

Analysing Offshore Wind Turbine Stop Events

Enabling Better O&M Decision-Making

Alistair Lee | May 2019

Introduction

The substantial and varied data streams generated in offshore wind to analyse turbine behaviour offer a significant untapped opportunity to enable better-informed operations and maintenance (O&M) decision-making strategies.

A detailed understanding of what data is available and the insights it can provide can be supplemented by a thorough and structured data analytics approach to provide significant added value to wind turbine and wind farm operators.

There are several challenges and perceived barriers around unlocking this value but the potential results make these worth overcoming.

This case study highlights the benefits of a structured and thorough approach to analysing wind turbine stop events for the purpose of establishing whether stop behaviour was anomalous or potentially damaging for the asset. Data was gathered from one turbine at an operational offshore wind farm in the UK that has been anonymised at the owner's request.

Key Findings

- Defining a clear scope is integral to delivering conclusions and outcomes in an efficient and structured way.
- Effective assessment of data quality is integral to producing accurate, high-quality analytics.
- Understanding all available data and interpreting the potential insights provided by this data is key before an analysis can be undertaken.
- Developing closer connections between analytics and operational staff can give a more rounded understanding of data sources available and potential insights to be gained from these sources.
- While having access to the greatest-functionality tools is beneficial to providing quality outputs, valuable analyses can be undertaken by using widely-available software tools.
- Visual and summarised outputs are useful tools in presenting granular, specialist information to a diverse audience.

Situation

Industry Background

Wind turbines generate vast quantities of data from a variety of sources. This data is generated with the intention of supporting and informing O&M strategies to maximise power production while maintaining key component integrity and reliability. While some subsets of this available data are routinely utilised by wind farm operators, the majority of data remains under-utilised and as such the full potential to inform asset owners remains locked.

One such category of available data is the turbine's alarm logs. Alarms are event indicators that are triggered to inform operators of specific events. For example, "gearbox oil high" is an alarm that triggers a shutdown of the turbine if the gearbox oil temperature sensor exceeds a defined threshold. The records of these alarms, known as alarm logs, are therefore treated as collections of indicators to support maintenance troubleshooting and are not typically analysed for trends and insights to inform long-term operational approaches and strategies. It is often the case that the full potential of this information source is untapped. This lost opportunity in utilising alarm logs to their full potential is common to exploiting many of the data sources available from individual wind turbines and whole wind farms.

Deriving maximum operational value from these data sources is a challenge for the wind industry. However, the potential for the industry to develop its approach to data analytics to realise significant rewards is underlined by the success of similar developments in more mature industries. Successful analysis of big data is commonplace in the automotive and aviation industries for a variety of purposes. An example of such success is the aviation industry's approach to improving spare parts management. By using collaborative data analytics across various operators to support process innovation, the industry has been able to improve spare parts availability by 30%.

While the opportunity for the wind industry is substantial, significant challenges remain if this potential is to be unlocked. For example, data structures from turbine systems are non-standard across different original equipment manufacturers (OEMs) and turbine models, meaning approaches to data analytics must also be non-standard. The lack of understanding of what data is available, and the insights that can be gained from analysing this data, creates uncertainty as to the commercial sensitivities that could be uncovered by data sharing. This sensitivity is heightened by uncertainty over data ownership between wind farm owners and turbine OEMs, and the pre-existing commercial sensitivities of these relationships.

Despite these challenges, the opportunity for the wind industry to gain increased insight from data analytics could be substantial and it is clearly acknowledged by many industry players. The challenge of how to unlock this potential is worth exploring.

The Challenge

Background

There are several potential barriers to successfully extracting the maximum potential value from turbine alarm log data. Like many others, the wind industry struggles with actual or perceived commercial constraints around the data generated by operational assets. Generally, turbine manufacturers hold a significant amount of power in this respect. Having designed the turbine and its associated data systems there may not always be full transparency of how data are recorded or categorised to the operator, or, for example such understanding can be lost following events such as staff turnover or ownership transactions. While it is reasonable for equipment manufacturers to seek to protect their intellectual property, this can become conflated with commercial confidentiality that may exist for contractual reasons – for example to prevent the release of information to third parties. It is relatively common for grey areas or mismatched expectations in the level of confidentiality to cause tension between an owner/operator and equipment manufacturer. This can create an artificial barrier to understanding and analysis of data generated by the assets including alarms and events.

The offshore wind industry is distinct from conventional power plants in both the nature and operating dynamics of the plant. This means that – as opposed to a conventional power station where a large amount of generating capacity is centralised or working in close proximity – the assets at an offshore wind farm are serially-produced individual systems that are spread over a large geographic area. In addition, assets can be in a variety of states concurrently that are distinct to the rest of the assets at the site. The resource itself is also inherently variable and hence output is expected to fluctuate significantly in short periods of time. All of these features can mean that the practices and ethos driving management of O&M (including data analysis), that have evolved from those used in conventional power plants, may not be optimally suited for offshore wind power. If maximum value is to be extracted from deeper analysis of data generated by assets, including alarm and event logs, then the potential value in these datasets should be communicated from leadership teams downwards. If value is placed on obtaining additional insights through analysis by the management team, it is likely that operational staff will find ways to do so. Empowerment and a sense of ownership of this untapped value potential are key ingredients to unlocking the potential of these datasets.

The type of skills required in offshore wind O&M may also be changing or at least broadening in scope. A particular area that is likely to see considerable expansion in the coming years is the application of new skillsets to this industry. Of particular relevance to this subject matter, it is expected that integrating capabilities, skills and experience from data science could increase the likelihood that an operator is able to extract insight from alarm and event log datasets. This may also increase the efficiency of the data analysis processes. Approaches that enable O&M staff to cross-skill and share experiences have been trialled in various guises by a variety of wind farm operators and can be a vital part of developing a rounded team that brings both relevant “domain expertise” of the equipment in question, and skills in data manipulation and analysis that can quickly and effectively derive insights.

It is also the case that operators with a centralised data analysis function may be missing vital hands-on or point-of-work experience at the time when analysis is conducted. As such, it is recommended that analytics is closely linked and encouraged to be in regular communications with operational staff, as opposed to periodic or “one-way” reporting. Site teams and engineers are likely to best understand the problems and challenges, but it is equally likely that data expertise will be most capable of conducting analysis to provide insights and solutions.

Specifics

Considering specifically the challenge and opportunity associated with alarm and event analysis; it is evident that while a large amount of data is being generated, and in many cases is used for its primary intended purpose of fault investigation, significant additional remaining value in this data is not always capitalised upon through analysis and insight.

When an analyst is presented with alarm and event data there are several fundamental points that should be considered, including:

1. How turbine stop events are identified in the data.
2. Whether stop events are recorded in isolation or form part of or are linked to a larger dataset.
3. Which categories are used to group stop events?
4. What documentation accompanies the dataset – for example, is a signal list or sensor map available?
5. Related to the above, are stop events well-defined, and are the limits of each categorisation clear and understood?

In the event that all of the information listed above is known and well understood by the team tasked with analysis, it will then be necessary to define the objectives of the proposed analysis. Having a well-defined problem statement, a written question to answer, and access to sufficient domain and equipment expertise will be essential if analysis is to yield maximum benefit.

What to look out for

There are a variety of potential pitfalls in any data analysis, including analysis of alarm and event data. General good practice can play a role in ensuring that any conclusions reached are sufficiently robust. For example, the time allowed for the analysis process should be sufficient, the analysts tasked should clearly understand the objective of the task they are being given, and should also be confident in the deadline by which results must be produced. Similarly, there should be an expectation placed upon analysts that they should provide regular feedback on progress and should flag any potential challenges, barriers or unknowns at an early stage, allowing additional resource, budget, time or other support to be brought in. The earlier a challenge can be identified and communicated to management, the more likely it will be that necessary support can be put in place ahead of the deadline.

Some observations that should be highlighted given the subject of this case study are discussed here. Firstly, the categorisation of alarm and event data may have been determined by the turbine OEM, or one of their suppliers. The decisions made in determining how to classify events may not be transparent to the person charged with analysis. For example, if turbine stop events are grouped into categories, it may not be clear to the analyst what the definition of these categories are or how the decision to classify events has been programmed. For example, if a shutdown event is categorised as "Grid" it may appear plausible to assume that any events in this category have been caused by events on the electrical network not in the control of the operator. In fact, it may be the case that sensors monitored by the control system have altered sufficiently to initiate a shutdown because a turbine component has influenced electrical behaviour. Unless there is a detailed understanding of what drives the classification of events, analysis may be incomplete.

Further, the flow and interaction of events with one another can cause misunderstanding. For example, in some cases the first alarm or event that is logged will be recorded as the cause of the entry in the log, whereas in others it may be that system logic records the highest priority, or most urgent alarm or event, as the cause of a turbine shutdown. Because it is likely that there will be significant interplay and interaction between events when a shutdown does occur, it is necessary to have some understanding of how these classifications are decided upon. Ideally, the entire logic model or decision tree for all events should be accessible to the analyst; however, accepting that for some of the reasons discussed earlier in this paper it may not be possible for the analysts to access this level of information, they should nonetheless have at least an understanding of what principles have been used to classify events. The risk if this understanding is not available is that implausible or incorrect conclusions will be reached.

It is very easy to assume that, because something appears as a numerical value or as timestamped data, that it is a representation of fact and is irrefutable and the only potential explanation. Since one event is highly likely to trigger another, if the sequence and interconnectedness of individual events is not understood then it may be that several events triggered by the same failure or subsystem may be categorised differently – an outcome that might not be identified in analysis conclusions.

Repeating an analysis for different time periods, peer review and checking, as well as drilling down and looking at detail – particularly outliers in results – are all methods that can help to increase confidence in findings.

The Approach

Pre-defined bins may pre-define limits of analysis

As discussed, the analyst may be challenged by the definition of categories, or bins, that have been predefined. They may either not fully understand the rules that define the categories used to group events, or if they are understood then such groupings may not meet the needs of the analysis proposed. It is advisable to study carefully and collect as much information as possible about the bins or categorisations used to group or classify alarms and events.

In some cases, it may be that several different systems constitute distinct parts of the chain that ultimately ends in the dataset the analysts have access to. For example, a subcomponent may have its own controller or some element of data logging and monitoring onboard that can be approached and interrogated as a more detailed subset of the whole turbine SCADA data. Moving up a level, higher-frequency data from a turbine controller will likely integrate some (but not necessarily all) of the data generated by turbine subsystems. How this data is then down-sampled and used to report lower-frequency SCADA data, alarms and events may not be defined, and again some analysis using higher time frequency data may prove illuminating, even if these are just occasional spot checks. The potential value of “go and see” should not be underestimated: while not always appropriate or possible, making an effort to expose operational staff to analysis and analysis staff to the physical hardware that is generating data should not be underestimated.

In some cases, it may be appropriate to collect all data routinely, including high time frequency and concurrent datasets from subsystems, and then analyse everything from all data sources together in a repeatable routine. In other situations, a small sample or individual delve into the detail may serve to validate, reinforce or answer questions about how the alarms and events that are to be used have been classified. This may then enable the analyst to revert to using the simpler alarm and event data on its own and allow others to conduct more straightforward analyses in future.

Consider bin definitions that are aligned to answering the question

To answer some of the questions that can be unlocked by alarm and event analysis, it may be necessary for the analyst to define their own set of categories or event bins. For example, should high-energy stop analysis (with the associated implication on asset integrity and remaining fatigue life of the asset) be of interest, it may be necessary to use the low detail alarm or event log to identify times where relevant events may have occurred – but then use high frequency data to determine the conditions, such as rate of change, that are associated with that stop. It may then be appropriate to count or analyse not only based on the event log but informing the classification of events. With scrutiny of high frequency data, studying how quickly events happen and how quickly subsystems respond is likely to lead to results that are of maximum value. If applied consistently such an approach serves, in theory, to contribute to an ability to accurately predict remaining useful life, which at present is a significant challenge for the offshore wind industry.

Counting data and tracking changes over time in their own event categorisation bins will enable the analyst to align their work as closely as possible to solving the fundamental unknows or questions that are driving the analysis. It may also be that as assets age and organisations change the requirements and motivations will change over time, where again a self-defined or configurable bin definition may offer significant value.

Consider the merits of 1Hz data

In writing this O&M Case Study, we were able to refine our analysis and develop a robust conclusion partly because we were able to access data sampled every second for the turbine in question. With the costs of data storage continually reducing, and technology easing the handling and manipulation of large datasets improving rapidly, the case for investment in storing more data at higher sampling frequencies is compelling. The ability to work effectively with the fundamentally larger datasets created when using higher sampling rates will be considerably enhanced by having access to data science skills and experience. It is commonly the case that engineering teams have adapted and taken a DIY approach to conducting individual and discrete pockets of analysis using tools such as spreadsheets. While useful, these are likely not the most efficient way of extracting learning and understanding from large datasets.

Results

Relevance of initial results

In our experience, it is very easy to rapidly reach conclusions from alarm or event analysis that appear to be a cause for concern. Some challenges contributing to this might be a lack of baseline or knowledge about whether the number (or frequency) of certain events (or shutdowns) is normal, or whether it should be a cause for concern – for example, are events being "over-recorded" (where benign situations are being recorded but need not be a cause for concern)? If this is not clear to the analyst, then incorrect conclusions can be reached. In conducting analysis of alarm and event data it will be useful to identify reference sources to inform what can be considered normal behaviour of the asset. It may be the case that an increased number of events is not normal and should drive reactive maintenance, or it may be that slight variances in the number or frequency of alarms should trigger action.

A deeper understanding of the asset is a benefit

In making attempts to analyse the performance and behaviour of an asset, other – perhaps less tangible –benefits can often result beyond answering the questions driving the analysis. At its most basic level, spreading experience and understanding of how assets operate and generate data is likely to deliver benefits to an organisation. For example, if further questions are posed in future, it is likely that subsequent analyses can be generated more quickly. It is also possible that increasing the understanding can contribute to stronger connections between operational and technical support staff that might in turn lead to the more timely identification of anomalies in performance.

Show two different results from one dataset

As an illustration of some of the concepts described in this paper, Figure 1 shows two different representations of the same turbine over the same period of time. It shows the situation that is common in wind farm analysis, where a raw data format (namely turbine alarms) can either be analysed directly or via some intermediary data monitoring and analysis tool.

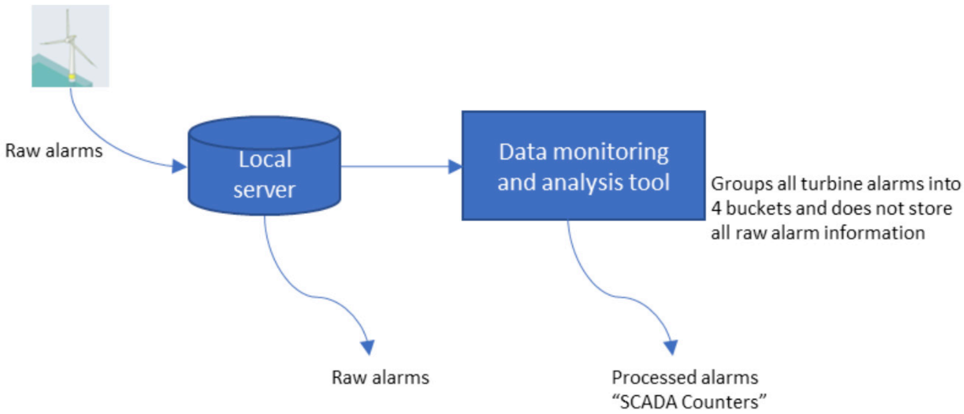

Figure 1: Turbine alarms being stored locally and being processed by a data monitoring tool.

In seeking to analyse the stop behaviour of the turbine in question, the analysis problem statement was whether a concerning amount of stop events were occurring. Using the same underlying dataset it is possible to present strikingly different results, in this case depending upon the classifications, groupings and bins that are used.

Turbine-Specific Investigation

Pre-defined binning

Alarms are counted and are grouped in a predefined binning, defined by the turbine manufacturer. There is little documentation as to what makes up these binnings, what alarms feed into them, and why they are categorised as so. When looking at these predefined groupings it may seem alarming that there is a large count of High Energy Stops, which could lead to incorrect interpretations as to what is happening at this wind turbine and therefore poorly-informed decisions. As seen from Table 1 and Figure 2, the High Energy Stops appear to make up over 50% of the total stops – an alarmingly high number.

Stop Category	Count
High Energy Stop	3192
High Medium	188
Low Medium	2138
Low	129

Table 1: Alarm count of automatically-defined stop groupings.

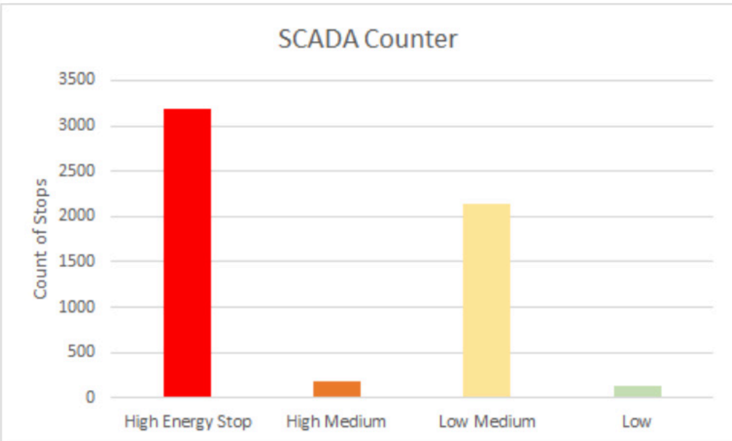

Figure 2: Alarm count of automatically-defined stop groupings.

Internally-defined binning

Alarms can be classified by the action the controller takes when it picks up that a certain alarm has been triggered. How the turbine reacts to different alarms changes how quickly the turbine stops generating. Some alarms may lead to a sudden stop in generation, where brakes are applied and aggressive stalling is undertaken, whereas some alarms may allow the turbine to slow down in a gentler manner. Knowing this information about the different alarms, site-specific binnings can be defined. As the user is aware of what makes up these binnings, more informed decisions can be made when it is felt there is a large amount of potentially damaging stops.

The same alarms were counted as before but binned by the newly-defined internal binning. As can be seen from Table 2 and Figure 3, what appeared to be high energy stops turned out to be more mild stops. If action was made on the previous information, it would have led to a poorly-informed decision. Now that more information is available, a well-informed decision can be made as to reducing the number of High Energy Fast stops.

Stop Type	High Energy	Low Energy	Total
Fast Stop	206	92	298
Slow Stop	2028	955	2983
No Stop	60	307	367
Relevant Event Missing	343	28	371

Table 2: Alarm count of internally-defined stop groupings.

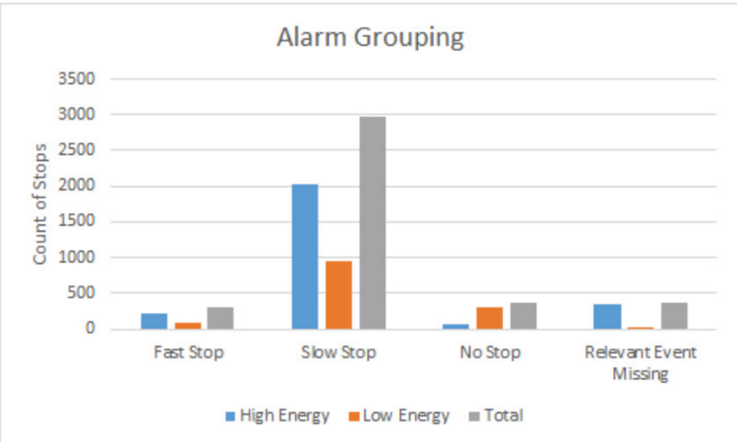

Figure 3: Alarm count of internally-defined stop groupings.

Lessons Learned and Best Practice

This case study has discussed how detailed and structured analysis of available wind turbine data sources can provide insight to improve O&M strategies. Through completion of the alarm logs analysis described in this case study, the following key lessons learned and best practice approaches have been realised or highlighted.

Scope definition

The first step in addressing any analysis is to start with a clear question – what are you trying to achieve? Defining a clear scope is integral to delivering conclusions and outcomes in an efficient and structured way. Poor scope definition will result in increased analysis time and less trusted and less insightful results.

Understand your data

It is possible that an analyst will have access to different and conflicting information from a variety of sources. Therefore, it is important to take a common sense approach to knowing which data can be trusted. Classification of data streams must be backed up with accurate documentation to remove uncertainty. If documentation is not available then we must use analysis or site visits to build confidence. Effective assessment of data quality is clearly integral to producing accurate, high-quality analytics.

Do it right, first time

A thorough and structured approach to analysing data sources is key to reaching the most accurate conclusions. Understanding all available data and interpreting the potential insights provided by this data is key before an analysis can be undertaken. If this approach is not followed, there is a risk of reaching an incorrect conclusion through delivering an incomplete analysis. It is also important to ensure that the analysis is fully complete and reviewed to ensure premature or inaccurate results are not distributed.

Relevant tools

While having access to the greatest-functionality tools is beneficial to providing quality outputs, valuable analyses can be undertaken by using widely-available software tools. Developing knowledge of packages such as Python and SQL may allow quicker analyses but the functionality of Microsoft Excel is such that it can still be used to provide value: albeit that we would recommend it as part of a diverse toolkit.

Presenting results

The method of sharing and presenting results is critical to the successful delivery of the key messages and outcomes. It is important to know your audience and to tailor the format and granularity of outputs to suit various levels of understanding and interest. Visual and summarised outputs are useful tools in presenting granular, specialist information to a diverse audience.

Wider value and developing connections

While addressing distinct use cases for analysing data provides direct value and outcomes, the experience and knowledge gained from these activities gives a more rounded understanding of data sources available and potential insights to be gained from these sources. It is recommended that analytics functions are closely linked, and encouraged to be in regular communication with, operational staff rather than periodic or "one-way" reporting.

Appendices

Author Profiles

Alistair Lee is a Graduate Engineer with the Catapult's Data and Digital team. He joined the Catapult in 2018, after graduating from the University of Strathclyde with a first in Naval Architecture and Ocean Engineering. Working primarily on analytics of offshore wind turbines, Alistair organises the offshore wind benchmarking service SPARTA, produces analytics reports, and provides engineering insight into analytics.

Disclaimer

While the information contained in this report has been prepared and collated in good faith, ORE Catapult makes no representation or warranty (express or implied) as to the accuracy or completeness of the information contained herein nor shall be liable for any loss or damage resultant from reliance on same.

Offshore Renewable Energy Catapult

Inovo

121 George Street
Glasgow
G1 1RD, UK
T: +44 (0)333 004 1400

National Renewable Energy Centre

Albert Street, Blyth
Northumberland
NE24 1LZ, UK
T: +44 (0)1670 359 555

Fife Renewables Innovation Centre

Ajax Way
Leven
KY8 3RS
T: +44 (0)1670 357 649

O&M Centre of Excellence

Ergo Centre
Bridgehead Business Park
Meadow Road, Hessle
HU13 0GD