TotalControl

Advanced integrated supervisory and wind turbine control for optimal operation of large Wind Power Plants

Title: SCADA-based condition monitoring and fatigue estimation

Deliverable no.: 2.2

Delivery date: 24.06.2019
Lead beneficiary: DNV GL
Dissemination level: Public

This project has received funding from the European Union's Horizon 2020 Research and Innovation Programme under grant agreement No. 727680

Author(s) information (alphabetical):		
Name	Organisation	Email
Christine Harkness	DNV GL	christine.harkness@dnvgl.com
Eleni Maria Tsirikou	DNV GL	eleni.maria.tsirikou@dnvgl.com
Francesco Vanni	DNV GL	francesco.vanni@dnvgl.com

Acknowledgements/Contributions:		
Name	Name	Name
Martin Evans, DNV GL	Ervin Bossanyi, DNV GL	

Document information

Version	Date	Description			
			Prepared by	Reviewed by	Approved by
1	24.6.2019		Eleni Maria Tsirikou	Christine Harkness	Francesco Vanni

Definitions

CMS	Condition Monitoring System
DEL	Damage Equivalent Loads
FFT	Fast Fourier Transform
IEC	International Electrotechnical Commission
LCoE	Levelised Cost of Energy
O&M	Operations & Maintanance
OEM	Original Equipment Manufacturer
OLE	Object Linking and Embedding
OPC	OLE for Process Control
SCADA	Supervisory Control And Data Acquisition
SCM	SCADA-based Condition Monitoring
WSM	Wind Sector Management

1. TABLE OF CONTENTS

1. Table of Contents 3
2. Executive Summary 5
3. Introduction 6
4. Definition of Condition 7
   4.1 Integrity 7
   4.2 Life 7
5. Measuring integrity 9
   5.1 Turbine alarms 9
   5.2 Condition Monitoring Systems (CMS) 9
   5.3 SCADA-based condition monitoring (SCM) 10
      Signal trending 10
      Frequency analysis 12
   5.4 Inspections 13
   5.5 Comparison of different integrity measurement approaches 14
6. Measuring Life 15
   6.1 Nominal life vs. design life 15
   6.2 Life estimation process 16
      Assumptions within the analysis 16
   6.3 Life estimation inputs 16
      Environmental and operational parameters 16
      Turbine model 17
      Control system information 18
      Turbine design information 18
      Operational strategies 18
   6.4 Fatigue life calculation 19
      Direct simulation vs loads database 20
      Deterministic vs probabilistic analysis 21
      One-off vs ongoing analysis 21
   6.5 Estimating uncertainty 22
7. Turbine condition and operational strategies 24
   7.1 Asset management strategies 24
      Turbine integrity and asset management 24
      Turbine loading, life and asset management 24
   7.2 Turbine control and wind farm control 25
8. Conclusions 27

9. References 28
10. Appendices 30
  10.1 Appendix A 30
  10.2 Appendix B 30

2. EXECUTIVE SUMMARY

In this deliverable we provide an overview of techniques to measure the condition of a wind turbine, which we define as a combination of its integrity, i.e. its ability to operate safely and effectively, and of its life, i.e. its ability to withstand fatigue loading.

After providing working definitions of integrity and life, we explore and compare the main approaches to monitoring integrity, including:

- Turbine alarms
- Condition Monitoring Systems (CMS)
- SCADA-based condition monitoring (SCM)
- Inspections

and the main approaches to measuring life, in particular:

- Direct simulation of load cases vs extraction from a loads database
- Deterministic vs probabilistic analysis
- One-off vs ongoing analysis

We explore how different approaches can fit within the optimisation of asset management strategies and of turbine and wind farm control to reduce overall LCoE. Particular focus is given to the benefits of utilising SCADA data, which enables:

- The deployment of signal trending and, if higher frequency data are available through OPC or IEC interfaces, or frequency tracking algorithms. These provide some of the predictive functionality of CMS at a fraction of the cost.
- The estimation of site-specific conditions and turbine operational information based on historical data rather than pre-construction assumptions. This provides a better measure of the conditions that have been experienced on site as it removes the uncertainty associated with wind modelling.
- Ongoing estimate of turbine loading and life, which can be used in a turbine control and wind farm control framework.

3. INTRODUCTION

The condition of an asset is defined as its state with regard to its appearance, quality, or working order. In the context of wind turbines, condition affects whether, how and for how long they can be operated, and has strong implications on their performance and on the frequency of the failures they experience. Turbine condition is clearly an essential characteristic to consider when defining a wind farm's operational strategy.

While visual inspections and monitoring of certain turbine parameters can provide insights into the condition of a turbine, some aspects of its condition are difficult to measure directly – a good example being its "life", intended as its capability to withstand fatigue loading.

In this document we will define the main aspects of a turbine's condition and we will discuss ways in which they can be measured and/or estimated, with particular focus on the use of Supervisory Control and Data Acquisition (SCADA) data [1].

We will further explore how these estimates can be used to design an operational strategy, whether it is an asset management one, implemented by humans, or a turbine or wind farm control one [2], implemented in digital systems.

4. DEFINITION OF CONDITION

The condition of a wind turbine is a concept that encompasses several different aspects. In this document we will focus on two in particular: the integrity of the wind turbine and its remaining life. These are the aspects that determine whether and for how long a turbine can operate.

Power performance, defined as the amount of power that the turbine produces, for a given wind resource, compared to a baseline expectation, and which has an impact on how the turbine operates, has not been included in the scope of this document.

4.1 INTEGRITY

The integrity of an asset is its ability to perform its required function safely and effectively [3]. In this document we use the term integrity to represent the overall turbine "health", a collection of characteristics and metrics which can help in providing an answer to the following questions:

- Can the turbine operate safely at this moment in time?
- Does the turbine require any repairs or corrective maintenance?
- Are there any signs indicating that a failure might be imminent?

Answers to the first two questions will generally come from application of yes/no test criteria, that can be implemented in control logic (alarms and status codes based on data from the turbine sensors) or in human procedures (e.g. inspection checklists).

An answer to the third question can be provided by the analysis of several alternative sources (visual inspection records, vibration signals, temperature signals) in order to determine the probability of a failure occurring in the future. Based on this probability, the asset manager may decide to continue or suspend operation, to request more detailed investigations or the replacement of certain components.

4.2 LIFE

The concept of turbine life is defined in the context of reliability, the inverse of probability of failure. A wind turbine's structural design life is the number of operating years after which the reliability level estimated for a set of design conditions is expected to drop below a target threshold.

The estimate of the probability of failure is based on a generic and simplified deterministic set of environmental conditions assumed to be experienced by the turbines, according to a standard definition of type class provided in design standards such as IEC [4] and GL [5]. These standards dictate that wind turbines must be designed to a nominal life of at least 20 years, and that the design process must consider both fatigue and extreme strength. Extreme loads relate to instantaneous stress on components, whereas fatigue loads relate to the cumulative damage to a component that increases over time.

The design process involves making changes to the turbine design specifications until extreme loads and fatigue loads (with appropriate safety factors) are such that, for the chosen type class conditions, an annual probability of failure below the target one. The latest IEC standard [4] defines a threshold for the annual probability of failure of one in ten thousand (10^{-4}), and specifies wind classes (e.g. class I, II and III) and turbulence class (e.g. class A, B and C).

A wind turbine is a complex system, comprising numerous components. The structural design life of a turbine is related that of its safety-critical components [6] (typically the major structural components such as the tower, nacelle mainframe, rotor hub and main shaft). Each component is normally designed with an individual target reliability level of 1 in 10,000 failures annually.

Note that in certain cases specific components (often the tower) can be certified to a different class than the rest of the turbine.

The blade is a structural turbine component which is considered non-safety-critical in certain contexts since a blade failure does not necessarily mean the total loss of the turbine. Nevertheless, a blade failure is a safety concern for maintenance personnel or local infrastructure and nearby people. Blades can be classified as replaceable components although practical supply-chain considerations should be applied when making this categorisation.

Redundant, or non-safety critical sub-components in the system may have a lower target reliability level since these can be replaced. For example, the “L10” criteria applied in the specification of bearings implies a cumulative probability of failure of 10% over the design life. It is understood however that such failure modes should not present a safety risk to the rest of the wind turbine system. The reliability of these components would form key inputs to an operational cost analysis. Key to the reliability of these components and sub-systems is how they are operated and maintained.

5. MEASURING INTEGRITY

As discussed in Section 4.1, assessing turbine integrity encompasses a range of approaches [7], some of which provide a binary response (the turbine is/is not able to operate safely, the turbine does/does not require repairs) and some of which provide an estimate of the probability that the turbine will fail in the future unless some corrective action is taken.

This section provides an overview of techniques and procedures aimed at determining the integrity of an operating turbine.

5.1 TURBINE ALARMS

One of the first indications of the integrity of a turbine is provided by the turbine's control system in the form of alarms and status codes. Alarms are raised by the control system as a result of a set of logical conditions being verified, of variables assuming a value outside of a given range, or a combination of these. Such conditions and variables can include:

- Measurements from the turbine sensors
- Internal status variables
- Alarms from subsystems (e.g. component status codes).

The turbine's control system has the objective of regulating the turbine's control parameters in order to optimise some benefit function (often energy capture) while maintaining the turbine in a safe range of operations. Because it acts as the last line of defence against faults and unexpected operating conditions, the supervisory logic implemented in a turbine's control system is generally simple and fail safe, designed to respond to, rather than pre-empt, integrity issues. The information provided by alarms and status codes is therefore generally of a binary nature, i.e. it states whether at a given time the turbine is ready to operate, without any indication of the probability of future failures.

5.2 CONDITION MONITORING SYSTEMS (CMS)

Condition Monitoring Systems are devices (or sets of devices) which help wind-farm owners and operators monitor the health of turbine components and electrical systems. Their purpose is to analyse data measured by sensors in order to identify patterns or trends that may be indicative of specific maintenance requirements or of incipient failures. This enables operators to reduce repair and maintenance costs by:

- conducting maintenance activities, repairs and replacements only when needed
- detecting faults and taking corrective action before they cause secondary damage.

The types of CMS normally used for wind turbines generally rely on one of the following:

- vibration sensors attached to gearboxes, generators and drive trains
- strain gauges attached to blades and towers
- oil-particulate systems which count metal pieces floating in lubricating fluid.

The first two types rely on frequency analysis to identify changes in the vibration signatures of a turbine component, whereas oil monitoring systems check that defined thresholds are not exceeded. In all cases, the analysis of the data recorded by the sensors can be done automatically and continuously, or off-line on a regular basis. Even when the analysis is done automatically however it is worth noting that integration of the CMS outputs within the algorithms running on a turbine's control system can be difficult, given that the two are often provided by different manufacturers. In most cases, this integration is limited to the turbine's control system reading binary alarms generated by the CMS.

Less common CMSs rely on acoustic or optical sensors.

5.3 SCADA-BASED CONDITION MONITORING (SCM)

SCM refers to the analysis of the signals contained in a turbine's SCADA (Supervisory Control And Data Acquisition) output dataset in order to identify patterns that may be indicative of an incipient failure.

This approach is becoming increasingly popular, one of the main reasons being that all modern turbines provide outputs via a SCADA interface so that there is no additional cost involved with collecting them. While the main purpose of the SCADA outputs is real-time monitoring of the turbine's production output and operational status, the SCADA datasets are becoming increasingly extensive and often include signals from a large number of sensors and sub-components, which can be used (either in real-time or retrospectively) for analysing the turbine's behaviour and performance.

Two common types of SCM analysis are signal trending and frequency tracking.

SIGNAL TRENDING

This approach relies on establishing a model which characterises the "normal" relationship between a number of selected input and output signals, based on the numerical analysis of a period of historical data which does not include any failures. The model is then used on an ongoing basis to predict what the value of the outputs should be, based on the measured values of the inputs. When the difference between the predicted outputs and the measured ones exhibits patterns (such as exceeding a threshold) that may be indicative of an incipient failure, an alert is raised.

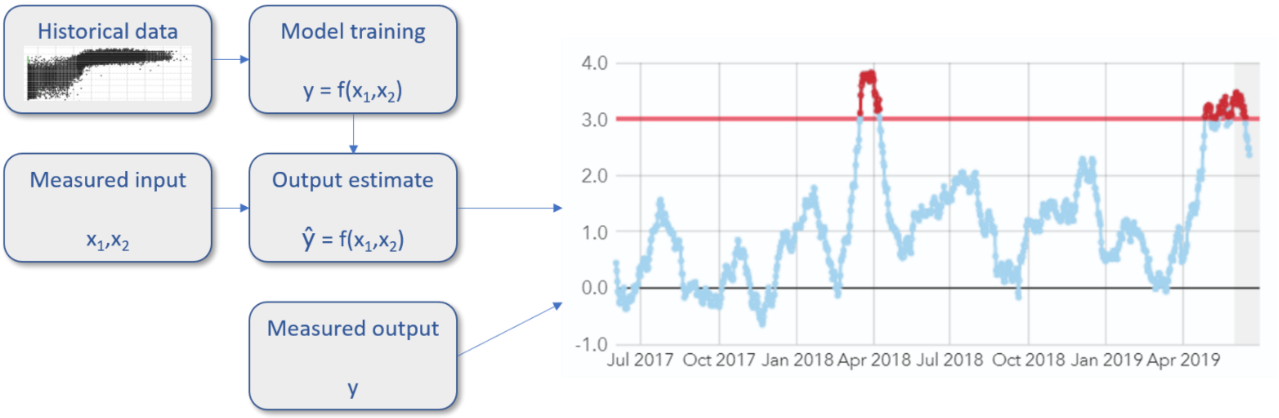

The SCADA outputs typically used in trending algorithms are the temperatures measured by thermocouples at different drivetrain components (gearbox, generator, bearings, oil sumps). The most basic models are based on a linear correlation between a single input, generally the turbine's output power or its generator speed, and the temperature to be monitored. Figure 1 and Figure 2 show an example of a single-input linear trending for the gearbox bearing temperature, where the model is obtained by correlating historical values of rotor speed and temperature. Once the model has been trained (i.e. the parameters have been identified) it is used to produce the estimated value, which is compared against the measured one.

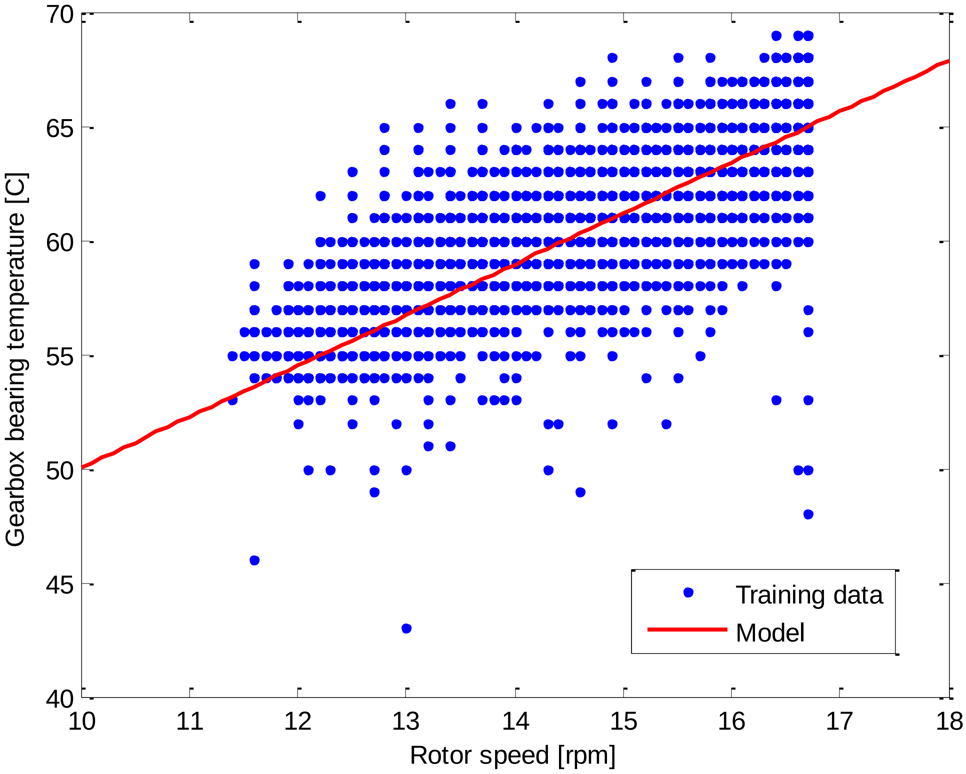

FIGURE 1: TRAINING OF SCM MODEL

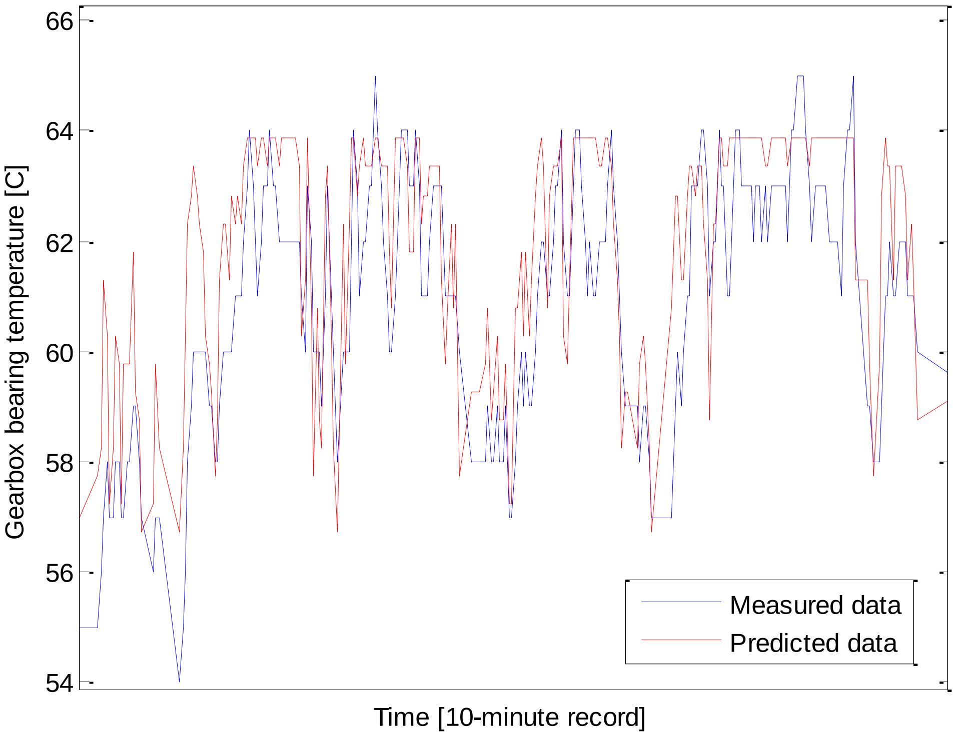

FIGURE 2: DIFFERENCE BETWEEN MEASURED AND PREDICTED TEMPERATURE DATA

More complex models may be based on higher order correlation, or on polynomial fitting, and can include multiple inputs, for example ambient temperature (to remove the seasonal effects that changes in the environmental temperature have on the measured outputs). The advantage of more complex models is that they allow a better characterisation of the nonlinearities in the relationship between the modelled variables. In DNV GL's experience however simpler models have the advantage that they can be "understood" more easily.

An example of this is a scenario where the training period used to generate the models contains data from a turbine that is already developing a failure. Simpler and more complex models would both be trained to "include" the abnormal behaviour, so that the turbine may fail without any significant difference being observed between measured and predicted temperatures. Simpler models however can be sanity checked more easily, and the models generated for different turbines on the site can be more easily compared to identify any clear outliers.

FREQUENCY ANALYSIS

While classic SCADA data outputs are generated in the form of 10-minute records, SCADA interfaces such as OPC and IEC allow turbine signals to be read at higher rates (generally between 0.1Hz and 10Hz). Collecting turbine data at rates above 1Hz enables analysis of the lower frequency structural mode such as the 1st tower mode and the rotor mode (1P).

This analysis is significantly more limited than that which could be achieved with higher frequency data (such as those read by a hardware-based CMS or a turbine controller), not only because it does not capture the dynamics of higher-frequency structural modes, but also because anti-aliasing filtering is seldom possible and the quality of the data is poor, often exhibiting irregular, repeated or missing timestamps. Nonetheless, given that the data are available without any cost for additional hardware, and that appropriate filtering can address some of the data issues above, frequency analysis of SCADA data can be a useful and relatively inexpensive tool to monitor turbine integrity.

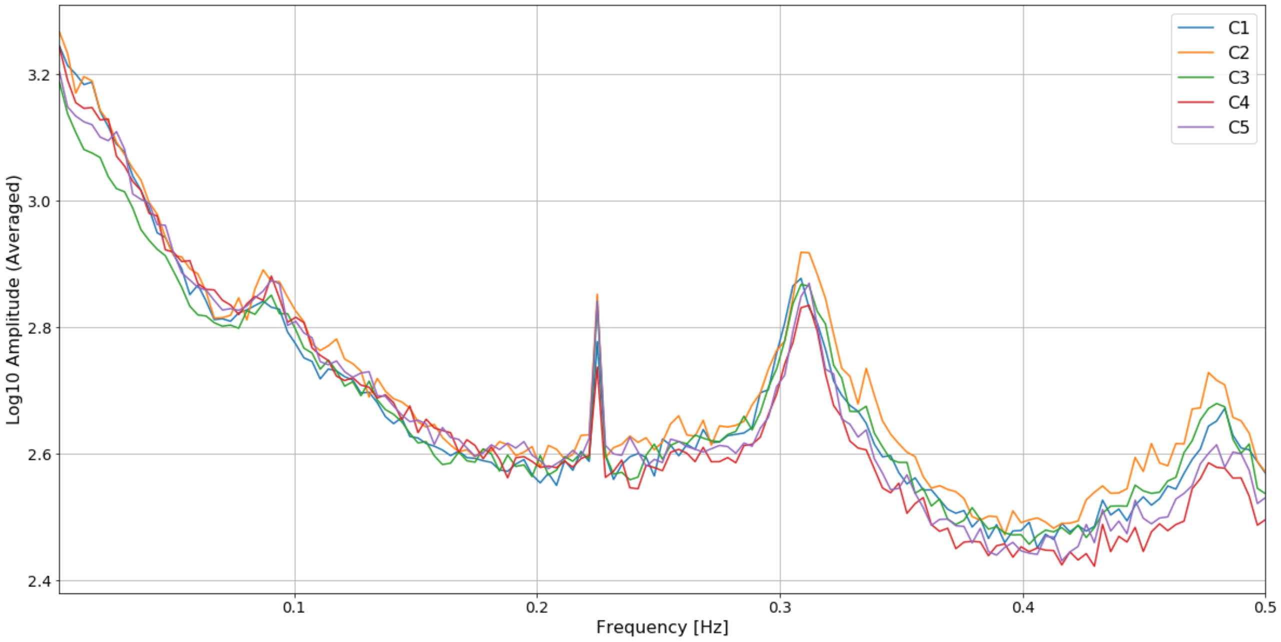

FIGURE 3: EXAMPLE OF FREQUENCY ANALYSIS ALGORITHM APPLIED TO 1HZ OPC DATA FROM 5 TURBINES.

Figure 3 shows the autospectra obtained by applying a Welch FFT algorithm to 1Hz OPC data from 5 turbines. Both the rotor mode (1P, around 0.22Hz) and the 1st tower mode (around 0.31Hz) are clearly visible. Good quality OPC data produce frequency analysis outputs which can be used to identify these frequencies and their amplitudes and track them over time. Shifts over time or across turbines can be indicative of issues such as rotor imbalance (due to pitch misalignment or significant differences between the 1st moment of mass of the blades) and softening of the foundations.

5.4 INSPECTIONS

Inspecting the condition of wind turbines is vital at various stages of a wind farm's lifecycle. It allows all stakeholders to reassure themselves of the quality of the turbines' fabrication, maintenance and performance. In some countries, regular inspections are a legal requirement; and many insurance companies insist on them when issuing a policy.

In general, inspections provide an opportunity for owners to obtain a detailed analysis of the condition of their wind turbines which can inform operational strategies. Moreover, inspections can help the identification of a problem that was not identified by other integrity monitoring approaches (e.g. loose parts, oil leakages, corrosion).

Typical milestones requiring inspections include:

- Manufacturing of individual components
- Wind farm commissioning
- Assessment of asset performance and condition during operation
- Damage investigations
- End of the turbine warranty period
- Extension of the wind farm's design life time

Inspections cover all components of wind power generation systems including the rotor, nacelle, tower, foundation and electrical system. Typical inspection procedures include:

- Visual assessment
- Lubricant sampling and analysis (gearbox oil, main bearing grease)
- Borescope inspection of gearbox, generator, main bearing
- Ground, rope-based or drone-based blade inspection
- Gearbox-generator alignment checks
- Generator and transformer electrical diagnostics
- Vibration analysis
- Thermography inspection
- Non-destructive testing (NDT) and foundation assessments

Depending on the turbine technology and operating conditions, some of these procedures can be of greater value than others. A wind farm's Operations and Maintenance (O&M) strategy specifies which inspection procedures should take place, and at which frequency, at each stage of the wind farm's lifecycle. Some of the procedures are mandatory (e.g. statutory inspections), whereas others are left to the discretion of the asset manager. In specifying the latter, the O&M strategy should aim, as much as possible, to optimise the trade-off between the cost of a given inspection procedure and the benefits that the procedure provides, which normally manifest themselves in a reduction of the probability of a failure (or degradation) taking place, or in a mitigation of the outcomes of that failure occurring.

5.5 COMPARISON OF DIFFERENT INTEGRITY MEASUREMENT APPROACHES

Table 1 below compares the different approaches to measuring integrity in terms of the components they can monitor, their capabilities to predict future failures (i.e. how long in advance they can identify a potential failure), their level of automation (i.e. how easily the system can work without any human intervention), their implementation costs and the potential to be integrated within a turbine and/or wind farm control system.

TABLE 1: COMPARISON OF DIFFERENT INTEGRITY MEASUREMENT APPROACHES

	Components monitored	Predictive capabilities	Level of automation	Cost	Integration w/ control
Turbine alarms	All turbine subsystems	Low	High	Low	High
CMS	Drivetrain Foundations Blades	High	Medium	Medium	Medium
SCM	Drivetrain Rotor Tower/foundations	Medium	Medium	Low	High
Inspections	All turbine components	Medium	Low	High	Low

6. MEASURING LIFE

6.1 NOMINAL LIFE VS. DESIGN LIFE

In Section 4.2 we explained how the definition of a turbine's design life is based on an assumed set of environmental conditions which are specified by design standards for a given type class. When designing a wind farm, in order to decide which type class is applicable to a certain site, the site-specific conditions (as recorded during a measurement campaign which generally relies on a meteorological mast) are compared to the conditions specified for the type class. This process introduces several sources of uncertainty, related for example to accuracy of measurements, to wind modelling, and to how the measurement period relates to a longer term, as well as to certain assumptions such as how much time the turbines are going to be spending in different operating states.

As a result, once a turbine is installed and commissioned, the average levels of structural loading which it experiences in the field can be different from those estimated at the design stage. In general, because of the conservative assumptions made during the design process, these levels of loading tend to be lower. This introduces the opportunity to operate turbines beyond their design life while still maintaining the reliability levels within the required thresholds.

When considering variations of a turbine's actual life from its design life, the expectation is that actual life will be curtailed by fatigue accumulation before it is curtailed by a failure caused by extreme loading events associated with a different environmental return period. Therefore, the fatigue limit state is assumed to be the primary limitation governing actual life, whereas extreme loads are often not considered in life estimation. It is important to note that this assumption does not hold when considering very large increases from the original design life.

With these considerations in mind, if a wind turbine is designed to a reliability level expressed as a nominal annual probability of fatigue failure of, for example, one in ten thousand (10^{-4}), then the nominal lifetime can be estimated as the number of years after which the expected annual probability of fatigue failure reaches 10^{-4}. Equally, a 20-year design life corresponds to an assumed level of accumulated fatigue loading on the main structural components of a wind turbine, so that a turbine's "life" and its "fatigue loading" can be used as proxies of each other.

Crucially, this means that fatigue load analysis can be used to estimate the amount of remaining (or depleted) life at any point during its operational lifetime. Estimates of turbine life can be used in asset management strategies (e.g. by extending operation of a turbine beyond its design life), in turbine and wind farm control (e.g. by designing algorithms that optimise the trade between energy capture and loading in order to achieve a lower LCoE). This chapter will describe methods to estimate remaining turbine life.

It should be noted that numerical life estimation analysis on its own is not enough to provide a comprehensive measure of life which can be use, for example, for the purposes of life extension. The DNV GL life extension standard [8] mentions that to pursue a practical life extension strategy, numerical results would need to be combined with practical methods such as targeted wind turbine inspections. However, given that the focus of this document is on uses of SCADA data, we will limit the discussion on the life estimation process to the numerical methods.

6.2 LIFE ESTIMATION PROCESS

The life estimation process involves three steps:

1. characterising (gathering or estimating) the site environmental conditions and the associated uncertainty
2. calculating site-specific and type class fatigue load using numerical turbine models
3. converting the ratio between site-specific and type class loading into an estimate of depleted (or remaining) design life, also with its associated uncertainty.

These steps are discussed in more detail in the remainder of this section.

ASSUMPTIONS WITHIN THE ANALYSIS

Depending on the level of detail of the analysis, various assumptions are applied. Commonly, these include, at minimum, the following:

- Turbines remain in good condition throughout their lifetime
- Turbines do not have structural/manufacturing defects
- Regular maintenance is performed according to the Original Equipment Manufacturer (OEM) guidelines/requirements

6.3 LIFE ESTIMATION INPUTS

The inputs required to carry out life estimation include environmental parameters and information on a turbine's models, control system and operational strategies. These inputs are necessary to characterise the operational history of the turbines, as well as the environmental conditions that they experienced on the site, both of which are intrinsically related to how much fatigue the turbine has accumulated, and in turn with its remaining life.

ENVIRONMENTAL AND OPERATIONAL PARAMETERS

Modelling the entire wind field around a turbine's rotor would require in theory the knowledge of a very large number of parameters. As explained in Section 4.2 however load calculations carried out according to certification standards are based on a limited set of site environmental inputs:

- Mean turbulence intensity at 15 m/s
- Standard deviation of wind speed standard deviation at 15 m/s
- Wind speed frequency distribution
- Wind shear
- Air density
- Flow inclination

These are generally referred to hub height. The operating state of the turbine also needs to be characterised. Input parameters related to the operating state of the turbine are:

- Number of hours spent by the turbine in different operating states
- Number of start-ups and shutdowns per year

PRE-CONSTRUCTION ASSESSMENTS

A relatively simple way to gather the environmental parameters required for life estimation is to extract them from the pre-construction assessments which summarise the results of the measurement campaigns carried out on a site when considering the development of a wind farm.

Data can be in the form of measurements given at mast locations, or of numerical results from wind modelling software. The uncertainty of these inputs is related to the accuracy of measurements, to the quality and representativeness of the wind modelling, and to how the measurement period relates to a longer term. Moreover, biases can be introduced when data are not measured at hub height, or when the distance between the location of the turbines and that of the masts used during the measurement campaign is significant.

INPUTS FROM SCADA DATA

SCADA data can be a valuable addition to pre-construction data, as they provide refined insight into the environmental conditions experienced by the turbine and into its operational state. The main benefit of using SCADA data is the fact that they provide historical information on the actual operational conditions of the turbine, thereby reducing uncertainty with respect to using a prediction of these conditions such as that provided by a pre-construction assessment.

Another advantage of using SCADA data is that they are produced in real time and can be used for ongoing life estimates (see Section 6.4) which can fit within a turbine or wind farm control framework.

Environmental inputs extracted from SCADA normally consist of wind speed and directional data measured by nacelle anemometers and wind vanes. Corrections need to be applied to remove the effects of turbine wakes in order to obtain free stream conditions. Turbulence intensity can also be estimated using wind speed mean and standard deviation data from the nacelle anemometer. A limitation of the SCADA measurements is that the quality of nacelle-mounted measuring equipment is significantly lower than that of the anemometers and wind vanes installed on a measurement mast, and the calibration process is less stringent.

A technique to overcome this limitation is to use the blade pitch, power and rotor speed average and standard deviation signals to estimate wind speed and turbulence intensity. This approach requires some knowledge of the turbine's control algorithm (and more specifically of its power curve) but generally it greatly reduces the uncertainty associated with estimating the wind conditions.

Some environmental inputs (e.g. shear, flow inclination) cannot normally be obtained from SCADA data and are derived from pre-construction assessments (or assigned standard values).

TURBINE MODEL

A numerical structural turbine model is required to estimate turbine loading based on the inputs specified in Section 6.3. Ideally, all calculations would make use of models of turbine-specific models provided by the turbine OEMs, who hold the most detailed information about their turbine models. This however is generally not a viable solution as OEMs are seldom willing to share information that contains significant amounts of intellectual property.

For this reason, the life estimation process often requires the creation of a model of the turbine. The level of detail of the model can vary from low (a generic turbine model which is broadly representative of the wind turbines under investigation) to high (a model that has been calibrated using site measurements in order to better capture the specific dynamics of the turbines). In general, a more detailed model will provide more accurate loading estimates but will also require more effort (and will therefore be more costly) to create.

When it is not possible to obtain or create a detailed model (e.g. because turbine-specific data are proprietary to the manufacturer) it may be necessary to rely on a generic model which can be used to represent a range of technologies. The loading and life estimates obtained with such a model will be characterised by a high level of uncertainty, often quantified through the application of pragmatic factors.

The accuracy of the analysis can be increased by making use of a turbine-specific model created using publicly available data: hub height, rotor diameter, blade dimensions and drivetrain technology are all information that are generally available; tower top mass and tower base diameter are sometimes provided in data sheets. Experience in turbine design experience can help turning a relatively limited set of design parameters into a representative model which will produce more accurate results and also enable more refined methods for calculating uncertainty (e.g. Monte Carlo analysis) as well as the calculation of additional loading outputs (e.g. foundation fatigue loads), neither of which could be undertaken with a generic turbine model.

The turbine model can be further refined using site measurements. This normally involves activities such as measuring wind turbine component dimensions, tuning the model to calibrate its outputs against load measurements, and validating the power curve. It should be noted that even extensive validation cannot reduce the uncertainty associated with specific areas of the model such as certain control system dynamics and local soil conditions.

CONTROL SYSTEM INFORMATION

Turbine loading response to environmental conditions is highly dependent on the control system of the turbine. A model of the control system that is more representative of the one implemented in the turbines under consideration will result in more accurate life estimates. Similar considerations to those made earlier in this section about turbine models apply to control system models.

In the ideal scenario the load calculations would make use of a controller model provided by the OEM. In reality, this will seldom be a viable solution, and a model of the control system will have to be developed along with the turbine model. The more detailed the turbine model, the higher the likelihood that the characteristics of the control system model will too be more representative of those of the control system implemented on the turbines.

The control system model too can be refined using site measurements (particularly generator torque and pitch angle histories, which can be used to characterise the speed control loops).

TURBINE DESIGN INFORMATION

The following design information is also required in order to translate loading results into a life estimate:

- Certification standard
- Certification wind and turbulence class
- Certification lifetime
- Wind turbine co-ordinates (used within the Frandsen wake calculation)

OPERATIONAL STRATEGIES

Different operational strategies can be adopted over time, for example to optimise operation or to satisfy environmental requirements and operational constraints. Some examples include:

- Curtailment (e.g. WSM, noise curtailment, shadow flicker, grid curtailment)
- Turbine uprating or downrating

- Deployment of power performance upgrades

These operational strategies have a direct impact on turbine loading and life and should therefore should be considered in life estimation analysis. Information about these strategies can be provided by asset managers or extracted from historical SCADA data (which however don't provide any visibility of future changes).

6.4 FATIGUE LIFE CALCULATION

In order to estimate the life of a wind turbine's structural components, site-specific fatigue loading is calculated and compared to type class fatigue loading. This involves running two sets of fatigue load calculations (one for the site-specific conditions and one for the type class ones) as defined by the turbine certification standards. These sets of calculations normally include different "load cases" which represent different operational scenarios such as power production, fault cases, normal starts, stops and idling.

This process can follow different approaches, depending on how and how many load calculations are run and on whether their outputs are interpreted and combined in a deterministic way or in a probabilistic one. Common to all approaches is the fact that the comparison between site-specific and type class fatigue loading is done using damage equivalent loads (DELs), that is constant loads that lead to the same damage as that caused by time varying loads over the same period. More information on how damage equivalent loads are calculated is given in Appendix A.

The DELs can be used directly to provide a measure of damage or they can be translated into a rate of accumulation of damage. The damage shows how much fatigue has been accumulated (or will be accumulated) during the period under analysis and can be converted into a percentage of design life, referred to as "DEL margin", by calculating the ratio of site-specific to type class DELs. The rate of accumulation of damage can be used to predict the lifetime of turbine as visualised in Figure 4.

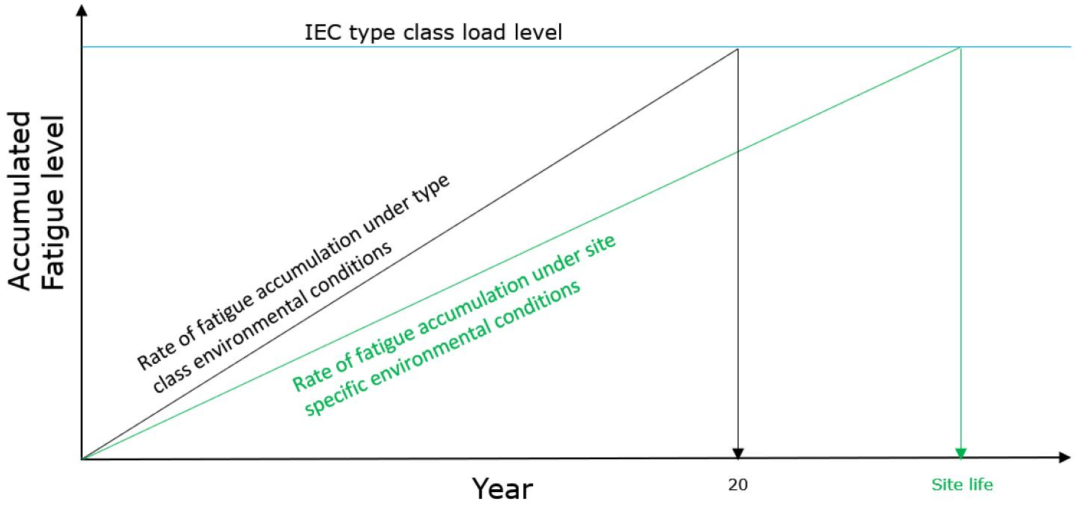

FIGURE 4: CONVERSION OF DELS INTO LIFE ESTIMATES

DIRECT SIMULATION VS LOADS DATABASE

A first possible difference in approach consists on whether each calculation in the load set is carried out running a full time-domain simulation with an aeroelastic model (direct simulation) or a loads database.

In this context a loads database is a collection of pre-calculated unweighted damage equivalent loads, created by running multiple sets of fatigue load sets, as defined by one of the turbine design standards, with varying input parameters. Each set of fatigue loads is run for a range of values for each of the user defined inputs (*i.e.* the external conditions) and the results are compiled into a multi-dimensional database by postprocessing the outputs, extracting the DELs for the main wind turbine load component locations, and collecting them into a multi-dimensional matrix.

DNV GL normally uses six-dimensional load databases with varying levels of wind speed, turbulence intensity, air density, flow inclination, wind shear. Stochastic effects can also be accounted for by running simulations with multiple wind seeds for each combination of external conditions and averaging the resulting DELs. Figure 5 provides a visualisation of a hypothetical 3-dimensional loads database containing DELs for four load component locations. More information on the fatigue loads database tool developed by DNV GL is provided in Appendix B.

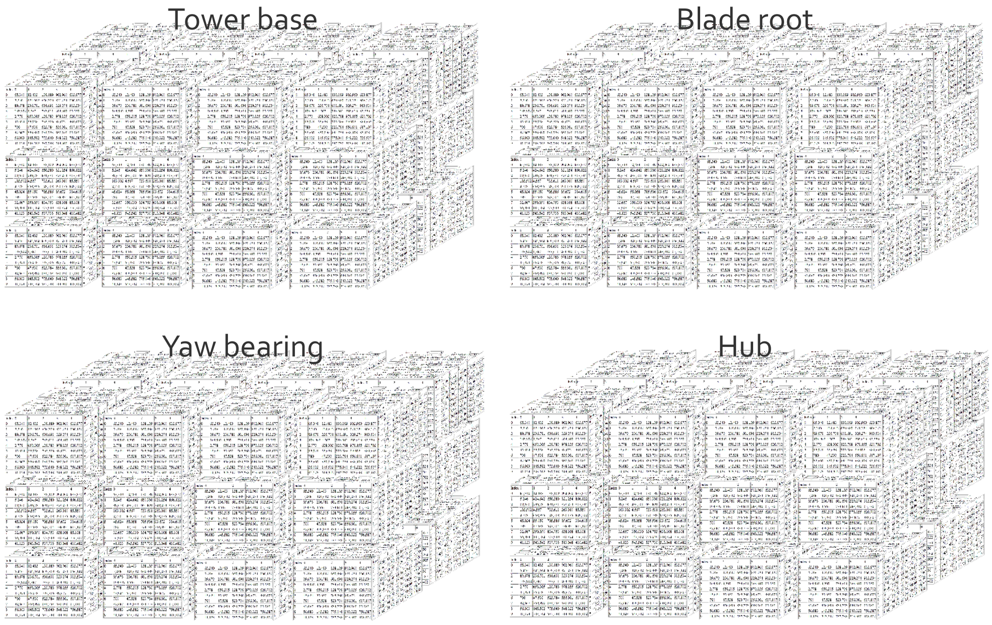

FIGURE 5: EXAMPLE OF LOADS DATABASE WITH 3 INPUTS.

A loads database can be interrogated, using a set of inputs which correspond to the estimated site-specific conditions, to extract DELs for the different load locations. Because the database is only created for a range of input conditions which vary in discrete steps, the extraction process generally involves an interpolation of the outputs. This introduces a level of approximation with respect to a direct simulation. The creation of a loads database is also a more onerous process than running a set of direct simulations.

The advantage of the loads database approach is that, once it has been created, it provides a relatively inexpensive way of calculating DELs for a given combination of inputs, which enables approaches which involve a large simulation of inputs, such as the probabilistic analysis and the ongoing estimate discussed in the next sections. While in theory these approaches are compatible with direct simulation, the computational load associated with it would not make it a practical solution.

DETERMINISTIC VS PROBABILISTIC ANALYSIS

One of the essential elements of life estimation is the determination of the associated uncertainty. This can follow two approaches, both of which are defined in certification standards.

In a deterministic approach, the load calculations are run for the site-specific and type class conditions extracted from a pre-construction assessment, from SCADA data or from a combination of the two (see Section 6.3). Where these inputs are characterised by a distribution (as in the case for example for wind speed) the outputs from the load calculations are weighted by the same distribution to obtain lifetime weighted DELs. Wake effects and any operational strategies such as WSM should also be taken into account when combining the DELs. Any uncertainty associated with the environmental inputs and with the models used for the simulations is then translated (normally through a sensitivity analysis) into an uncertainly that is applied to the final outputs. Deterministic analysis is compatible with both the direct simulation and loads database approaches discussed above.

In a probabilistic approach, uncertainties - both in the inputs and the methodology itself - are taken into account directly. The distributions describing the uncertainty in the site conditions and those describing the numerical modelling uncertainty are sampled at thousands of times in a Monte Carlo framework and used as inputs to the load calculations to generate a discrete distribution of loading outputs, which is then fitted with different mathematical distributions to establish a ‘best-fit’ continuous DEL distribution. This DEL distribution can then be transformed into a probability distribution of lifetime, from which an expected lifetime value can be obtained.

Because of the large number of load calculations required to model a probability distribution for each individual input, the probabilistic approach is more suitable to a loads database than to direct simulation.

ONE-OFF VS ONGOING ANALYSIS

Another possible difference in approach to life estimation stems from whether the estimate is based on a single set of load calculations, based on inputs (or input distributions, in the case of probabilistic analysis) intended to represent the best estimate of the long-term average site conditions, or on individual load calculations, each one of which models the site conditions at a specific point in time.

A one-off analysis requires, as a first step, the determination of the long-term average site conditions. If the input conditions are extracted from a pre-construction assessment this process is straightforward, as the conditions are already provided as central estimates within a long-term context. If the input conditions are extracted from operational data (with the advantages and disadvantages discussed in Section 6.3) then the 10-minute values collected through the operational history will have to be aggregated and, particularly where the operational period is relatively short (indicatively below 5 years), adjusted so that they are representative of the long-term. This is normally achieved by correlating site measurements with concurrent values from an appropriate source of long-term reference, often reanalysis data such as MERRA-2[1] and ERA5[2].

One of the advantages of using SCADA data for extracting the site-conditions inputs is that this can be done on an ongoing basis. Using this approach, every time a new data record becomes available, a DEL for that period is produced through deterministic or probabilistic load calculations. The DELs from each 10-minute period are accumulated over time to provide an ongoing estimate of the turbine's depleted (or remaining) life which is automatically updated. While in theory this approach is compatible with both direct simulation and use of a loads database, in reality the latter is particularly suitable because of its relatively low processing overhead, particularly where each 10-minute DEL is calculated via probabilistic analysis.

The ongoing analysis approach has the limitation, with respect to a one-off analysis, that it is not possible to calculate a long-term adjustment. It has however the advantage that estimates are updated in near real-time and in a way which gives consideration to the actual site-specific conditions experienced by the turbines and of their operational strategies.

6.5 ESTIMATING UNCERTAINTY

Regardless of the specific approach adopted to calculate the accumulated fatigue loading, an essential step in estimating turbine life is the estimation of the uncertainty. The main sources of uncertainty are associated with the environmental inputs and with the parameters which characterise the turbine model.

The uncertainty associated with environmental inputs is a result of varying quality of measurements, modelling and assumptions. In addition, there is a natural variability associated to the parameters that needs to be accounted for.

The numerical models used to calculate loading are also subject to a degree of uncertainty. Limitations in the simplified engineering relationships that are used to represent a complex system, as well as the difficulty in estimating all the model parameters, lead to uncertainty in the simulated results provided by the numerical models. The use of generic representative wind turbine models (see Section 6.3) also introduces additional uncertainty.

Depending on whether the loads calculations are carried out following a deterministic approach or a probabilistic one, two approaches exist to translate the uncertainty associated with environmental inputs and models into output uncertainty.

In deterministic analysis, inputs and model parameters are assigned a range of variation (generally based on a combination of measurements and experience), which is then translated into output uncertainty through a sensitivity analysis.

---

[1] The Modern Era Retrospective-analysis for Research and Applications, Version 2 (MERRA-2) data set has been produced by the National Aeronautics and Space Administration (NASA) by assimilating satellite observations with conventional land-based meteorology measurement sources using the Goddard Earth Observing System, Version 5.12.4 (GEOS-5.12.4) atmospheric data assimilation system [9]. The analysis is performed at a spatial resolution of 0.625° longitude by 0.5° latitude.

[2] ERA-5 is the fifth generation of the European Centre for Medium-Range Weather Forecasts' (ECMWF) atmospheric ReAnalyses of the global climate [10]. ERA5 incorporates vast amounts of historical measurement data, including both satellite-based, commercial aircraft, and ground-based data to produce a description of the state of the atmosphere, including wind speed. Hourly analysis fields are available at a horizontal resolution of 31 km and include wind data at 100 m above ground level, as well as surface air temperature and air pressure.

In probabilistic analysis, inputs and model parameters are assigned a probability distribution, which is propagated onto the outputs using a Monte Carlo approach.

In a one-off analysis, uncertainty will only be calculated once and then applied to the final fatigue (or life) outputs (i.e. the central estimates).

In an ongoing analysis, the uncertainty will be calculated for each 10-minute record and then accumulated over time as shown in Figure 6.

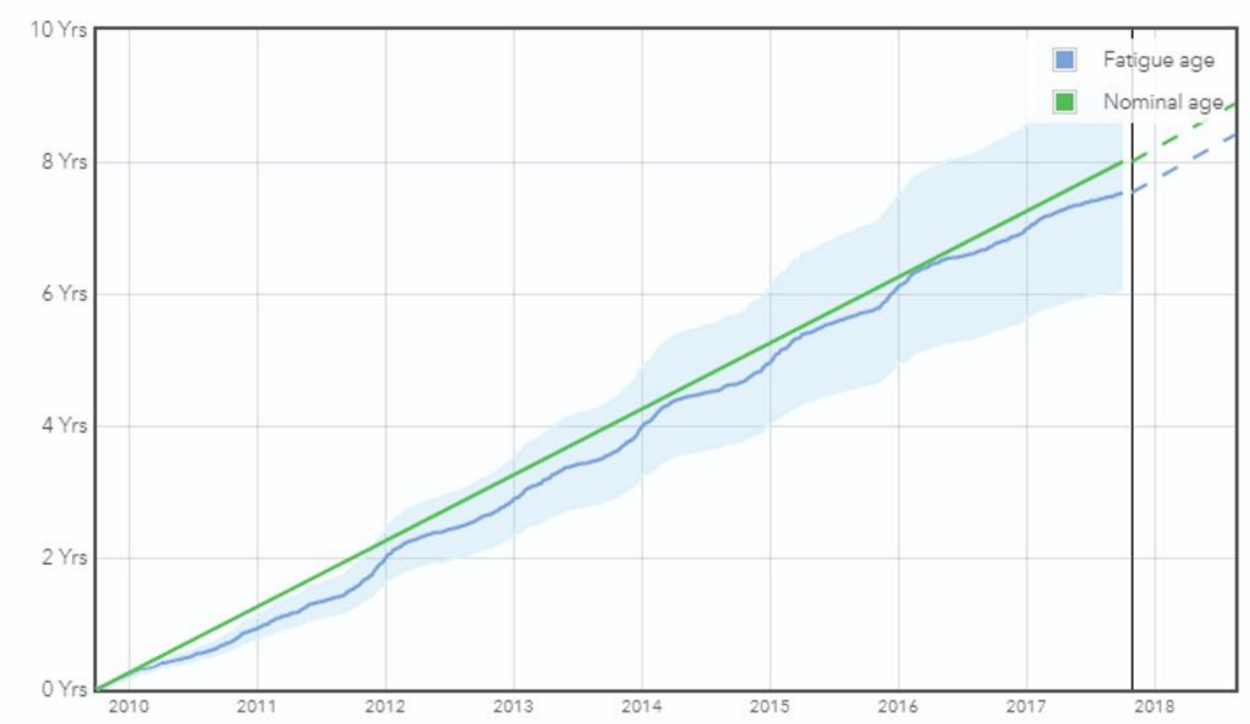

FIGURE 6: UNCERTAINTY ASSOCIATED WITH LIFE ESTIMATE PRODUCED THROUGH ONGOING ANALYSIS. UNCERTAINTY IS REPRESENTED BY THE BLUE SHADED AREA.

7. TURBINE CONDITION AND OPERATIONAL STRATEGIES

We will now explore how the estimates of integrity and life obtained through the approaches discussed in Section 5 and Section 6 can be used to design an operational strategy, whether it is an asset management one, implemented by humans, or a turbine or wind farm control one, implemented in digital systems.

7.1 ASSET MANAGEMENT STRATEGIES

Asset management strategies can be informed by the outputs of integrity and life measurements. Knowledge about the condition of a turbine can be used to explore trade-offs between different performance indicators such as energy production, life, failure rates and inspection and maintenance costs.

TURBINE INTEGRITY AND ASSET MANAGEMENT

Knowledge about turbine integrity can be used to optimise inspections and to reduce O&M costs.

A basic strategy is to deploy CMS and/or SCM analysis to predict failures. The ability to identify patterns that are indicative of incipient failures results in a reduction in the downtime and costs associated with component replacements: the time to source spare parts and to mobilise maintenance teams is greatly reduced. Moreover, inspection and maintenance costs can be reduced by focussing some of the activities only on turbines that are identified by CMS/SCM as having a higher risk of failure.

More advanced strategies can involve trading the probability of a failure with inspection and maintenance costs and with energy production. Such strategies require that the CMS/SCM analysis provides not only a binary indication of potential failures (“at risk” vs. “not at risk”) but also an indication of the probability of failure and of the impact that uprating/derating a turbine has on it. The asset managers can then decide, based on parameters such as inspection, maintenance and component replacement costs and revenue per MWh, to reduce (or increase) energy production and/or time between inspection and maintenance activities in order to reduce (or increase) the probability of a failure.

TURBINE LOADING, LIFE AND ASSET MANAGEMENT

Information about turbine loading and life too can be used to optimise asset management strategies. Like in the case of turbine integrity, a range of strategies are available to capitalise on turbine life estimates.

A basic strategy consists in targeting inspections on turbines which are subject to the highest levels of loading. This can have a significant impact on O&M costs, particularly for offshore wind farms and larger onshore wind farms.

Another basic strategy, where life estimates suggests that the lifetime of the turbines is significantly longer than the design lifetime, is to update financial models to consider an increased operational period. As an example, in January 2019 the publicly listed asset management company Greencoat UK Wind announced to the London Stock Exchange that it was changing the life assumption of its wind farm portfolio from 25 years of life to 30 years, based on the results of a life extension analysis carried out by DNV GL [11]. This resulted, within 24 hours, in an increase of approximately 5% of their market value.

More advanced strategies can involve trading fatigue loading (and therefore life) with energy production. Such strategies require the ability to model the impact that different operational strategies (uprating, derating, application of WSM) have on energy production and on fatigue life, as well as the impact that extending or reducing life has on the overall LCoE. This enables the selection of an optimal strategy.

The relatively long decision-making times associated with asset management strategies mean that these are only modified on a relatively infrequent basis. This in turn means that the information which supports the decision process can be updated only when needed, which is compatible with both the one-off analysis and the ongoing analysis discussed in Section 6.4.

7.2 TURBINE CONTROL AND WIND FARM CONTROL

The trades discussed in Section 7.1 can be explored not only to optimise a human-supervised asset management strategy, but also a control algorithm implemented in the turbine or wind farm controller.

The ability to consider integrity and life estimates when designing a turbine controller extends the space over which its algorithms can be optimised, increasing the potential to reduce the overall LCoE.

The outputs of CMS/SCM analysis can feed directly into the supervisory logic, so that turbines at elevated risk of failure are stopped, or de-rated.

The outputs of life estimation can be used in two ways. The results of a one-off analysis can be used to optimise the design of the turbine controller, using the knowledge about the site-specific rate of fatigue accumulation to increase (or decrease) energy production setpoints, either across the board or for specific ranges of environmental conditions and operating scenarios. Strategies like WSM can also be optimised, e.g. by reducing them based on the fact that site conditions are more benign than type class conditions.

If it is possible to implement an ongoing life estimation algorithm within the turbine controller, then its outputs can be used to regulate the setpoints dynamically in response to the estimated turbine loading. This has the benefit, with respect to using the results of a one-off analysis, of removing any approximation associated with considering a single central estimate of the site conditions to optimise turbine operation. The site-specific conditions (and resulting loading) are estimated for each control cycle and the control outputs are adjusted accordingly. Because of the relatively fast rates of execution of turbine control cycles (normally between 10-100Hz), such a strategy would not be based on 10-minute SCADA but rather on the signals used by the controller.

Turbine control design is normally based on the “selfish” objective of maximising a benefit function for a turbine, considered as an individual machine operating in isolation. Turbine wakes however cause other turbines operating downstream of the controlled turbine to experience a lower wind speed, causing a reduction of power output, and a higher level of turbulence, resulting in higher mechanical loads. There may also be an increase in asymmetrical loads due to velocity gradients across a downstream turbine rotor when it is partially immersed in a wake.

Because of the importance of wake effects in reducing energy production and increasing loads, there has been much interest in recent years in the concept of minimising wake effects through wind farm control [12]. Instead of allowing each turbine to behave as designed, i.e. selfishly, to achieve the best combination of energy production and loading for itself, the concept is that the wind farm controller commands changes to the operation of individual turbines in order to achieve the optimum performance for the wind farm as a whole [14].

The main approaches to wind farm control are:

- Traditional sector management, where some turbines are switched off to avoid wake effects on closely-spaced turbines [13]. The production from some of the turbines is lost but loads are reduced on all turbines, and overall production may increase.
- Induction control, where the power setpoint of some turbines is reduced to weaken wake effects on downstream turbines. The production from some of the turbines is reduced but loads are reduced on all turbines, and overall production may increase.
- Wake steering, where the yaw misalignment of some turbines is regulated in order to deflect wakes away from downstream turbines [15][16]. The production from some of the turbines is reduced and their loading may increase, but loads are reduced on downstream turbines, and overall production may increase.

Life estimation fits the wind farm control framework very well. Wind farm control can be designed to reduce loading by distributing it across turbines, and more in general to optimise the trade between the overall lifetime of a project and its lifetime energy yield.

As in the case of turbine control, the outputs of life estimation can be used in two ways. The results of a one-off analysis can provide a useful indication of the potential benefits of wind farm control and, at a later stage, they can be used to optimise the wind farm control design setpoints.

Furthermore, the outputs of an ongoing life estimator can feed directly into the wind farm control algorithm, so that the production and yaw setpoints of all individual turbines are regulated dynamically to achieve optimal LCoE at wind farm level. Because of the relatively slow execution rate of a wind farm controller (when compared to the individual turbine controllers) an ongoing life estimator based on 10-minute SCADA is suitable for operating within a wind farm control framework.

8. CONCLUSIONS

We define a wind turbine’s condition as a combination of its integrity, i.e. its ability to operate safely and effectively, and of its life, i.e. its ability to withstand fatigue loading. The quantification of these two metrics can follow several approaches.

In this document we explored and compared the main approaches to monitoring integrity, including turbine alarms, Condition Monitoring Systems (CMS), SCADA-based condition monitoring (SCM) and inspections, and the main approaches to measuring life. We also discuss how the outputs of these types of analysis feed into the optimisation of asset management strategies and of turbine and wind farm control algorithms to reduce the overall LCoE.

Some of these monitoring and measurement approaches are only possible when operational SCADA data are available. SCADA data, which are inexpensive to access and do not require the installation of additional sensors, enable:

- The deployment of signal trending and, if higher frequency data are available through OPC or IEC interfaces, or frequency tracking algorithms. These provide some of the predictive functionality of CMS at a fraction of the cost.
- The estimation of site-specific conditions and turbine operational information based on historical data rather than pre-construction assumptions. This provides a better measure of the conditions that have been experienced on site as it removes the uncertainty associated with wind modelling.
- Ongoing estimate of turbine loading and life, which can be used in a turbine control and wind farm control framework.

9. REFERENCES

[1] IEC 61400-25-1 Part 1 to 6: Communications for monitoring and control of wind power plants

[2] Knudsen, T., T. Bak and M. Svenstrup (2015). 'Survey of wind farm control: power and fatigue optimization'. *Wind Energy* 18.8:1333–1351

[3] UK Health and Safety Executive Report, 'Key Programme 3 Asset Integrity Programme: A report by the Offshore Division of HSE's Hazardous Installations Directorate', 2007

[4] IEC 61400-1 edition 4, 2019, Part 1 : Design requirements

[5] Germanischer Lloyd, Rules and regulations, IV - Non-marine Technology, Part 1 – Wind Energy, Regulation for the certification of the Wind Energy Conversion Systems. Chapter 1 – 10, 2010

[6] DNV GL Service Specification, DNVGL-SE-0441, Type and component certification of wind turbines

[7] Tchakoua, P. et al., 'Wind Turbine Condition Monitoring: State-of-the-Art Review, New Trends, and Future Challenges'. Energies 2014, 7, 2595-2630

[8] DNV GL Standard, DNVGL-ST-0262, Lifetime extension of wind turbines

[9] National Aeronautics and Space Administration, MERRA-2, MDISC website: http://disc.sci.gsfc.nasa.gov/mdisc/,, MERRA-2 tavg1_2d_slv_Nx: 2d,1-Hourly,Time-Averaged,Single-Level,Assimilation,Single-Level Diagnostics V5.12.4 (M2T1NXSLV), 1980-present.

[10] The European Centre for Medium-Range Weather Forecasts (ECMWF); ERA5; website: https://www.ecmwf.int/en/forecasts/datasets/archive-datasets/reanalysis-datasets/era5

[11] Keir Harman 2019, 'Why investors should now be banking on longer life', website: https://www.windpowermonthly.com/article/1579890/why-investors-banking-longer-life

[12] Ervin Bossanyi 2018, 'Combining induction control and wake steering for wind farm energy and fatigue loads optimisation', Journal of Physics: Conference Series

[13] Bossanyi, E. and Jorge, T. (2016) 'Optimisation of wind plant sector management for energy and loads.' *IEEE European Control Conference*

[14] Vali, M., Vollmer, L., Petrović, V. and Kühn, M. (2017) 'A closed-loop wind farm control framework for maximization of wind farm power production'. *Wind Energy Science Conference*

[15] Dijk, M. T., van Wingerden, J.-W., Ashuri, T. and Li, Y. (2017) 'Wind farm multi-objective wake redirection for optimizing power production and loads'. Energy 121:561 569

[16] Grant, I., P. Parkin and X.Wang (1997). 'Optical vortex tracking studies of a horizontal axis wind turbine in yaw using laser-sheet, flow visualisation'. *Experiments in Fluids* 23.6:513–519

10. APPENDICES

10.1 APPENDIX A

Damage equivalent loads are used to equate the fatigue damage represented by rainflow cycle counting data to that caused by a single stress range repeating at a single frequency. The method is based on the Miner's rule. The damage equivalent stress is given by the following formula:

$$L_N = \sqrt[m]{\frac{\sum L_i^m n_i}{N}}$$

where

L_N	is the equivalent stress for N cycles
L_i	is the stress range bin i.
n_i	is the number of rain flow cycles at stress range bin i.
m	is the negative inverse of the slope on the material's Wöhler curve (also referred to as the S-N curve slope).
N	is the number of cycle repetitions in the turbine lifetime.

The S-N curve slopes (m) are integer values where 4 typically represents steel and 10 typically represents glass reinforced plastic.

The stress, L_i, depends upon the geometry of the structure under consideration. It is assumed that stress is proportional to load, therefore it is acceptable to use load instead of stress in the above equation.

10.2 APPENDIX B

The Site Suitability Tool (SST) is a loads database tool developed by DNV GL to provide an efficient solution to the challenge of site suitability assessments by enabling site-specific loading to be evaluated at minimal cost. It was developed based on decades of experience in turbine design, turbine numerical modelling, aero-elastic software development (Bladed) and site suitability expertise.

The SST derives the site-specific environmental conditions at each turbine location (including WSM and neighbouring wake effects using the Frandsen effective turbulence methodology as per IEC 61400-1 Annex D), and it generates the site-specific fatigue loading at all main turbine load component locations through linear interpolation of the values within the fatigue loads database. It then compares this with type class loads also extracted from the database.

The SST includes an intuitive user interface (shown in Figure 7) which provides quick and easy access to visualisations and tabulations of loading that facilitates comprehensive understanding of site-specific loading across a project taking into consideration ambient conditions combined with the effects of neighbouring wakes and WSM.

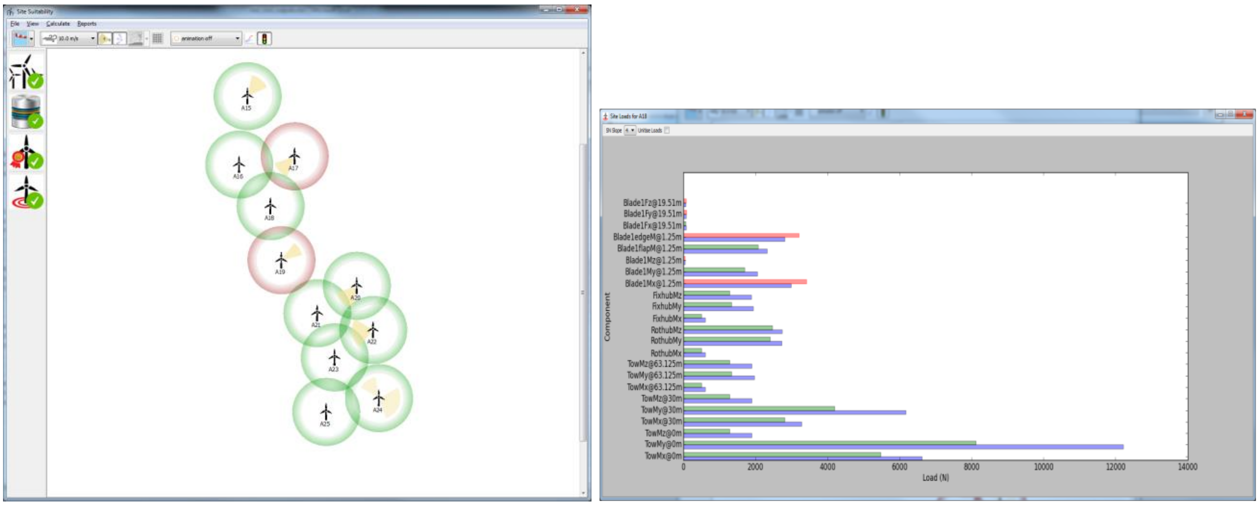

FIGURE 7: SST ‘TRAFFIC LIGHT’ VISUALISATION OF SITE-SPECIFIC LOADING VS TYPE CLASS LOADING AND SITE-SPECIFIC VS TYPE CLASS ‘BAR CHART’ OUTPUT – ALL DATA ARE ALSO REPORTED IN SPREADSHEET FORMAT.