated to be the expected maximum in ETSU-R-97);
• the foregoing metric was considered to be highly susceptible to the detection of 'false positives', even in the absence of any wind farm noise being present and quite apart from any EAM being present;
• the foregoing metric has not been substantiated in any way in terms of any dose-response relationship.</td><td>No, but:
• the metric referred to opposite has been demonstrated to be highly susceptible to the detection of 'false positives', even in the absence of any wind farm noise being present and quite apart from any OAM being present, effectively requiring difficult subjective consideration in its application;
• objective metric(s) methods which can effectively detect modulated wind turbine noise with a sufficiently reduced amount of 'false positives' have now been developed as part of this project;
• the dose-response relationship for the aforementioned metric(s) was found to be complex, although a significant effect due to the presence of certain levels of modulation was demonstrated, although there was relatively little additional effect beyond a certain level of modulation.</td></tr>
</table>

<table>
<tr><th>PAST KNOWLEDGE</th><th>PRESENT KNOWLEDGE</th></tr>
<tr><td colspan="2">Can an objective planning condition presently be drawn up for EAM/OAM?</td></tr>
<tr><td>Yes, but:
• this was not considered necessary given the recognized low probability of occurrence;
• whether or not this would survive the tests for a valid condition was debatable;</td><td>Yes, as listed opposite, but also:
• arguments concerning necessity are finely balanced;
• tonality conditions are retained as standard even though it is argued that tones are not expected to occur, although gearboxes can fail and there is then a clear method for assessing the effects; however,
• current evidence suggests that some of the earliest proposed conditions may fail the tests of precision and reasonableness;
• although more objective and reliable methods have been considered as part of this project, there is still limited experience of their practical application and associated consequence;
• It is possible that OAM mitigation measures may be developed based solely on the operational characteristics of the affected turbines (i.e. modifications to the turbine control systems, and in particular the blade pitch control systems) or through physical modifications to the turbines themselves, as opposed to operational curtailment of turbines; however,
• The lack of definite technical experience in this regard, in the current state of the art, represents a key uncertainty. This may therefore necessitate operational curtailments in the meantime;
• ice throw from blades, which is unusual and part of technical design and is not conditioned, can be considered for parallels;
• it may still be argued that the known incidence of occurrence and frequency may not satisfy the necessity argument .</td></tr>
</table>

ANNEX B – EVALUATION OF ADDITIONAL PROPAGATION EFFECTS ON AM MODEL RESULTS

The model presented in WPA1 does not take into account the longer-distance propagation effects, described in WPA2 or Chapter 3 of [9], which are known to affect the long-range propagation of noise over larger distances. The model was effectively validated at distances of 3 RD [6], but not at further distances which are relevant to the far-field region in which wind farm neighbours can generally be found. Over these distances, the potentially most significant propagation effects are:

- spherical spreading of sound (which will reduce overall noise levels equally in all directions and is not considered further);
- refraction effects due to wind speed gradients in upwind/downwind;
- atmospheric absorption effects at high frequencies.

The effects of the foregoing factors were accounted for to some extent in the similar model developed by Boorsma *et. al.* [14], but, as flow separation noise was not considered, we can instead take the results of WPA1 and try to estimate potential influence of the additional effects identified in WPA2.

Dominant frequencies and atmospheric absorption

The ETSU-R-97 report [3] described blade swish as dominated by frequencies at and above the 800 – 1000 Hz range, which was typical of turbines at this time[19]. As turbine designs have evolved, and have increased in size, it can be interesting to consider how this may have affected the characteristics of the amplitude modulated noise from the turbines.

The research undertaken by Delta as part of a large research project for Danish Energy Authority[20] analysed a large variety of tested sound power spectra for turbines of different size categories (both less and more than 2MW generating capacity), to determine if certain types of turbines produced additional levels of low-frequency noise. It concluded that the noise emission spectra for the larger category of turbines, as averaged across all turbines tested, was only marginally[21] more dominated by the frequencies in the range between 100 and 200Hz than the smaller category of machines. However, the observed difference was less than the variability found between turbines falling within the same size categories, such that a certain small turbine may have larger relative low frequency components than a larger and vice versa. This was characterised as a potentially noticeable but not an essential change.

Whilst this relates to the overall tested[22] sound power results, we can consider further how the evolution of turbine designs may be considered in standard sound emission models of the modulated part of the turbine noise. With the increase of the size of turbine rotor diameters, the blade chord dimensions tend to increase (which will increase the thickness of the turbulent boundary layer over the aerofoil) whilst tip speeds may remain similar by design. This effectively results in a decrease in the ratio between the flow velocity and the boundary layer displacement thickness (and therefore the Strouhal number[23]), which will tend to reduce the dominant frequency in the modulation spectrum. The figures A1(a) and (b) illustrate this potential effect based on the observations made for two different turbines types and the comparison of model predictions and experimental results. The peak frequency of the spectra for the larger turbine have approximately halved. The results of the more general study by Delta suggest this may be affected by a degree of variability, due to differing designs, which means the comparison between the 2 turbines may not necessarily be of general significance.

[19] Wind turbine measurements for Noise source identification - ETSU W/13/00391/00/REP - Dr A J Bullmore, Mr J F Lowson, Dr J H Bass, Dr P Dunbabin, 1999.
[20] Delta Acoustics for the Danish Energy Authority, EFP-06 Project - Low Frequency Noise from Large Wind Turbines, 2010.
[21] By approximately 2 dB in relative terms. In the range 200-500Hz this reduced to approximately 1dB.
[22] In accordance with the IEC 61400-11 standard
[23] $S_t = f\,\delta^*/U$, where δ^* represents the boundary layer displacement thickness, and U the flow speed.

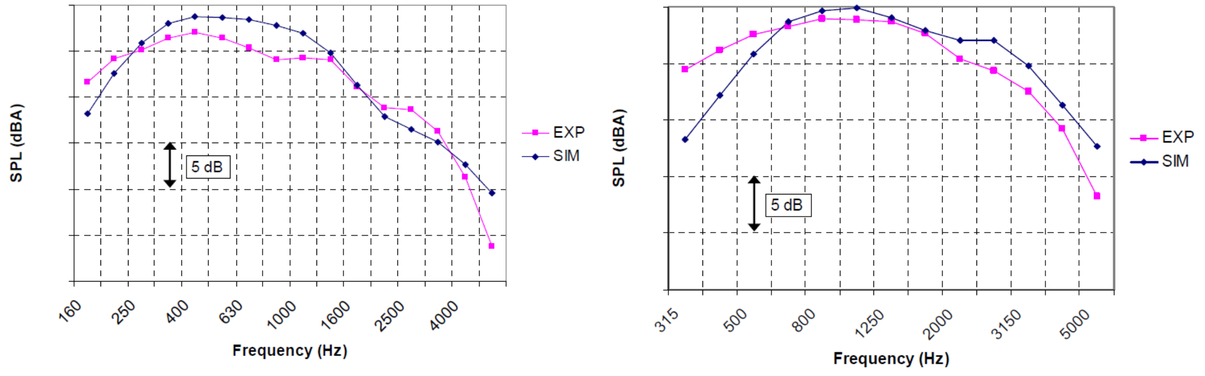

(a) 94m rotor diameter turbine (b) 58m rotor diameter turbine

Figure A1 - Predicted and measured spectra for two different turbine models (from [5]).

This would be consistent with the observations of Legarth[24,25], which noted that the most relevant frequencies for the modulation was in the 350-700Hz band, in comparison with frequencies centred around 1kHz in a 1989 study. This was based on a study of recordings of turbines of rotor diameters of 30 to 80 m, made at distances of between 1.5 and 3 hub heights from the turbines.

The effects of atmospheric absorption are highly frequency dependent, and higher frequency sound experiences greater atmospheric attenuation than lower frequency sound (with a rapid increase above 1kHz). For modulation spectra dominated by the 500Hz region, the effect of atmospheric absorption will therefore not be as significant as when dominated by the 1kHz region.

WPA1 notes that, for the situation of a detached flow in the downwind direction, the peaks of the modulation are dominated by the stall noise, whereas the troughs are dominated by the trailing edge noise. This effect is illustrated in Figure 25 of the WPA1 report, reproduced in Figure A2 below. As the periods of stall have increased low-frequency content, it is natural to consider whether atmospheric absorption effects may enhance the difference between peaks and troughs and therefore the levels of modulation.

Figure A2 shows the same spectra after applying a correction for these effects based on standard data[26] presented for conditions of 10 degrees Centigrade and 70% humidity (corresponding to relatively low absorption) for a propagation over a distance of the order of 10 RD (say 800 m). It can be seen that as the A-weighted spectra in this case is dominated by frequencies close to 500 Hz, and the relative difference between the different instantaneous spectra is not significantly affected. The main difference between the periods of stall and non-stall in the modulation, the higher relative prominence of the spectra below frequencies of 200 Hz (of the order of 5 dB), is only marginally increased, but this would not tend to affect the overall A-weighted levels to a noticeable degree.

The effect of air absorption on the spectra of turbines dominated by frequencies closer to 1 kHz would be more significant.

[24] Auralization and assessments of annoyance from wind turbines, Wind Turbine Noise 2007 Conference, Lyon, France.

[25] Methods for assessment of the characteristics of wind turbine noise, T. H. Pedersen, S. Von-Hunerbein, S. Legarth, Fourth International Meeting on Wind Turbine Noise, Rome, Italy, April 2011.

[26] ISO 9613-1 Acoustics – Attenuation of sound during propagation outdoors, part 1: Calculation of the absorption of sound by the atmosphere.

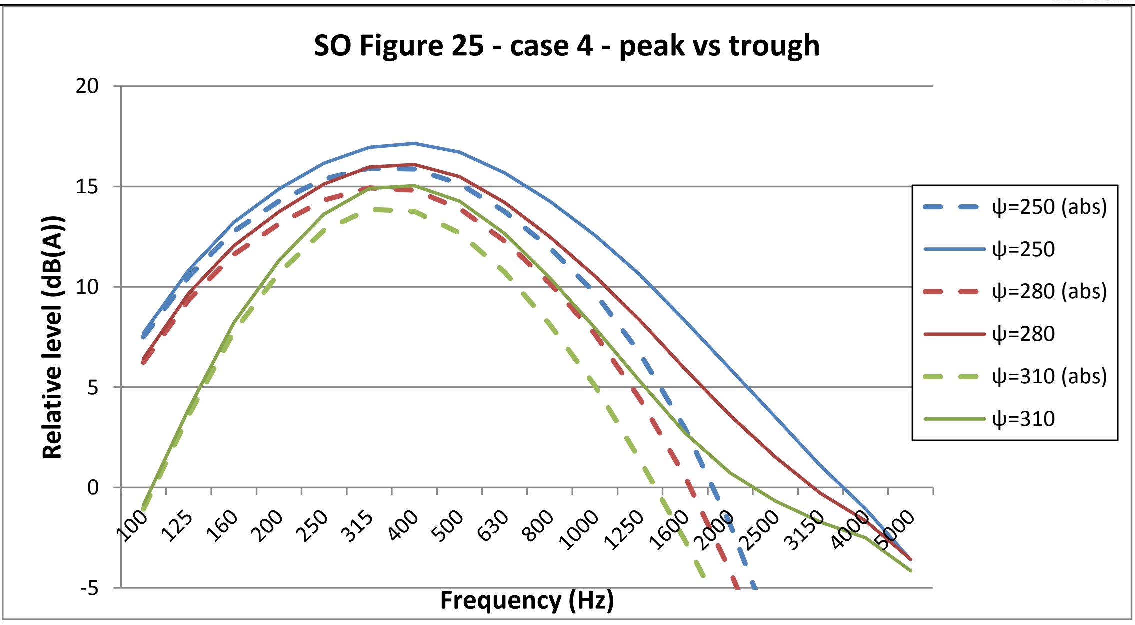

Figure A2 – Instantaneous spectra for 'case 4' as studied in WPA1 (detached flow, m=0.3), modelled in the downwind direction 10 RD away. ***Dotted lines show modified spectra accounting for air absorption.*** *ψ = 250 corresponds to the peak of the modulation (stall) and ψ = 310 to the trough.*

<u>Additional considerations – refraction effects</u>

We can extend the results of the WPA1 model, which effectively assumed a neutral atmosphere (for propagation purposes), by accounting for far-field refraction effects occurring in different wind directions.

It is well-known (see WPA2] or Chapter 3 of [9]) that the wind speed gradients[27] present in downwind conditions will correspond to favourable propagation conditions (Figure A3), whilst inverse gradients found in upwind conditions will represent unfavourable conditions. These effects tend to become significant in the far-field region only, between 5 and 15 hub heights[28] .

For broadband signals, the former favourable conditions will represent an increase up to around 3dB relative to neutral conditions; in upwind conditions, measurement studies have shown that levels can be 10 dB or more lower than the received noise levels under neutral conditions. It is also generally accepted that such downwind conditions may effectively be present with a wind speed vector component of 2 m/s from source to receiver, which may occur from angles of 10 degrees from downwind propagation.

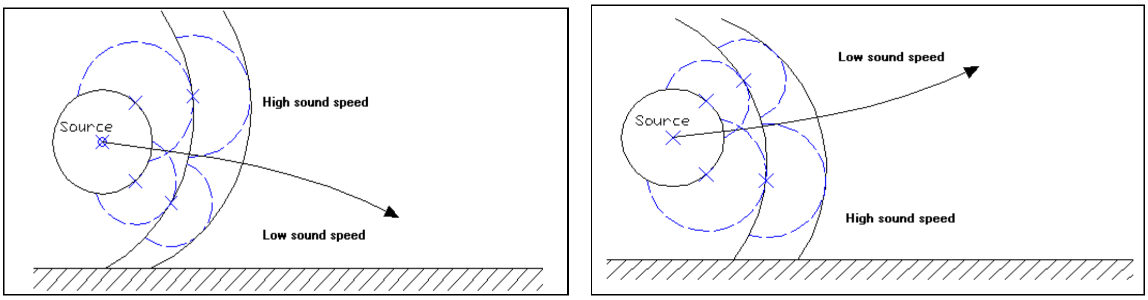

(a) downwind conditions (b) upwind conditions

Figure A3: sound curvature under different wind speed gradients

[27] Although temperature gradients (under temperature inversion and lapses) can results in comparable sound speed gradients, wind-related effects are generally dominant for situations relevant to wind turbine noise.
[28] Shepherd, K. P.; Hubbard, H. H. Wind Turbine Acoustics, NASA Technical Paper (TP) 3057, Dec 1990.

These effects can be (simplistically) modelled through a noise propagation directionality function, with effects ranging from +3 dB downwind of a turbine to -10 dB upwind. These corrections can be applied to the far-field corrected directivities presented in WPA1: this is done for the 'case 2' and 'case 4' studied, as presented in Figures 17 and 19 of the WPA1 report, which correspond to a case of moderately high wind shear (m=0.3), both without and with detached flow respectively. In Figure A4, the calculated directivity function for the overall noise levels is shown for both cases as shown in the report (dotted line); in addition, solid lines show a corrected directivity function which accounts for the additional effects of refraction for the far-field directivity, based on the model described above.

It can be seen that, once the effect of the upwards refraction of the sound and the associated 'shadow zone' in upwind conditions are taken into account in addition to the intrinsic directivity of the source, that: outside of a zone of +/- 45 degrees from downwind of the turbines, the overall level of noise experienced in the far field decreases significantly. This means that this direction is key when considering noise emissions from turbines, as it will tend to be masked by other sources as the overall levels reduce.

However it should be noted that the model used of reduction in upwind conditions is crude and that, as noted in WPA, the effects of source height and turbulence could have significantly more complex effects. A clear summary of the situation highlighted in WPA2 is presented in [9], page 126:

'There are some propagation effects that theory suggests may not only moderate AM …. but might actually cause it. One of these is an upwind effect. The reason sound is not heard over long distances upwind is because of the effect of wind shear. This has the effect of bending sound waves upwards and away from the ground so that beyond some distance there is a sharp reduction of sound level. ….. However, the higher the noise source, the greater the distance at which it can be heard. So with turbine noise it is possible that, at say 600 metres distance, the sound of the turbine blade at the top of its trajectory could be heard, but not the sound at the bottom or half way down. So a form of AM could be generated in this way [note: the effects of refraction would also cause the relative contributions of the lower frequencies to be enhanced relative to the higher frequencies, which is one of the observed characteristics of OAM]. *It is also possible that at certain frequencies the ground effect (interference of the direct sound and the sound reflected at the ground) could reduce turbine noise at the bottom of the trajectory more than at the top or vice versa. This could cause a form of AM in certain frequency bands.*

It nevertheless seems significant (as highlighted in Figure A5 which overlays the modelled swish amplitude) that, in the case of detached flow, significant modulation occurs in the downwind sector. This is because of the significantly different directivity of the stall noise (a dipole in perpendicular to the rotor plane) compared to the trailing edge noise (a cardioid in the direction of blade rotation), as shown in Fig. 12 of the WPA1 report. This is particularly evident when comparing instantaneous noise footprints between the two cases (see Figure 2 in the main body of this report).

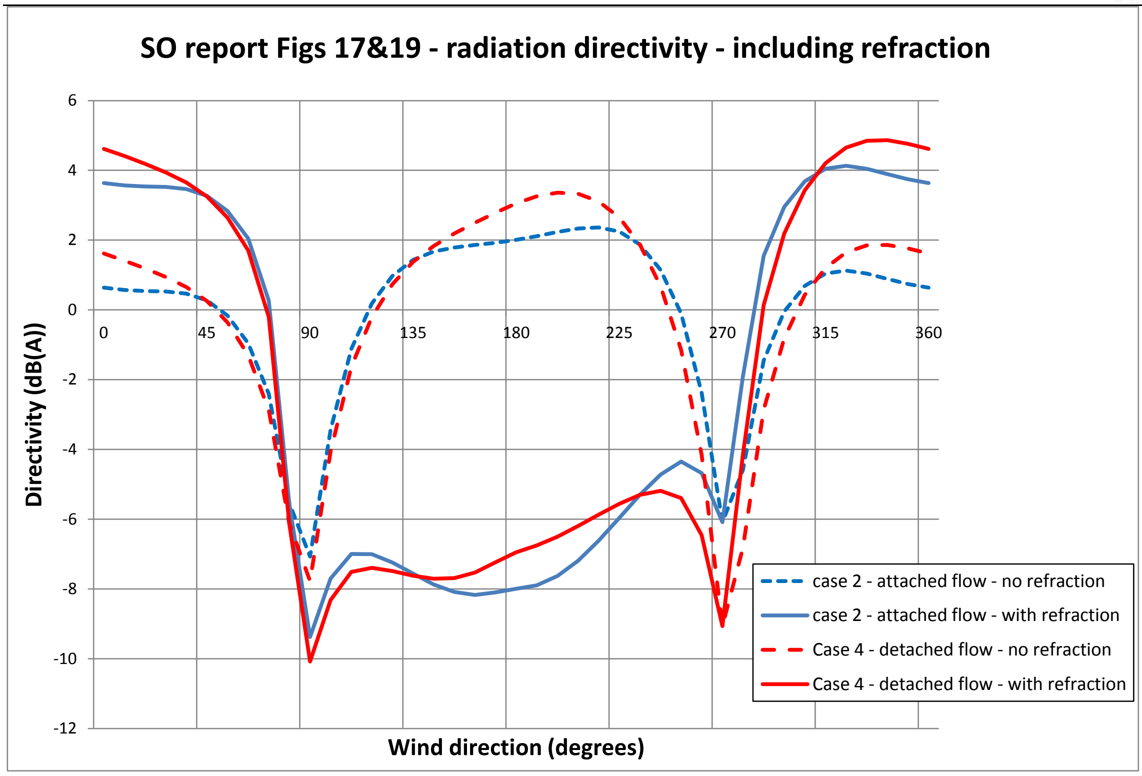

Figure A4 – Modelled trailing edge noise directivity from WPA1 (dotted lines), cases for m=0.3 with both detached and attached flow (dotted lines); corrected directivity accounting for far-field refraction propagation effects (straight lines)

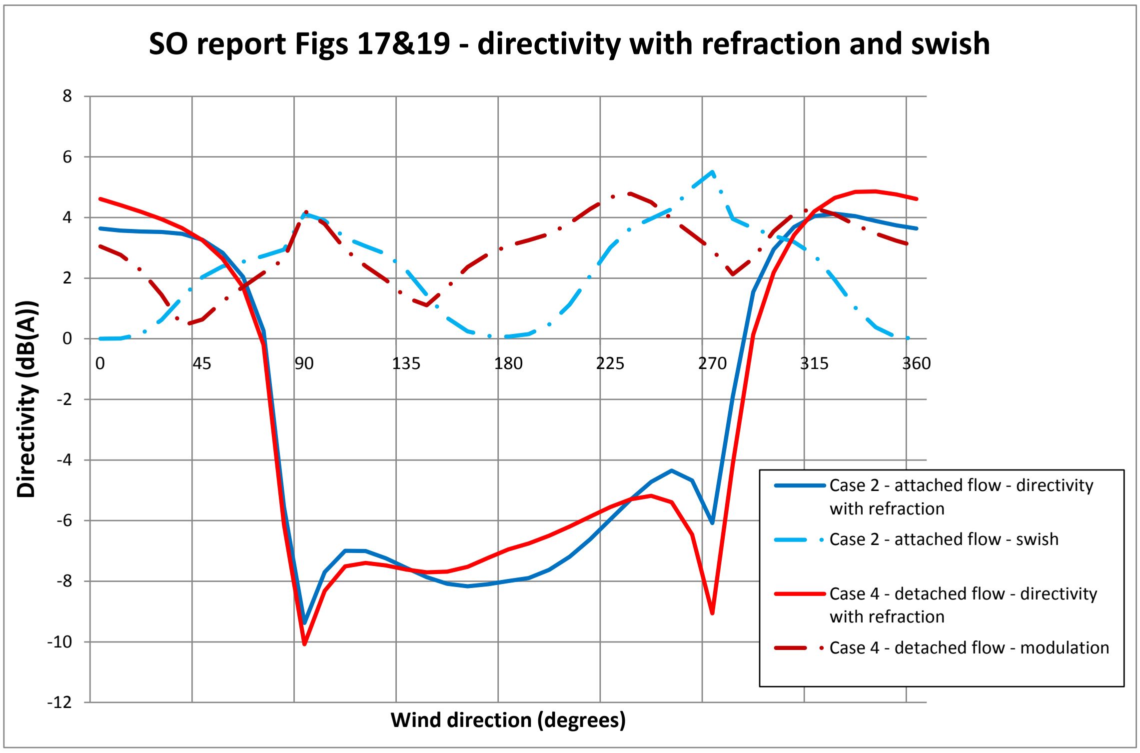

Figure A5 - Modelled trailing edge noise directivity from WPA1, corrected for directivity

ANNEX C – APPLICATION OF AN AM METRIC ROUTINE TO A RANGE OF SIGNALS

A principal AM metric routine has been developed and applied to a variety of signals. It is similar to the first category of methods described in WPB1, i.e. Fourier analysis of the acoustic signal envelope. This was based on an implementation done in MATLAB originally developed by RES. This method is described in section 4.3.

C.1 Application to simple test signals

This implementation has first been applied to a set of simplified test signals to assess the normalisation process used and the effect of different input parameters.

A simple sinusoidal signal is defined directly in terms of the $L_{Aeq,100ms}$ signal envelope directly, and the analysis applied to this signal. The modulation frequency f_m is fixed to 0.5Hz. The peak-to-trough modulation amplitude is defined when generating the signals. For this pure signal, one clear peak at f_m can be seen in the spectra: see Fig. C1. The amplitude of this peak in the modulation spectrum can be compared to that expected from the artificial signal: this is done for a range of modulation depth in Fig. C2. A good agreement is obtained both in terms of the absolute values obtained and the general trend observed.

When considering the 20s signal period, a standard Blackman-Harris windowing function was used to produce the results in a first instance. A rectangular window could be considered instead: as expected, the resulting spectrum has less accuracy in the spectral domain, with the apparition of side lobes: Fig. C3. But the peak at the modulation frequency in the raw Power Spectral Density (PSD) is finer than when using the Blackman-Harris window, and with a higher amplitude; however, the use of integration over the PSD over a frequency window (of width set to 10% of f_m) means that similar amplitudes are obtained for the amplitude of the integrated modulation spectra at f_m with either window function: see Fig. C2.

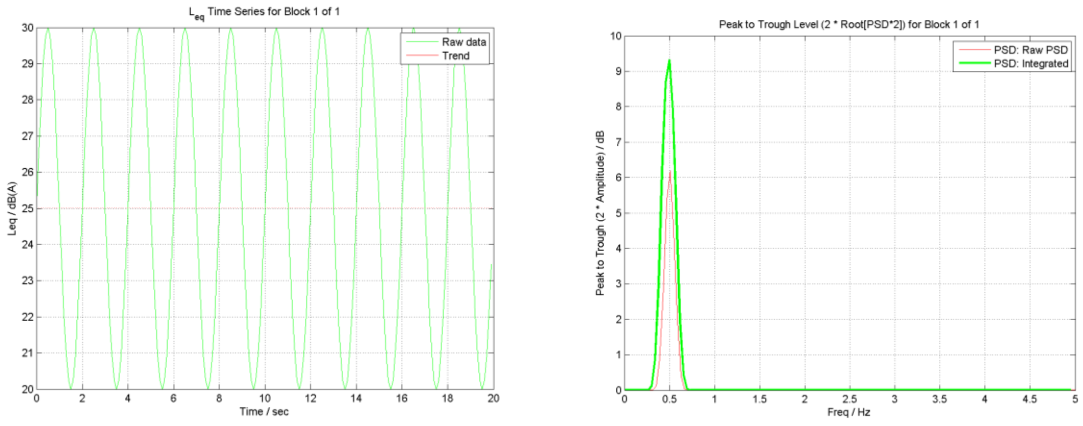

Figure C1 – Time history ($L_{Aeq100ms}$, left) and resulting modulation spectra (right) of artificial sinusoidal signal envelope (with Blackman-Harris window) – example of 10dB(A) modulation depth

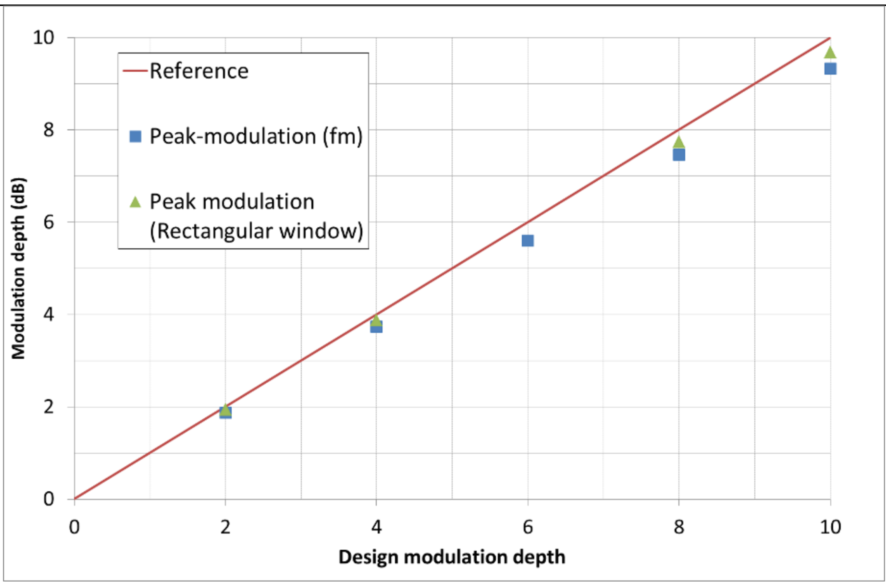

Figure C2 – Amplitude (dB) of peak modulation in spectra for different sinusoidal signals of defined modulation depth

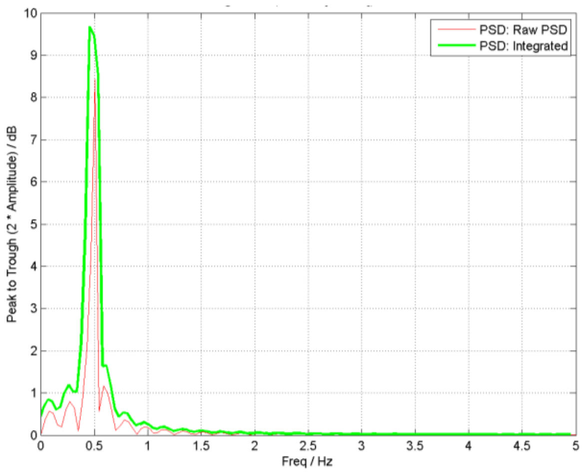

Figure C3 – Resulting modulation spectra for signal of Fig. C1, using rectangular windowing

For a stimuli of 10dB(A) modulation depth, similar integrated peak values are obtained for either windowing method, with differences of 0.3dB. In contrast, the peak of the raw PSD is 6dB for the Blackman-Harris window, and 8dB for the rectangular window. In effect, the spectrum integration process compensates for the diffusion of the reduction in frequency resolution introduces by the windowing effect, and therefore the significance of the windowing function used is reduced.

As discussed in WPB1, modulation signals might not necessarily be sinusoidal, and this introduces some additional considerations. This can be illustrated by considering the (unrealistic) extreme example of an ideal saw-tooth signal, modulating at a similar rate (f_m=0.5Hz): see Fig. C.4. The resulting modulation spectrum contains secondary peaks at harmonics of the main frequency f_m. WPB1 notes that, if we maximise a metric that includes the sum of the amplitudes of the modulation spectrum at this frequency (f_m) and integer multiples of it (k f_m), we can determine the fundamental frequency of the modulation of such signals in an optimal way. Whilst this may be a good way of identifying the modulation *frequency* in certain applications, this is less of a consideration for wind turbine noise applications in which the modulation rate is related to the rotation of the turbine(s) which is more likely to be well-known. This does not tell us what is the correct way (if any) of determining the modulation *amplitude* corresponding to these more complex signals.

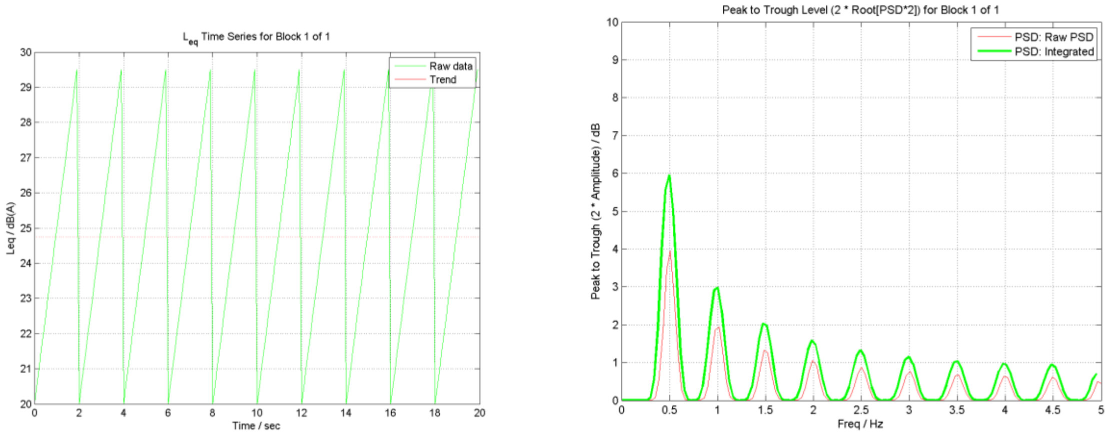

Figure C4 – Time history ($L_{Aeq100ms}$, left) and resulting modulation spectra (right) of artificial sinusoidal signal envelope (with Blackman-Harris window) – example of 10dB(A) modulation depth

We can examine what this means for the example of the saw-tooth modulation. It can be seen that the amplitude of the peak at the fundamental frequency in this case does not match the overall peak-to-trough amplitude of this modulated signal envelope. For this ideal signal, it can be shown that multiplying the first harmonic by a factor of π/2 results in a good estimate of this total amplitude. A reasonable approximation is obtained by adding the amplitude of the first two harmonics, although it can be seen that adding further harmonics (3 and more) will result in artificially high results: this is because this summation does not take into account the phase information of the spectral analysis. It should also be noted that for real signals (as will be seen below), high-order harmonics will become increasingly ‘swamped’ by noise. As above, the use of a rectangular window will result in a marginally higher value: for a design modulation depth of 10dB, after multiplication with π/2 factor, results in a total amplitude of 9.8dB instead of 9.4dB).

It should be noted that there is no reason why a metric which accounts for the harmonics in the signal would correspond better with the subjective response: it is equally conceivable that the main harmonic would have the most significant effect. Importantly, the consideration of any subjective response should be made consistently with the consideration of the metric normalisation, including the way signal harmonics are accounted for. We therefore next consider the artificial stimuli used in the work for WPB2.

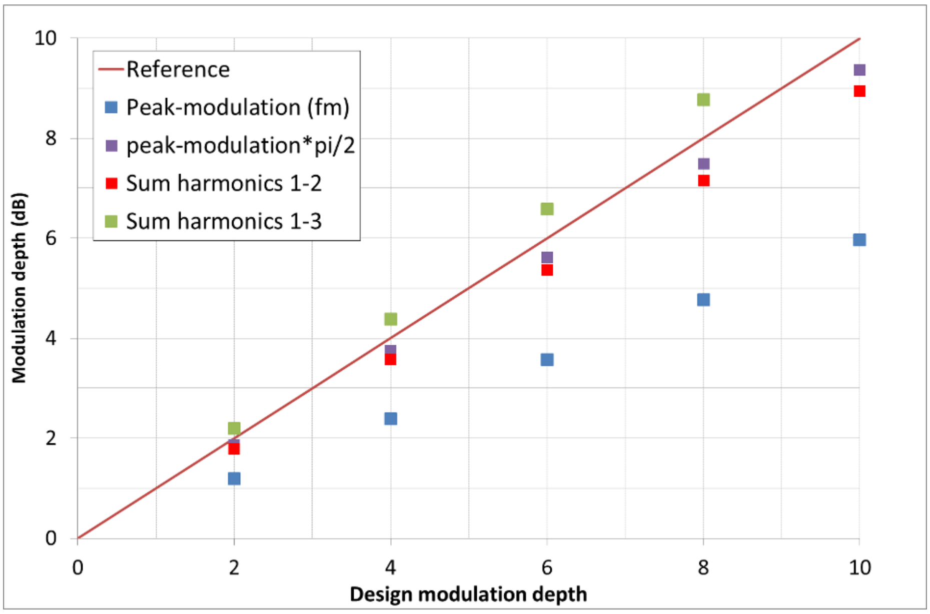

Figure C5 – Amplitude (dB) of peak modulation in spectra for different saw-tooth signals of defined modulation depth

C.2 Application to artificial stimuli developed for subjective testing under Work Package B2

The analysis method was applied to a representative selection of the artificial AM stimuli signals used in the study of subjective response as part of the current project, for the final tests described in Appendix V of WPB2. The artificially generated stimuli were played back in the listening room and recorded as audio signals to allow post-processing and analysis.

As discussed in section 17.4 of WPB2, the peak-to-trough variation in the short-term RMS levels of the signal was used as a guide to aid the design of a range of stimuli, with typical variations between 0 and 12 dB(A). This simple modulation depth metric is designated as 'MD' for clarity and consistency with WPB2. It was reasonably well-defined (through the use of averaged peak and trough levels) for these artificial signals, but this type of analysis becomes more difficult to apply in practice to realistic signals, and therefore the application of more precise analysis methods will be compared to the design metric MD.

The modulation peaks (with a defined frequency content) were overlaid on top of a masking signal (un-modulated 'wind turbine noise' in this instance). The spectrum of the modulation signal was centred on the 300Hz region, with a 180 Hz bandwidth and a slight low-frequency bias. The envelope of the pulses was also asymmetric in time, and assumed a Gaussian profile: the modulation was therefore effectively non-sinuisoidal: see Figure C6.

This translates in a clear modulation spectra, with several peaks visible at integer multiples of the modulation frequency of f_m = 0.8Hz which was used. The modulation analysis was done using the main metric routine for a range of stimuli, after A-weighting the recorded stimuli using the dBFA software from 01dB. The resulting modulation spectra corresponded to the entire duration of the stimuli signal (20s).

To make the analysis scale-invariant and eliminate the 0Hz component of the spectrum, a de-trending analysis is made, by calculating an average value for the spectrum and subtracting it: this can be done for example using a 5th order polynomial, overall average or moving average (the 'optimal' method will depend on the evolution of the signal).

The modulation analysis is first made of a full set of stimuli, all normalised to 40dB $L_{Aeq,20s}$. It was subsequently verified that the analysis results were similar for stimuli designed to different normalised levels, which is consistent with the stimuli design procedure (in which the overall gain was changed), and confirms that the procedure is scale-invariant. The lack of substantial component near 0Hz shows that the stationary component of the signal has effectively been supressed: see for Figure C6.

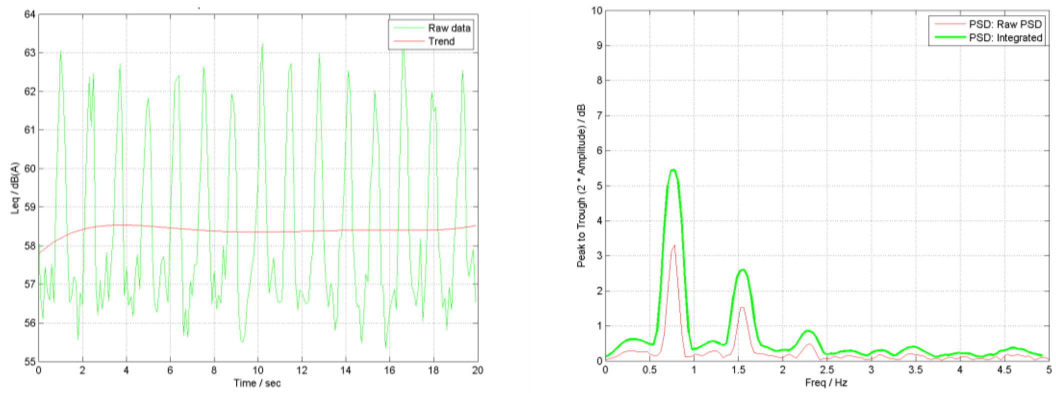

Figure C6 – Time history ($L_{Aeq100ms}$, arbitrary scale, left) and resulting modulation spectra (right) of artificial AM stimuli – example of 6 dB(A) design modulation depth

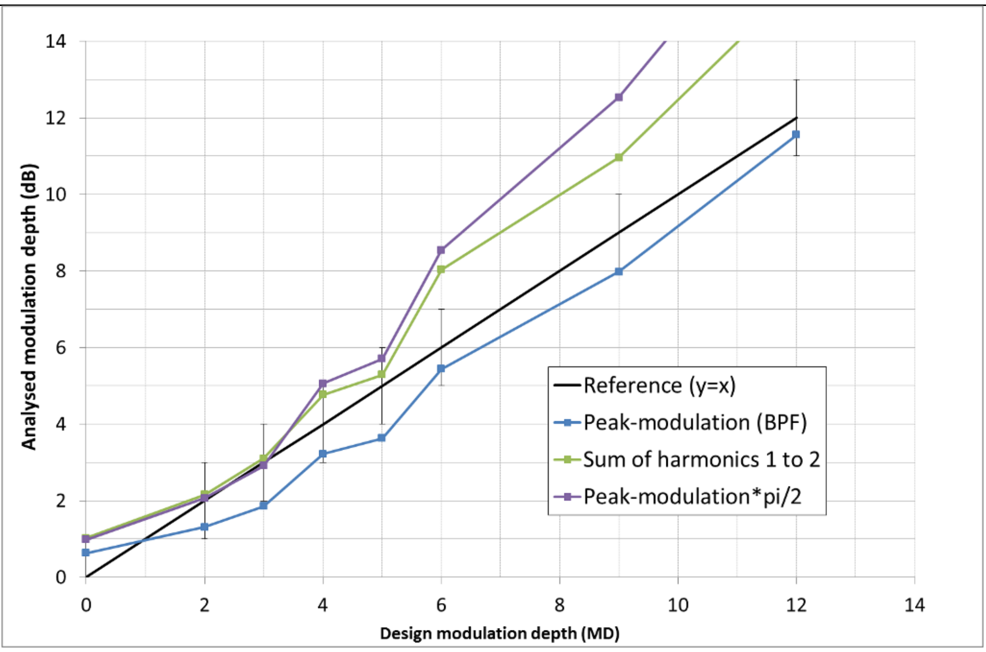

Figure C7 - Blackman-Harris window, 20s analysis period, de-trending using 5-th order polynomial – harmonics analysis

The calculated amplitude of the peak in the modulation spectra can be compared to the design (MD) value for the range of stimuli considered: see Figure C7. It can be seen that for the un-modulated signal, some residual modulation appears to be detected, and this was at frequencies close to the modulation frequency used in the other stimuli (peak at 0.7Hz instead of 0.8Hz). This could be a combination of the inherent characteristics of restricted-band signals (see WPB2), or to the inherent noise present in the analysis itself. For the actual modulating stimuli (MD ≥ 2dB), the calculated modulation amplitude peak at f_m increases in relation to MD, with a similar slope, but giving results which are consistently of the order of 1dB below MD. It should be noted that this difference of 1dB is comparable to the uncertainty in the determination of the MD quantity. As shown in Figure C8, the peaks and trough are determined as mean values of peaks of short-term dB(A) values that may vary by this amount in any case (because of masking effects): this is represented by error bars on Figure C7.

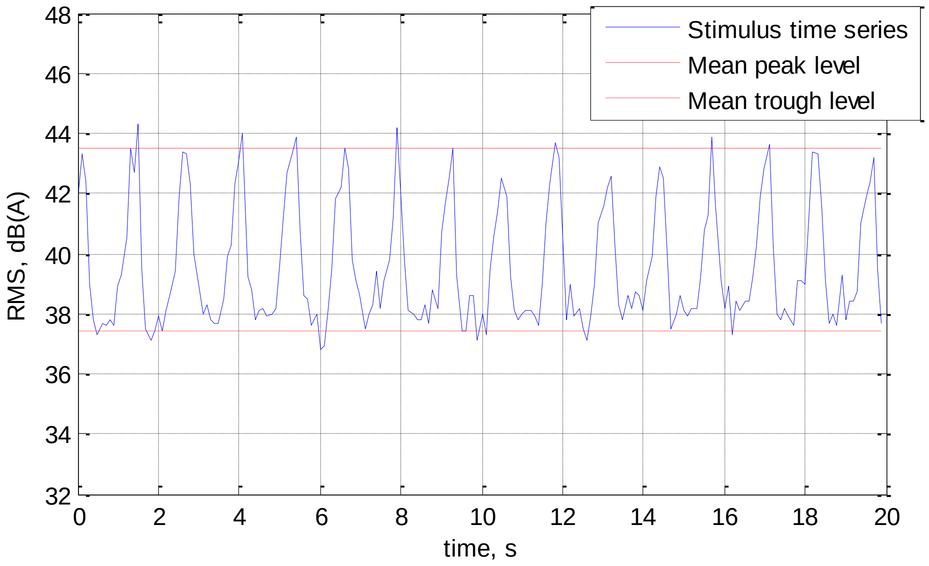

Figure C8 - WPB2 Figure 17.9: evolution of measured short-term A-weighted RMS values for an artificial stimuli at MD = 6dB and representation of mean peak and trough levels used to determine this value of MD

Because of the non-sinusoidal nature of the signal, as evidenced by the presence of harmonics in the spectrum, it is natural to seek to reflect this in the overall modulation metric. As the WPB2 stimuli pulse shape is asymmetric, we can make reference to the analysis of the pure saw-tooth signal above: firstly, a factor of π/2 was applied to the calculated amplitude of the fundamental; secondly, the sum of the amplitude of the first two harmonics was also considered (with reference to the above analysis and as harmonics of higher order were often difficult to distinguish in the calculated spectra). The resulting corrected values (Figure C7) are close to MD for the lower modulation values, but then diverge strongly for MD>5dB. The slope of the change with increasing modulation for these alternative metrics is also different from the one obtained when considering only the fundamental peak, and does not seem consistent with the results of the subjective testing shown in WPB2, which do not exhibit such a dramatic increase in annoyance with increasing MD.

These results suggest that corrections which were relevant to idealised saw-tooth signals may not necessarily be applicable in the same way to more general stimuli which were based on actual recordings of wind turbines AM. We can also consider the results of the sensitivity tests of WPB2 (see report section 7), which determined that the degree of asymmetry in the shape of the modulation pulses, as defined by a ratio of rise time versus fall time, did not have a significant effect on the subjective response to the stimuli. Finally, considering increasing harmonics of the fundamental modulation frequency will in practice be increasingly affected by noise in the signal and the analysis process.

In the Figure C9, a similar analysis of the calculated modulation amplitude of the first fundamental at f_m is undertaken for a range of different parameters. A different signal de-trending method is used, with a moving average over a window of 10s centred on each interval period (instead of a polynomial fit). A rectangular windowing function is also used. Finally, a filtered version of the signal is analysed by applying a band-pass filter covering the 315Hz 1/3 octave band, as this is the frequency region which dominates the modulating part of the signal (as is known in this case because these are artificial stimuli).

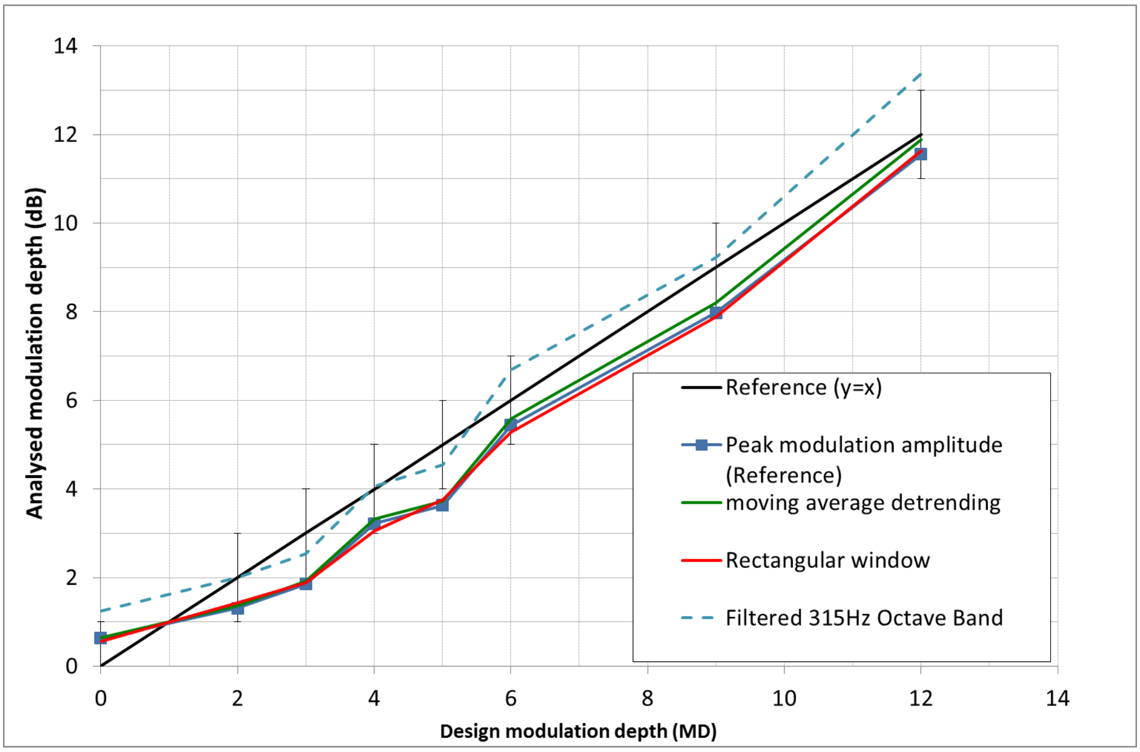

Figure C9- Analysis of the artificial stimuli using different modelling parameters

The changes in the windowing or de-trending method used in the analysis do not appear to significantly affect the analysis outcomes. Considering a filtered signal (315Hz 1/3 octave band) results in consistently higher values of peak modulation at BPF, which appear closer to the MD metric although this not necessarily significant and is likely specific to the present stimuli signal. The increase can be attributed in part to the reduced level of masking present in the filtered signal, although a contribution from the increased irregular nature of the resulting modulation spectra (see Figure C10) seems to be a factor.

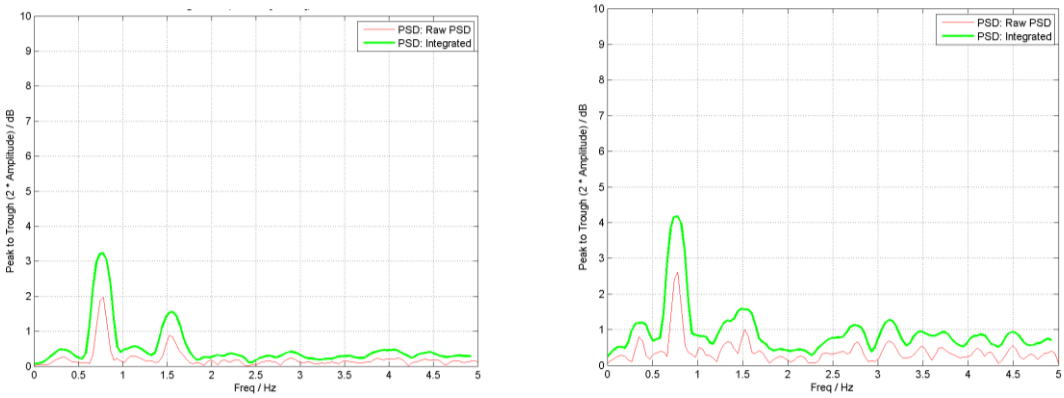

Figure C10- Comparison of modulation analysi for an artificial stimuli sample (MD = 4dB) for both the unfiltered signal (left) and that obtained filtering only the 315Hz 1/3 octave band (right).

WPB2 shows the results of the analysis of the final tests of subjective response for a range of different metrics, including the peak modulation amplitude (first harmonic) given by the implementation developed by RES. This highlights that the results of the subjective testing must be considered in relation to the appropriate metric.

<u>Application to recorded noise samples from operational turbines</u>

As a result of Work package C WPC, a selection of samples was collected. Specific periods corresponding to worst-case levels of modulation experienced during quiet periods were selected and were analysed using the technique used above. The goal of this study was to evaluate the application of such techniques to real data and evaluate its robustness, practical outcomes and the effect of range of analysis input parameters on the outcome metrics.

Some of the key feature of realistic recordings as opposed to simulated or idealised signals are:

- The presence of spurious sources of noise from background sources present in the environment, which either produce masking effects (which ‘drown out’ the modulation) or produce changes in noise levels which can affect the modulation analysis itself
- Even when dominated by wind turbine noise and/or wind turbine AM, the signals will tend to be highly variable, and levels of modulation experienced will change significantly with time

Regarding the latter point, the modulation spectra analysis using the techniques based on Fourier analysis of the signal envelopes effectively represent an average over the analysis period. The whole period will be split in segments of varying length, over which the modulation analysis is made. The selection of longer analysis periods may average out spurious effects but conversely may dampen the calculated amplitude of the periods of highest modulation.

It is useful to consider the evolution of the whole modulation spectra over the entire analysis period. This helps to highlight the range of frequencies experienced and may highlight for example the presence of harmonics of a fundamental modulation frequency, changes in modulation strength as well as periods where the analysis was affected by spurious noises (see WPB1), and the spectrum is unlikely to correspond modulation at the turbine blade passing frequency. For example: see Figure C12.

This highlights the need to isolate modulation components which may be related to the operation of the turbine (see WPB1) by concentrating on the region located close to the BPF: in the following analysis, the following will be considered separately:

- the overall peak in the modulation spectrum (as in WPB1)

- local peaks in the modulation spectrum which are located within 10 spectral lines of the expected BPF for the turbine, which was set as an input of the analysis

From the data collected as part of work package C, the analysis focused on periods in which clear modulation was identified. These represent particular periods which may not be representative of a typical situation, even at the properties identified, but were considered representative of periods of AM of interest.

The following files will be considered, with a naming convention consistent with WPB1: Table C1.

Sample File	Name	Duration (s)
1	DTI LFN [2] study - Site 1	570
2	File 1 - extract, last 20s	20
3	DTI LFN study – Site 2 - external	83
4	DTI LFN study – Site 2 - internal	81
5	Van Den Berg sample (JSV article)	173
6	Web-sourced audio (extract)	23
7	Other AM example	60

Table C1 – sample file list obtained from WPC as input to modulation analysis

<u>File 1 - DTI LFN study - Site 1</u>

The period selected (9’30’) was a period identified in the DTI LFN report [2] as one for which turbine noise modulation particular audible. This was measured during the middle of the night/early morning which was relatively quiet. Noise levels were measured externally outside the complainants’ property at a free-field location.

The overall time history of Figure C11 was divided into 10s analysis blocks which are then analysed using the above methodology. The resulting set of modulation spectra for the entire data period is Figure C12 for the following parameters:

Parameter	Value
BPF fundamental modulating frequency (f_m)	1.3Hz
Detrending Method.	5th order Polynomial

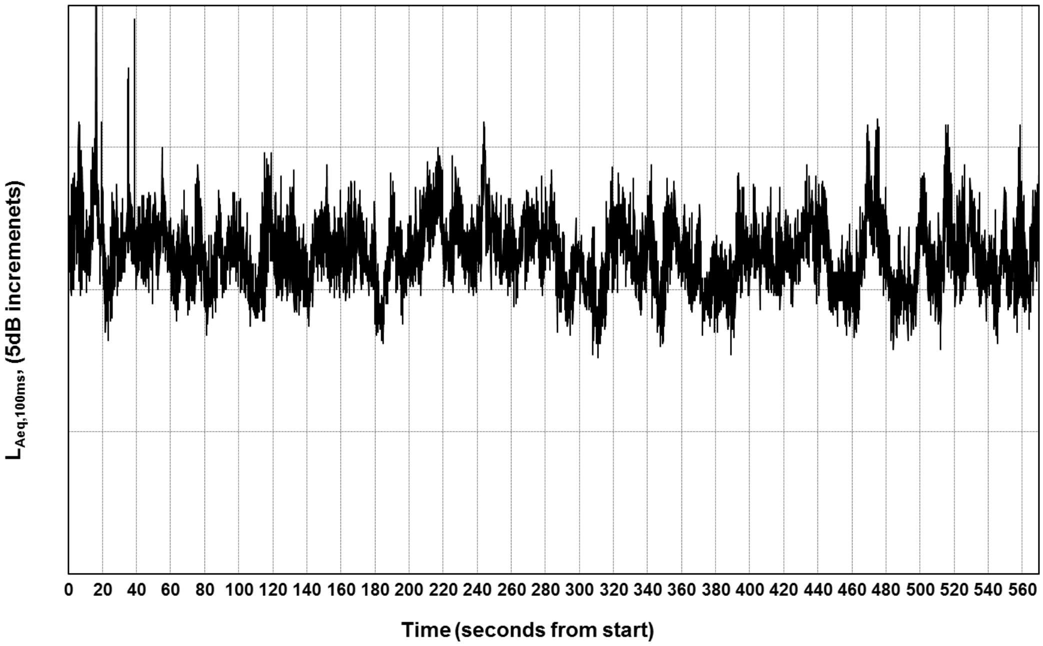

Figure C11- Evolution of the time-history of short-term $L_{Aeq,100ms}$ level over the entirety of file 1.

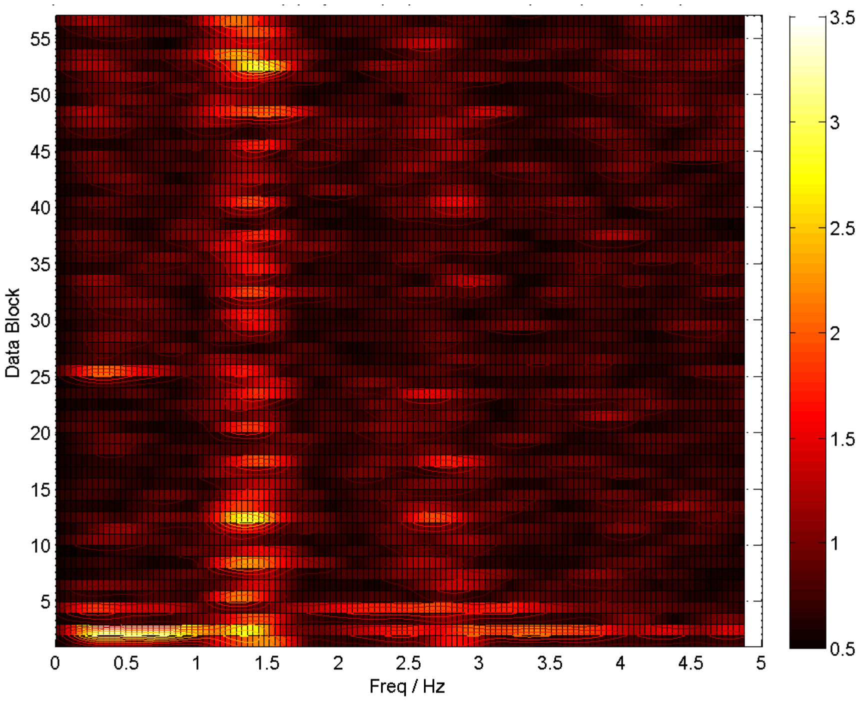

Figure C12 - Evolution of the modulation spectra for file 1 for different 10s analysis blocks (dB scale shown on the right)

In Figure C12, a clear trend highlighting the presence of fundamental modulation close to 1.3 Hz can be seen as a vertical line of varying intensity. This frequency is consistent with the reported rotational speed of the turbines of 26 RPM [2], and the turbine model operating in a fixed speed operation. A secondary, fainter line can just be distinguished: it represents the second order harmonic close to 2.6Hz. An example of a period of relatively clear modulation is shown in Figure C13. Other periods also appear to be affected by extraneous sources of noise, as identified by an unrepresentative modulation spectrum which contains components at the lowest frequencies or over a wide range of the spectrum (horizontal lines): an example is also shown in Figure C14. This justifies the choice of f_m which was made for the analysis. This is illustrated in a graph of the frequency of the peaks in the spectrum: Figure C15. When the frequency of overall peak in the spectrum deviates from f_m (as in the example of Figure C14), then the peak modulation does not represent modulation of the wind turbine noise, it is just an artefact. A metric which sums up the amplitude of the integrated spectrum at the first two harmonics (f_m and 2 f_m) is also shown, although this does not necessarily provide a particularly relevant quantity, particular when considering that the component at 2 f_m is not necessarily visible in the spectrum (Figure C12).

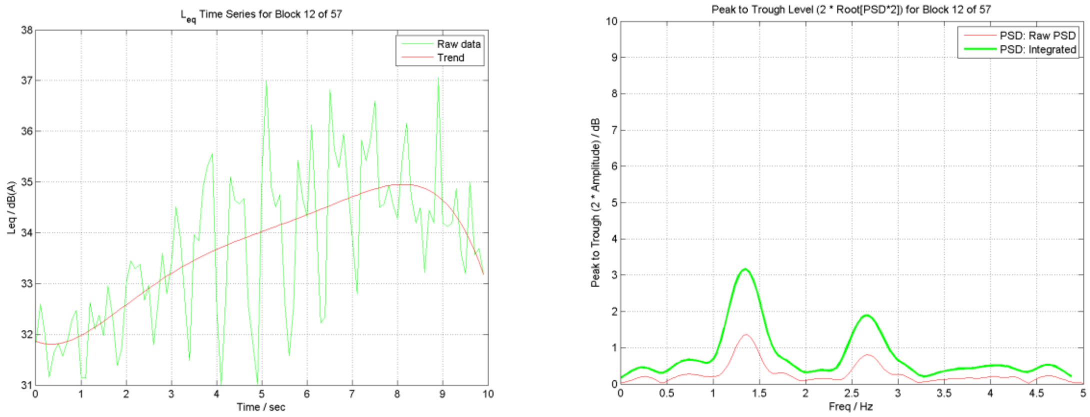

Figure C13 - Example of a modulation spectrum corresponding to clear modulation at BPF (block 12)

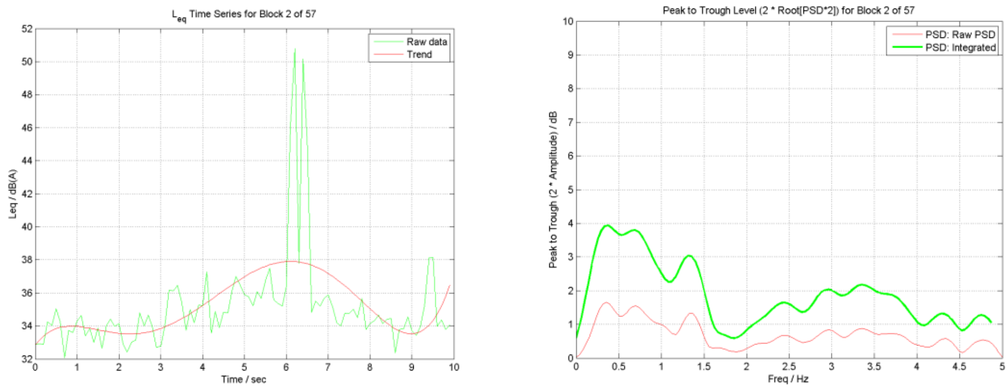

Figure C14 - Example of a modulation spectrum contaminated by spurious data (block 2, unidentified 'whistling' noise) - in this case the peak at 1.3Hz is lower than that content <1Hz

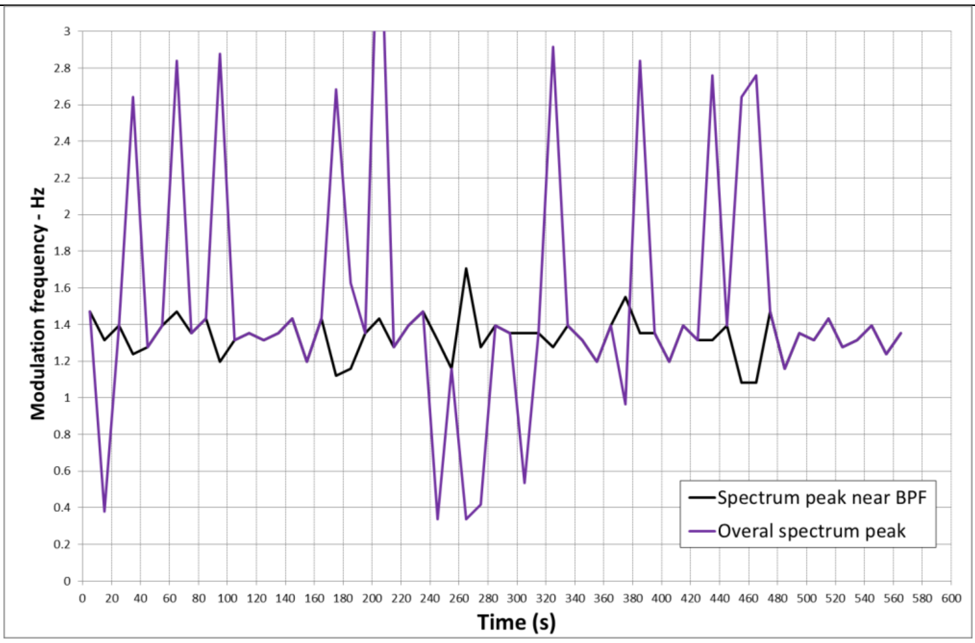

Figure C15 - Frequency of peaks in the modulating spectra: peak of maximum amplitude overall or local peak in proximity to the expected BPF

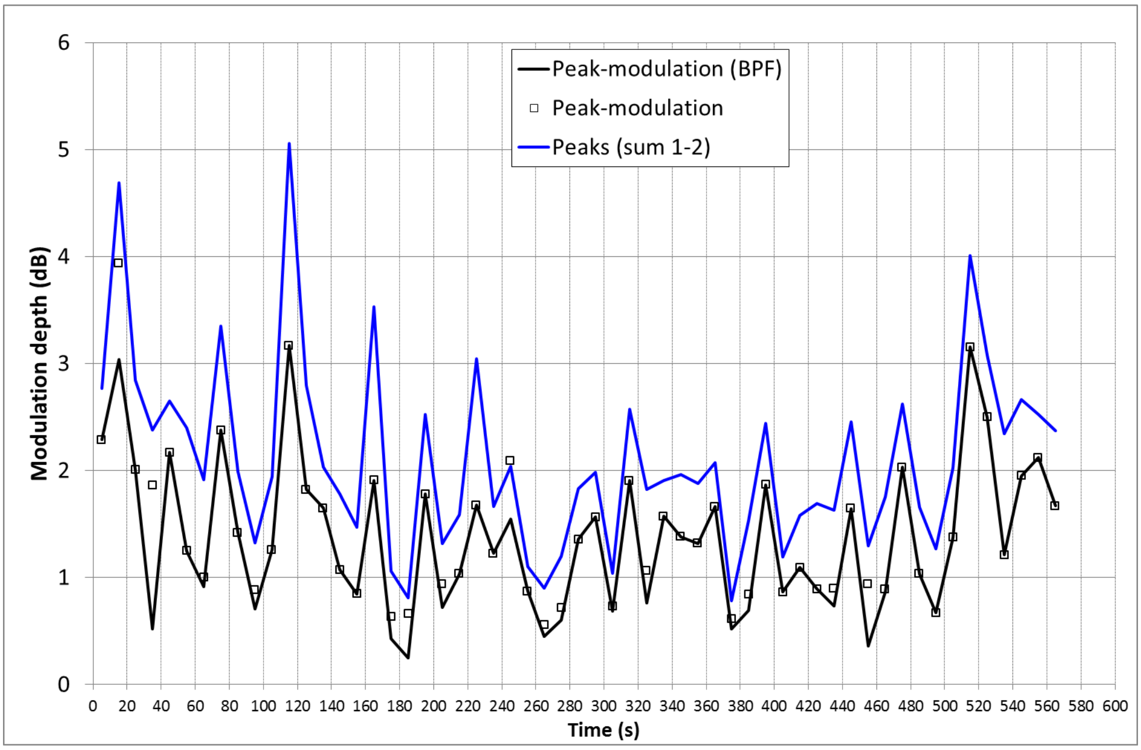

Figure C16 - File 1 – 10s analysis period – peaks in modulation spectrum shown both overall and close to fm. The sum of the amplitude of the first two harmonics of fm is also shown.

Furthermore, considering a sample of several minutes means that it is possible consider different durations for the segments or blocks over which the analysis is undertaken, as discussed above. For example, Fig. C17 presents the evolution of the modulation spectra (comparable to Fig. C12) when using 1 minute time blocks: this results in a coarser analysis. Although the modulation frequency and its harmonic are still clearly identifiable, the amplitude is affected due to an effective averaging process. Fig. C18 presents the evolution of the calculated peak modulation amplitude (close to f_m) for analysis over different time scales.

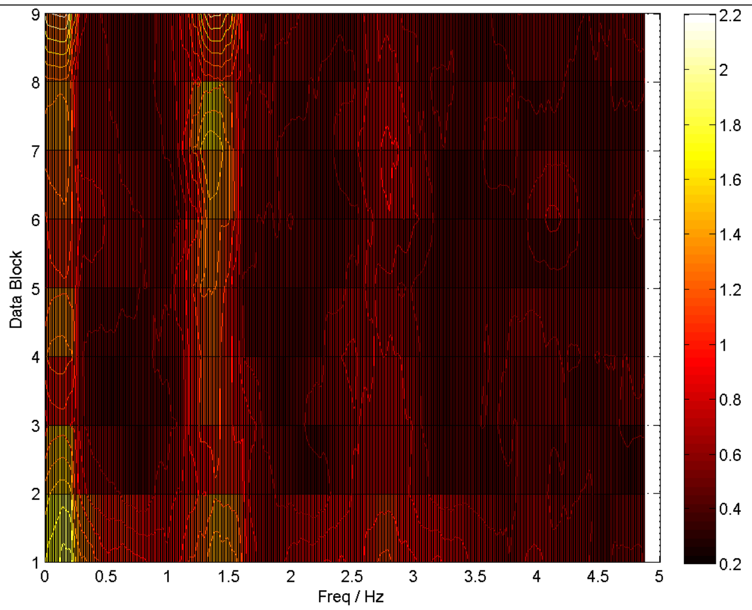

Figure C17 - Evolution of the modulation spectra for file 1 for different 1minute analysis blocks (note change in horizontal scale)

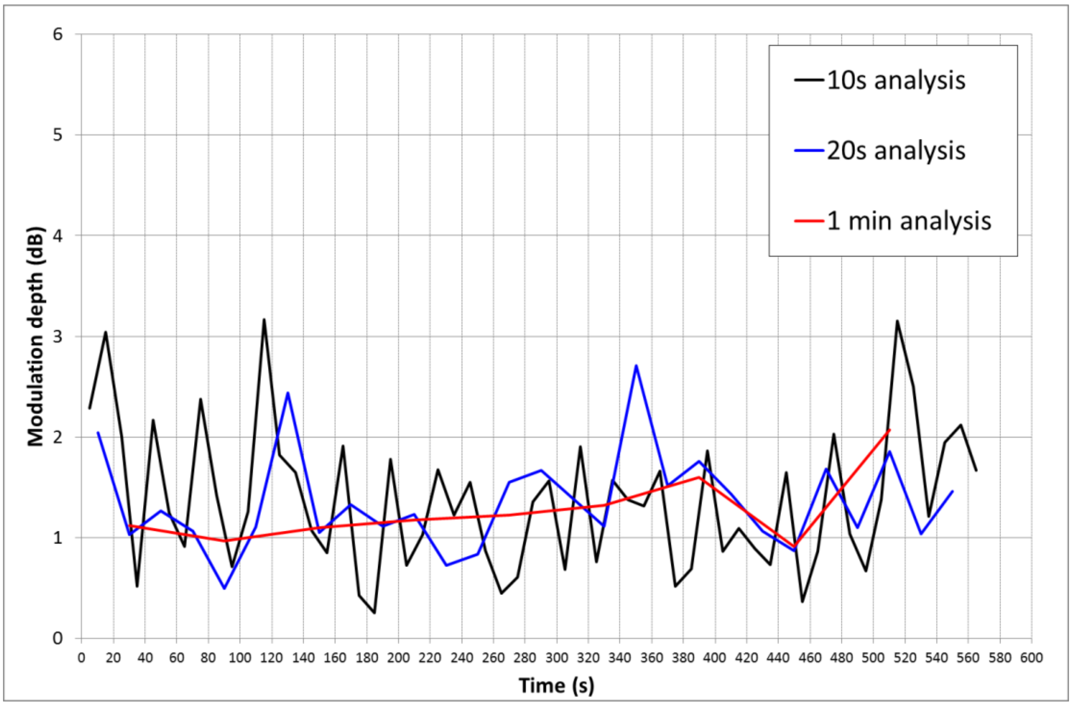

Figure C18 - Amplitude of modulation spectrum peak (near fm) for analysis of the file 1 sample, when analysing the time series in blocks of different size: 10s, 20s and 60s.

<u>File 2 – DTI LFN study - Site 1 (extract of file 1)</u>

This considers in more detail the results above, as this sample corresponds to the last 20s of the sample above. The evolution of L_{Aeq} with time (including the effect of filtering for different frequency ranges) is shown in Fig. C18, and along with the resulting modulation spectra in Fig. C19. The calculated modulation depth at f_m is 3.6dB.

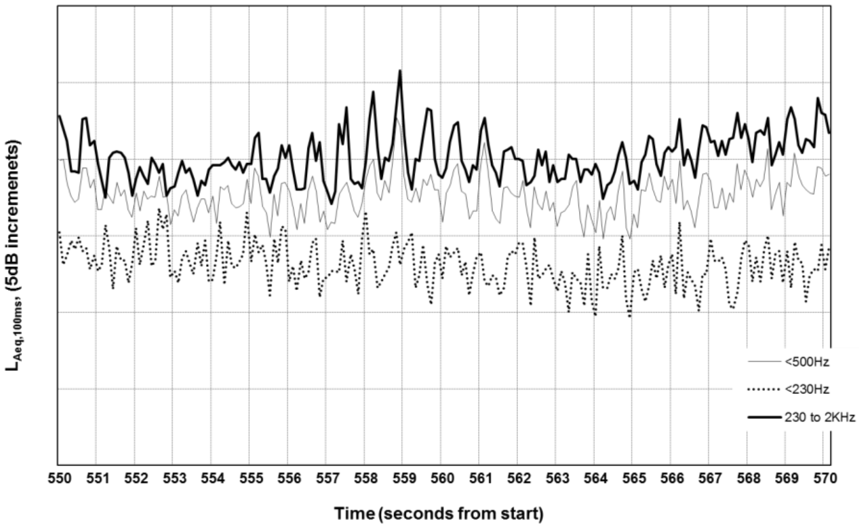

Figure C19 - Evolution of short-term Laeq,100ms level for file 2 (i.e. extract from end of file 1).

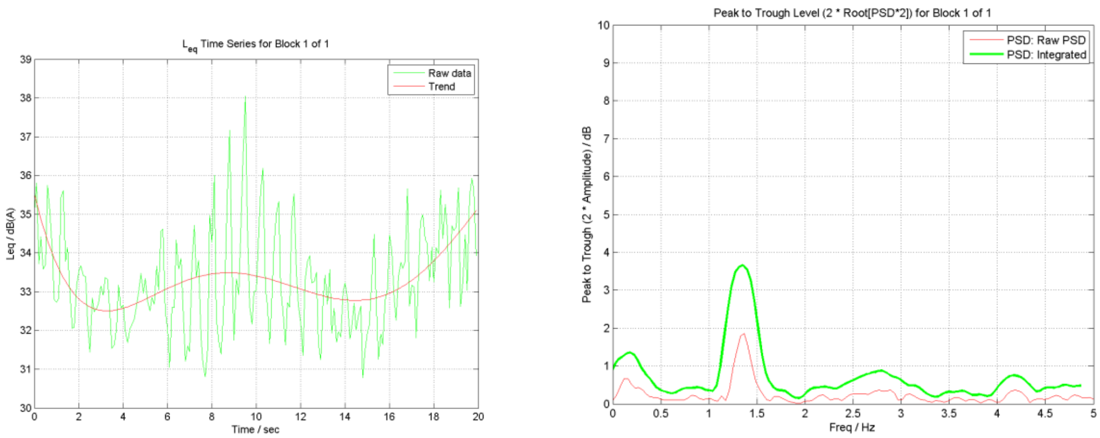

Figure C20 - Leq and resulting modulation spectra for file 2 (over entire 20s block time).

If the 20s file is separated into two 10s blocks, the peak amplitude for each block reduces to approximately 1.5dB. This demonstrates that the choice of the duration and position of the analysis period has non-trivial consequences on the calculated modulation amplitudes for cases in which it varies significantly with time.

Further analysis – files 1 and 2

WPB1 notes that using a low-pass filter on the recorded audio signal to exclude all frequencies above 1kHz may assist in eliminating many sources of spurious noises present in rural environments and would not significantly affect the wind turbine noise signal. For the signal in files 1 or 2 this does not make a significant difference to the overall analysis as the spurious noise sources cover a large frequency band.

But as noted in [2] and WPC, the modulation is dominated by the frequencies in the region 400-800 Hz. We can therefore undertake the analysis after filtering for this region, considering for example the 630 Hz 1/3 octave band. The resulting modulation spectra evolution for file 1 is given in Fig. C21, and the change in the BPF peak in modulation amplitude with and without filtering is also shown in Figure C22. When looking at the 20s extract in particular (file 2), the resulting modulation depth at f_m increases marginally from 3.6 dB to 4.4dB, when analysing over the entire 20s period.

Filtering the audio signal in a narrower octave band has therefore assisted the exclusion of spurious noise sources, but this requires a knowledge of the dominant frequency content of the modulation signal. This can be determined through a spectrogram of the signal such as shown in WPC, Fig. C2a. The effect on the modulation amplitude was also noted but this did not appear to be unrepresentative.

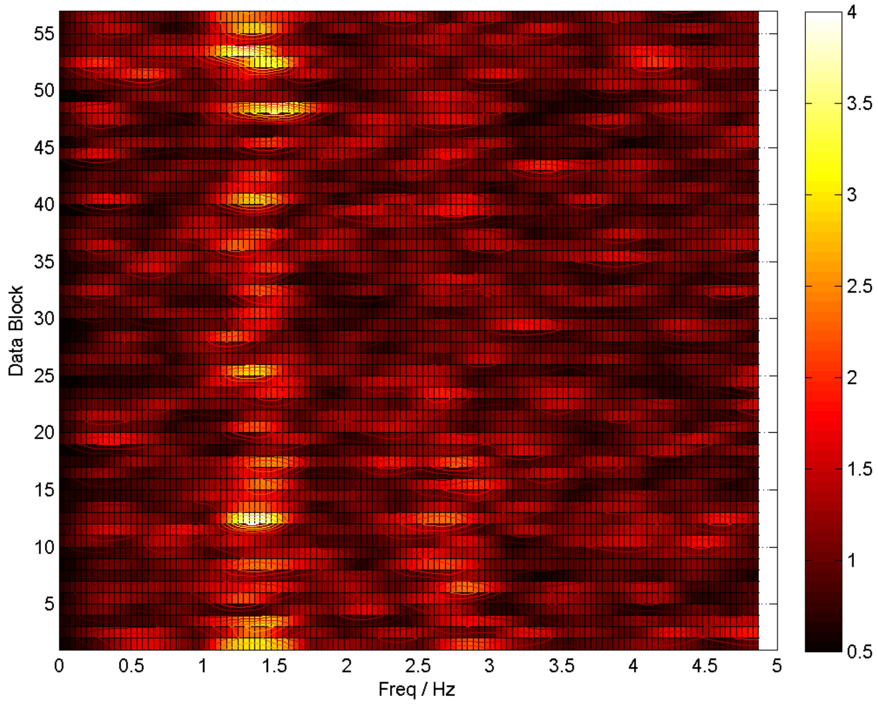

Figure C21 - Evolution of modulation spectrum for file 1, 10s analysis periods, signal filtered in the 630Hz 1/3 octave band

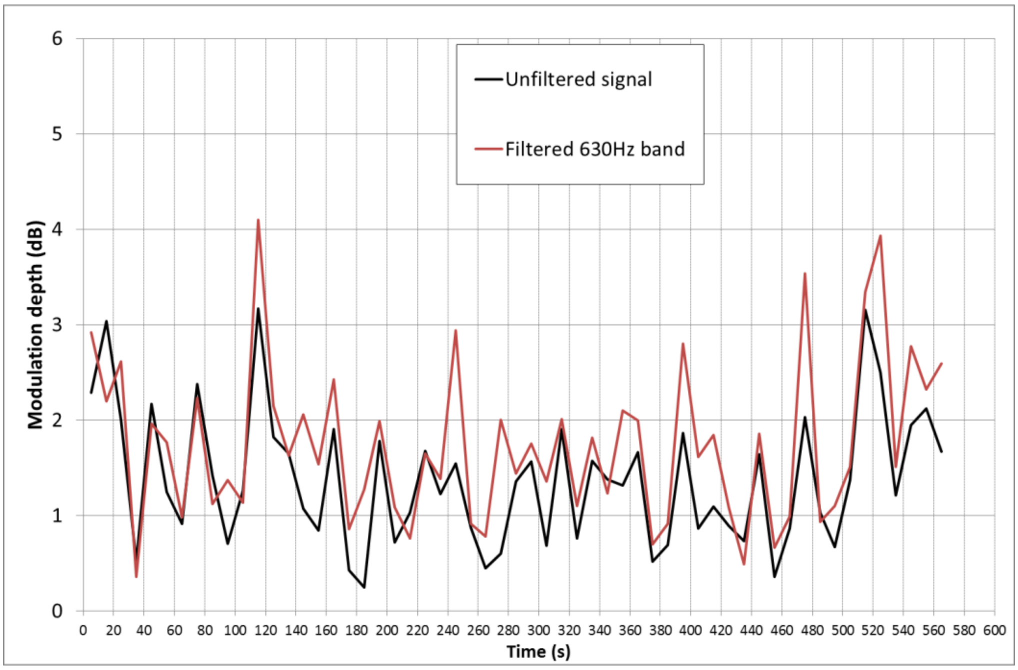

Figure C22 - Evolution of modulation spectrum amplitude at BPF for file 1, for 10s analysis periods, with the signal filtered in the 630Hz 1/3 octave band

Files 3 and 4: DTI LFN study – Site 2 - measured externally and internally

The available measurements from this site were made in parallel both at a free-field location outside the dwelling (file 3) and at an internal location (file 4) in a bedroom facing the wind farm (windows open). The period for the extract of file 3 and 4 correspond to a period of relatively marked modulation. A detailed time-history is shown below, as well as results of the modulation analysis undertaken with the following parameters:

Parameter	Value
BPF fundamental modulating frequency (f_m)	1.35Hz
De-trending Method.	5^{th} order Polynomial

The results show the presence of modulation at the BPF, which then dominate the modulation spectrum, except at some periods where peaks are below <0.5Hz and are therefore not related to the wind turbine. The value of the peak at BPF is shown for different analysis periods in both cases: it can also be seen that the values calculated outdoor and indoor are comparable. Further detailed analysis in 1/3 octave bands was not conclusive because of the lack of clarity in the frequencies dominating the modulation.

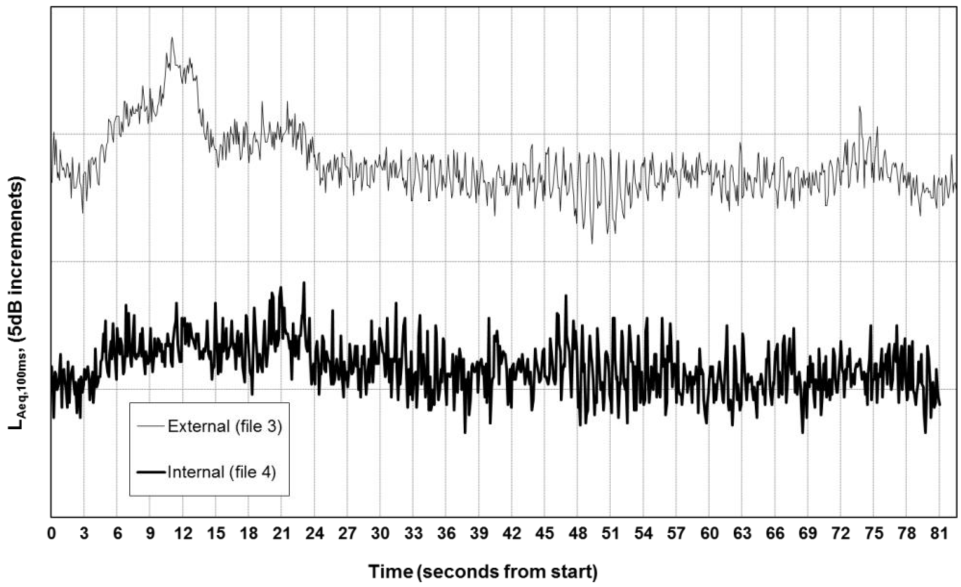

Figure C23 - Evolution of the time-history of short-term $L_{Aeq,100ms}$ levels for both the external and internal levels (file 3 and 4 respectively) at Site 2 DTI LFN study

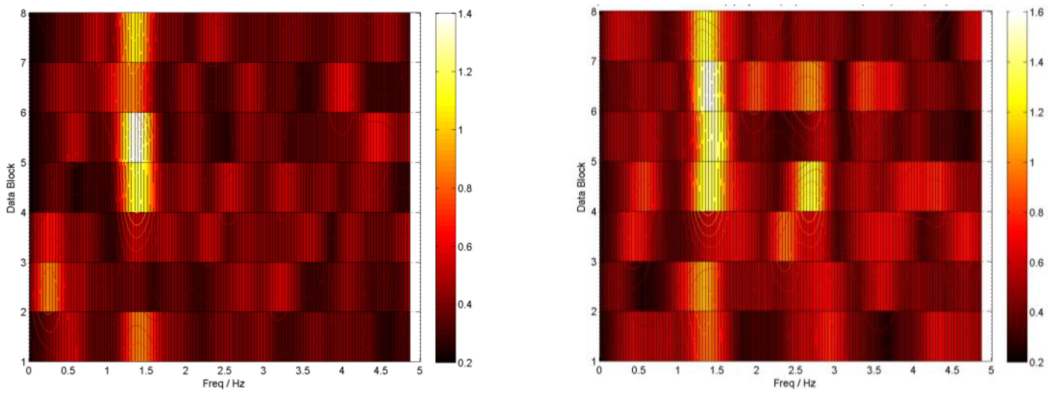

Figure C24 - Evolution of the modulation spectra for file 3 and 4 for different 10s analysis blocks

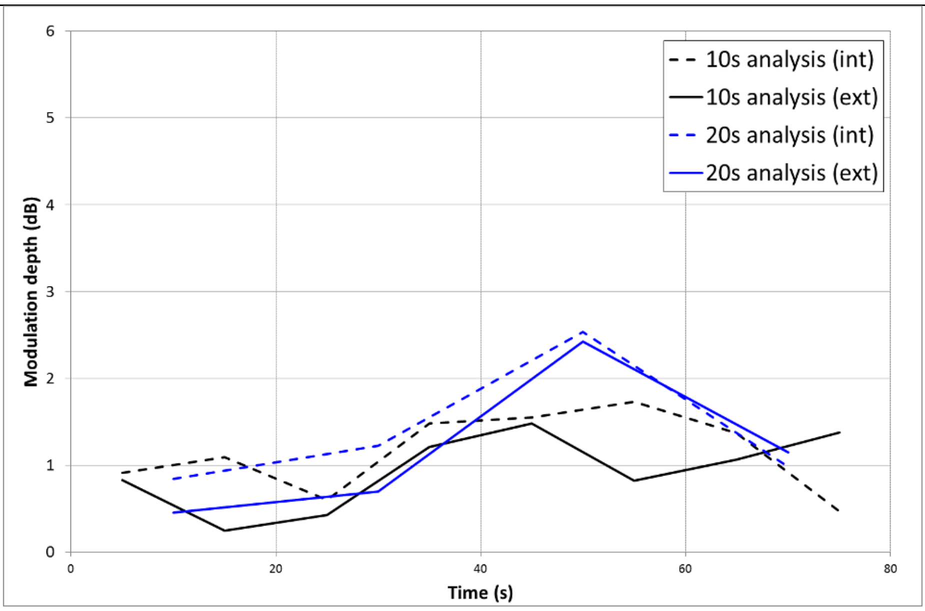

Figure C25 - Amplitude of modulation spectrum peak (near f_m) for analysis of the file 3&4 samples, when analysing the time series in blocks of different size: 10s and 20s.

File 5 - Van Den Berg sample (JSV article)

A sample of approximately 3 minutes was extracted from the data supplied, in which some variable levels of AM are present as described in WPC. The data was already supplied filtered to exclude frequencies above 1kHz. Following an analysis in 10s blocks, the resulting modulation spectra evolution (Figure C26) exhibits a clear trend as a vertical line, highlighting the presence of fundamental modulation close to the f_m=1 Hz (likely to be BPF).

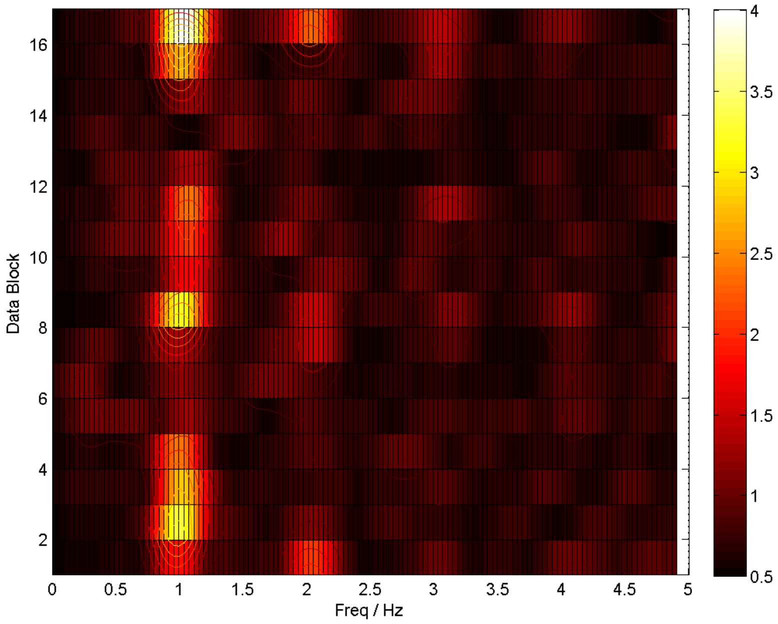

Figure C26 - Evolution of the modulation spectra for file 5 for successive 10s analysis blocks

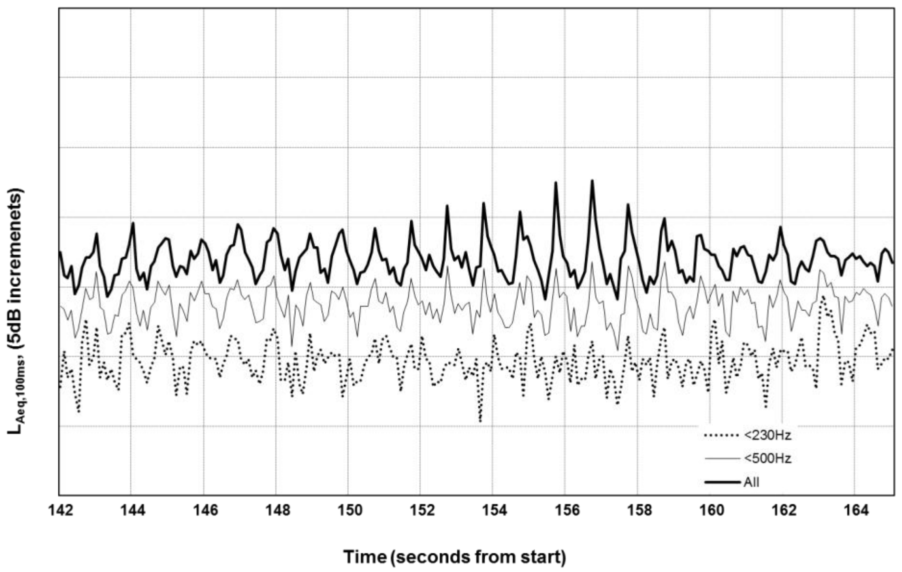

Figure C27- Evolution of the time-history of short-term $L_{Aeq,100ms}$ level for file 5 (extract).

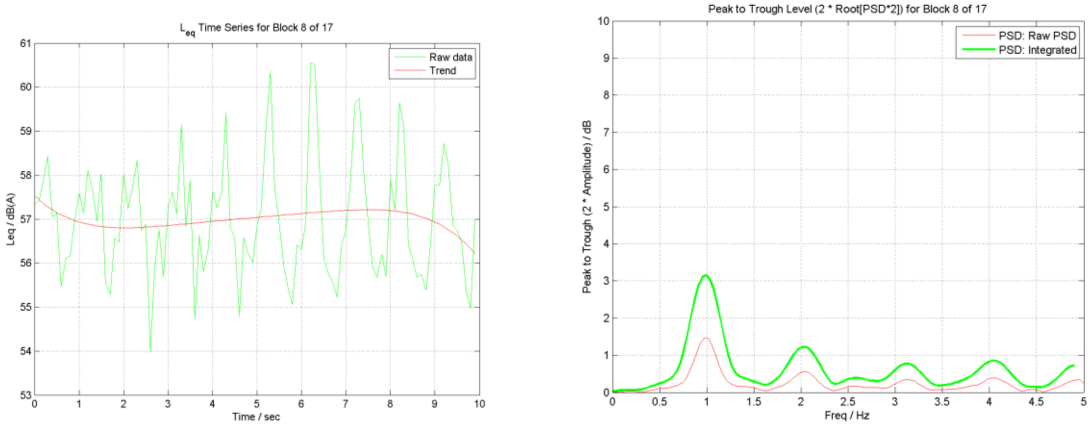

Figure C28 - Example of a modulation spectrum corresponding to clear modulation at BPF (block 18) (please note that noise levels are represented on an arbitrary scale, as the audio data was not calibrated)

Figure C27 shows a time history for a period of clear modulation towards the end of the recording, and Figure C28 another example in which several harmonics appear visible in the spectrum (as is also apparent in Figure C26).

File 6 - Web-sourced audio (extract)

There is limited information available on the recording, but it appears to be an example of turbine AM noise measured in the far-field. The A-weighted energy time history and the corresponding modulation spectrum for a representative 20s extract from the recording are shown in Figure C27. A modulation peak is apparent close to 1Hz, which could be associated with the operation of neighbouring turbines, however there are other apparently spurious peaks which could be due to artefacts in the recording.

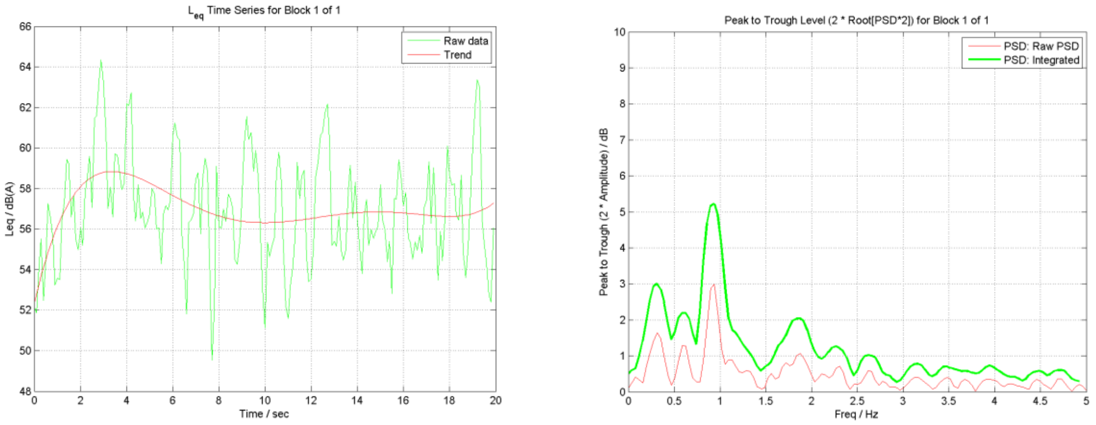

Figure C28 – Time history of 100ms noise levels (arbitrary scale, non-calibrated data) derived from an extract of the recording and corresponding modulation spectrum

File 7 – OAM example

Figure C29 shows the time history for a sample period of ‘other’ AM described in WPC, measured at a free-field location in the far-field of a wind farm site (>800m from the turbines), during a quiet period where little other sources of ambient noise were present. The corresponding modulation spectrum when undertaking the analysis is shown in Figure C30, with peak modulation levels of up to 5.5 dB present at a frequency of approximately 0.8Hz. The analysis of this data sample in 20s results in marginally lower peak modulation amplitudes in this example: Figure C31.

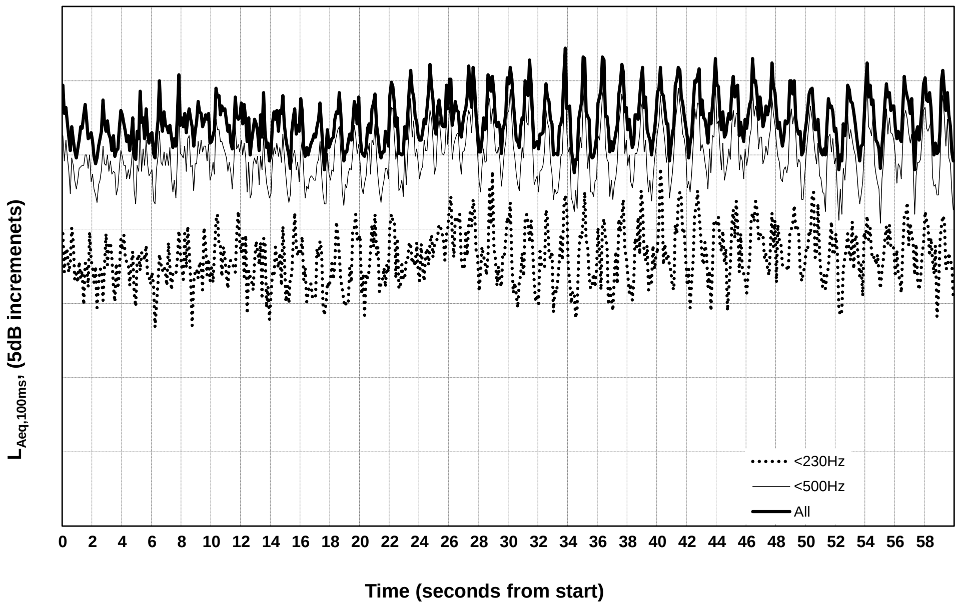

Figure C29- Evolution of the time-history of short-term $L_{Aeq,100ms}$ level for file 7

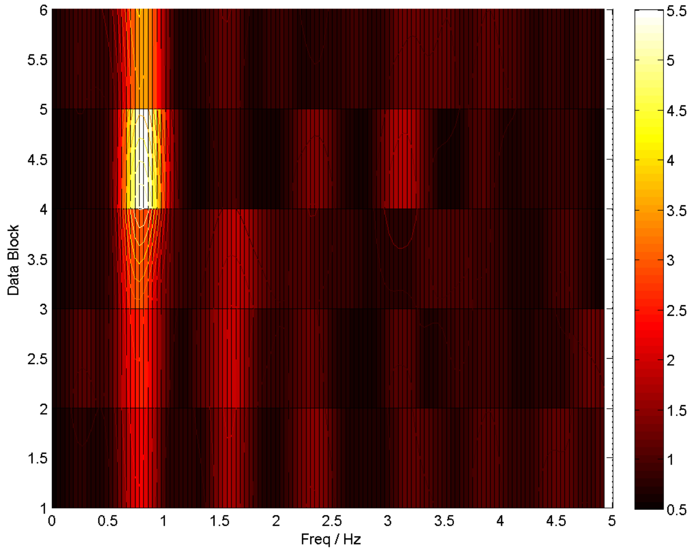

Figure C30- Evolution of the modulation spectra for file 7 for successive 10s analysis blocks

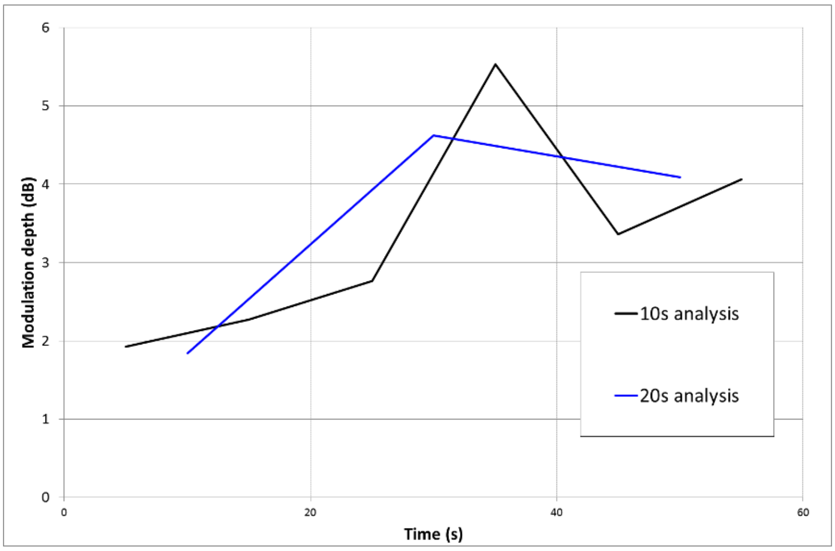

Figure C31- Comparison of peak modulation amplitudes at 0.8Hz between an analysis of file 7 over contiguous periods (10 or 20s).

Background noise analysis

A sample of audio recordings made in the absence of any wind turbine noise at a typical rural locations were analysed using the Fourier techniques described earlier. These were based on the recordings which were used for a systematic analysis of proposed AM investigation techniques [15]. Each of the recordings was approximately 10minutes long. A selection of recordings was selected at random but reviewed to represent a range of conditions, from quiet periods with or without the presence of bird noise, to noisier periods in which winds were higher or which were affected by agricultural noise. The ‘modulation’ analysis was made based either on the 100ms L_{Aeq} levels or on the 315Hz 1/3 Octave band (for consistency with the above analysis).

Quiet conditions

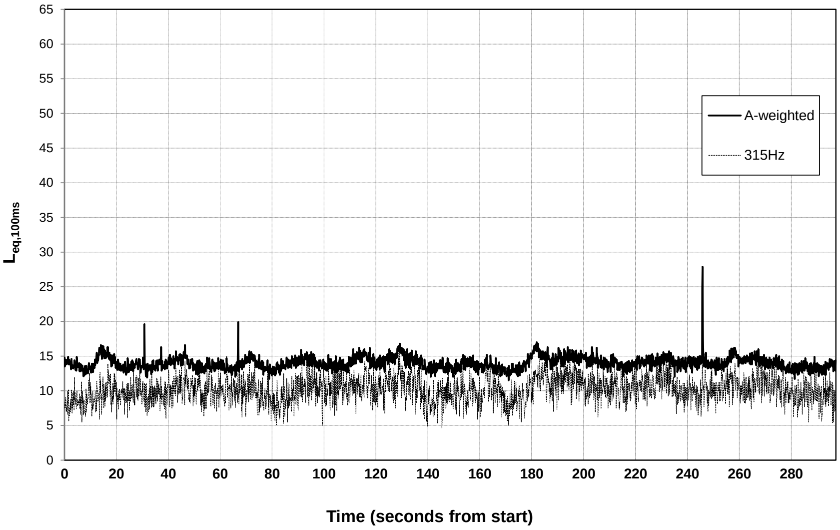

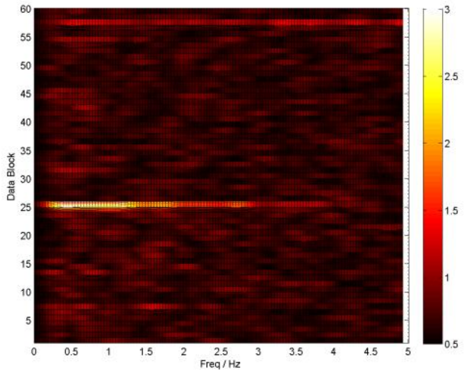

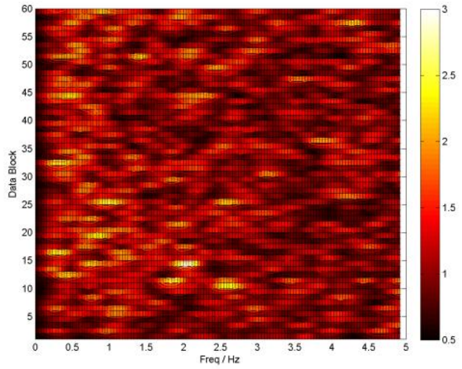

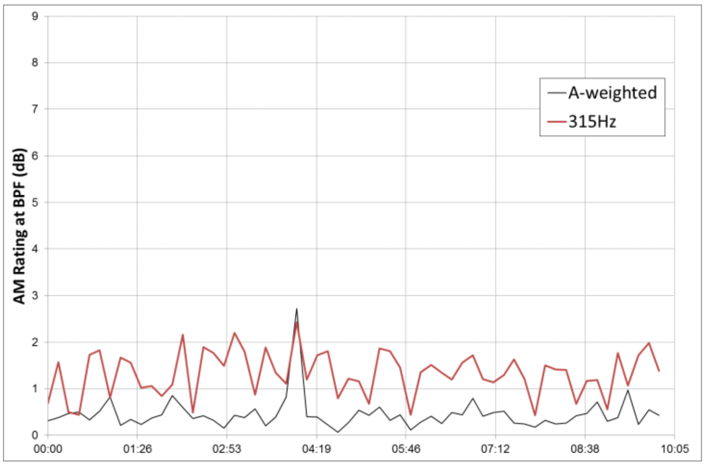

Windier period

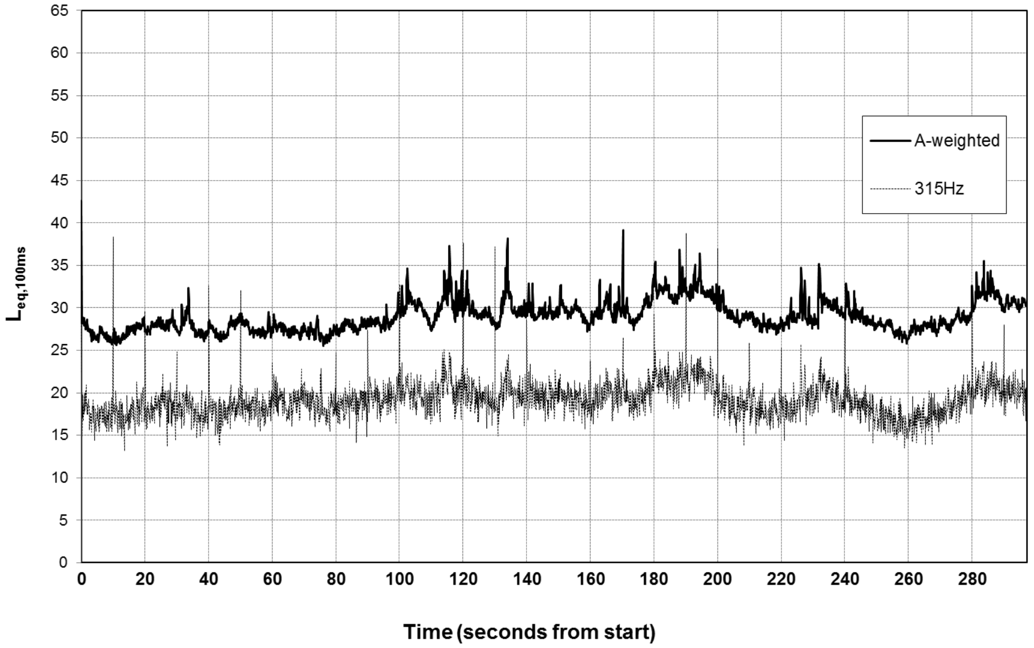

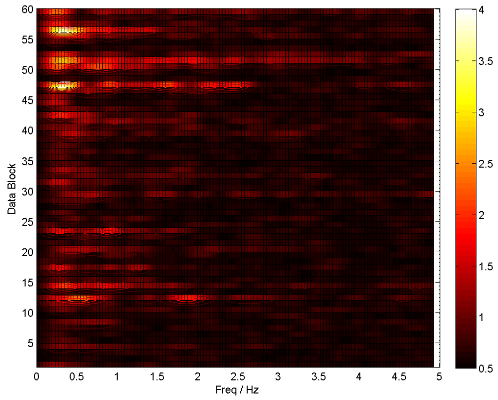

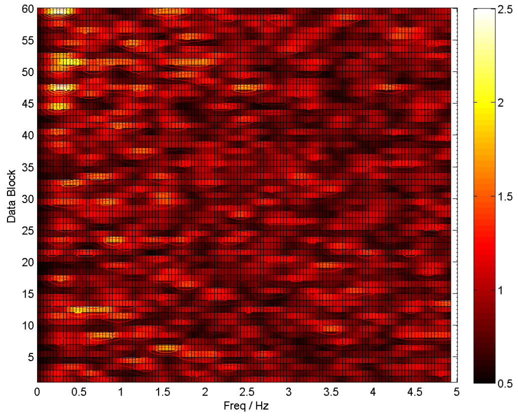

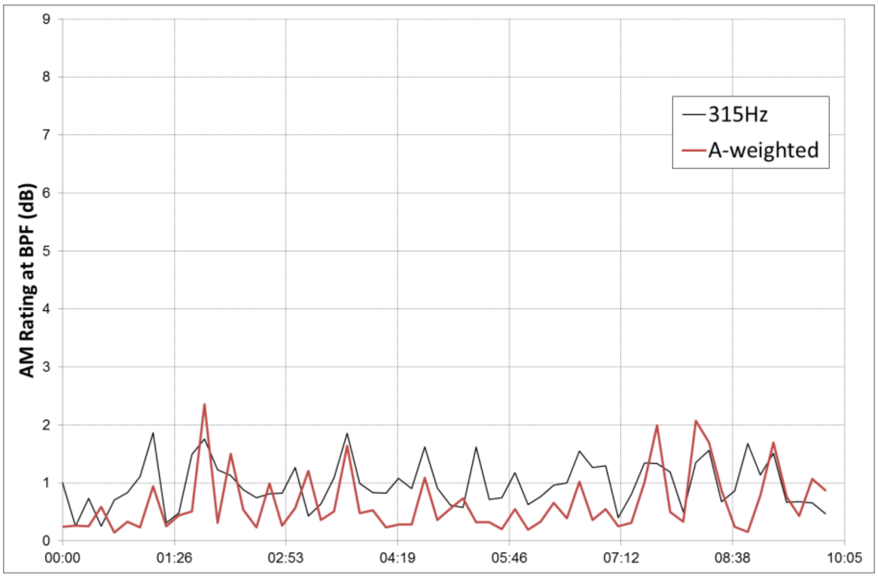

Bird noise

Bird noise can be highly variable in time and sometimes dominant in some rural locations, particularly during some early morning periods. Because of its regular nature it may sometimes appear to modulate at frequencies characteristic of wind turbine noise. Because it is generally dominated by high-frequencies, it can be suitably filtered, for example by considering the 315Hz octave band considered above.

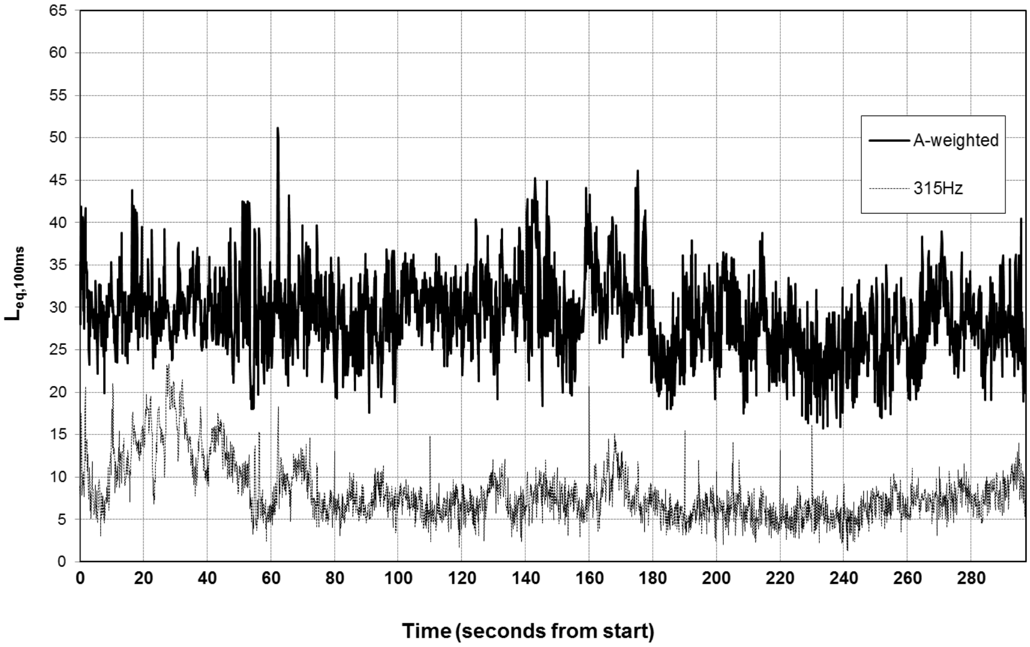

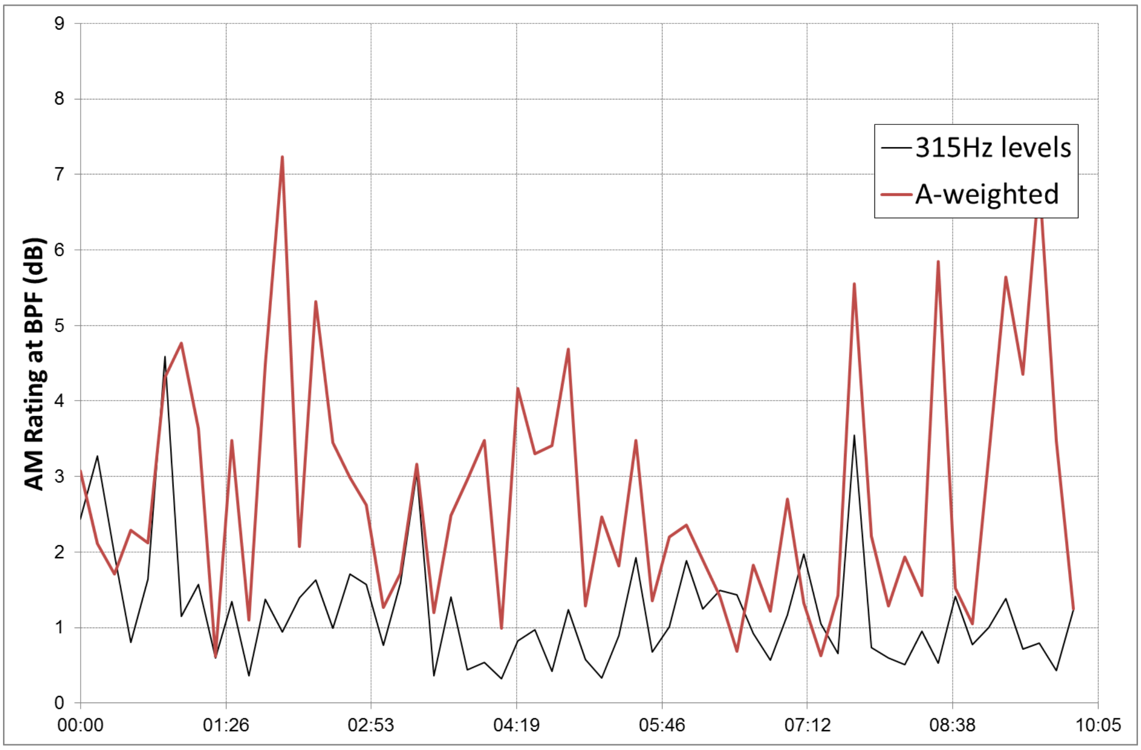

Agricultural activity

This type of noise is an example of one more difficult to filter out, but will tend not occur for limited periods only. Whilst some high values are detected, their isolated nature is not similar to that observed for WTN.