



Total Control

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

A Hierarchical Wind Power Plant Supervisory Controller

D4.2

Delivery date: 04.11.2020

Lead beneficiary: SINTEF Energy Research

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
Valentin Chabaud	SINTEF Energy Research	valentin.chabaud@sintef.no
Paula B. Garcia Rosa	SINTEF Energy Research	paula.garcia-rosa@sintef.no
Konstanze Kölle	SINTEF Energy Research	konstanze.kolle@sintef.no
Ander Madariaga	ORE Catapult	ander.madariaga@ore.catapult.org.uk
Karl O. Merz	SINTEF Energy Research	karl.merz@sintef.no
John Nwobu	ORE Catapult	john.nwobu@ore.catapult.org.uk

Acknowledgements/Contributions:

Document information

Version	Date	Description	Reviewed by	Approved by
1.0	04.11.2020	Prepared by Karl O. Merz	All the authors	Approved by Gunner C. Larsen

Contents

1	Introduction	6
2	Control objectives, requirements, and design philosophy	7
2.1	Interface between the plant controller and other subsystems	7
2.2	Active load control and its limitations	9
2.3	Desired response characteristics	10
2.4	Design philosophy	10
3	Control architecture	11
3.1	Duelling controllers	12
3.2	Scheduling as a function of damage rate	15
4	Estimation of wind, loading, and fatigue	15
4.1	Design philosophy	15
4.2	Interface variables	16
4.3	Architecture	17
4.3.1	Turbulent wind	19
4.3.2	Rotor aerodynamics	20
4.3.3	Ocean wave forces	21
4.3.4	Available power	21
4.4	Observer gain matrix	21
4.5	Estimating the weather and turbulence	25
4.5.1	Clusters of turbines	26
4.5.2	Estimation of the cluster wind speed and direction	26
4.5.3	Time series of turbulence	29
4.6	Turbulence and response spectra	29
4.6.1	Measured turbulence spectrum	29
4.6.2	Theoretical turbulence spectrum and parameters	29
4.6.3	Rotationally-sampled turbulence spectra	30
4.6.4	Scaling of a reference spectrum	31
4.6.5	Ocean wave force spectra	31
4.6.6	Transfer functions and response spectra	31
4.6.7	Response spectra in fixed and rotating coordinates	32
4.7	Component wear	36
4.7.1	Material stress	37
4.7.2	Material fatigue	38
4.7.3	Actuator wear	49
4.8	Tuning and performance	49
4.8.1	Setup for testing the observer	49
4.8.2	Full and reduced-order observer models	50
4.8.3	Tuning and dynamic response	51
5	Examples of controller tuning and performance	51
5.1	Tuning, frequency response, and step response	56
5.2	Clustering, weather, and turbulence cascade	59
5.3	Spectra for estimating component wear rate	61
5.4	Controlling the TotalControl Reference Wind Power Plant	61

6	Power systems design and analysis with PSCAD	63
6.1	Electrical System	64
6.1.1	Collection Grid	64
6.1.2	Export system	66
6.2	Wind turbine model	66
6.2.1	Switch Gear	66
6.2.2	Electrical power train	67
6.3	Analysis of WPP model	71
7	Conclusions and Outlook	74
7.1	Outlook	74
7.2	Integrated model for case studies	74
7.3	Case studies on grid transients and faults	76
7.4	Case studies on WPP control and grid interactions	77

ATTACHMENTS




Figure 1: Architecture of the wind plant control law. The starred block provides anti-windup features.

Executive Summary

Operating a wind power plant (WPP) in the future will aim to maximise the revenue rather than maximising the power production. The reduction of structural loading and actuator wear and the provision of ancillary services will become more important in this context. There is a need for plant control strategies that trade these control objectives and respect the industrial practice, where the different levels in the hierarchy may be provided by different equipment vendors. A setup for a baseline plant controller is developed that uses only typically available data from the turbines' SCADA system. The plant controller is coupled to an observer that provides state estimates at the turbines.

Figure 1 shows the control architecture. This is based on a proportional-integral (PI) controller that provides power command tracking, and a proportional controller with low-pass filter that provides thrust compensation. The two objectives may conflict, but the structure of the controller ensures that power command tracking will prevail and dominate the response; the plant is therefore sure to meet grid code requirements for active power control. At the same time, the fact that the plant consists of many wind turbines allows the thrust compensation to be effective. The resulting performance is summarized in Fig. 2, where it is seen that the standard deviation in low-frequency thrust fluctuations can be significantly reduced for over half the turbines in the plant: specifically, those experiencing the highest damage rates, who need it the most. This comes at the cost of increasing the thrust fluctuations on a few turbines, but it is precisely these turbines, experiencing a low damage rate, that can best cope with it.




Figure 2: Trends in the mean and standard deviation of thrust and power fluctuations, as a function of damage rate. Each point represents one turbine in the TotalControl Reference Wind Power Plant. Nominal results for the case with a constant, uniform power command sent to all the turbines are shown in the background as gray points.

1 Introduction

Operating a wind power plant (WPP) in the future will aim to maximise the revenue rather than maximising the power production. The reduction of structural loading and actuator wear and the provision of ancillary services will become more important in this context. There is a need for plant control strategies that trade these control objectives while respecting the existing industrial practice, where the different levels in the hierarchy may be provided by different equipment vendors.

This report establishes a baseline, hierarchical wind power plant controller for the TotalControl project: “baseline”, because the controller relies only on standard sensors associated with existing SCADA (Supervisory Control And Data Acquisition) systems; and “hierarchical”, because the plant controller is subservient to the individual turbines’ controllers. The turbine-level controllers have complete command over the actuators (generator power, blade pitch, nacelle yaw), and provide the supervisory functions and limits needed for safe operation of the turbines. The interface between the plant and turbine controllers is limited to an active power command¹ sent from the plant to each turbine, and sensor measurements returned from the turbines to the plant.

The design of a control system can be split into, first, the specification of control objectives and requirements, and afterwards, the synthesis of a control architecture that fulfils the objectives. We have followed this approach, with Section 2 describing the specifications and Section 3 the synthesis. In short, the objective of the plant controller is to track an operator power command while providing coordinated active load control, reducing the severity of loading on the turbines in a “smart” way. It was found that a simple control law will suffice; but some rather sophisticated state estimation is needed in order to estimate the severity of loading on each turbine, based on limited sensor measurements. Section 4 describes the development of a state observer providing the required load metrics.

The performance of the plant controller is demonstrated in Section 5, for operation of the TotalControl Reference Wind Power Plant under curtailment – that is, the operator has provided a power command that lies under the maximum possible production, given the present wind conditions. The analysis is based on the electromechanical model of TotalControl D1-5 (STAS, Merz *et al.* 2019), to which the plant controller has been linked. It is found that there is a meaningful synergy in coordinating the control of the turbines in a large wind power plant: loads can be reduced on those turbines that need it the most, without excessively burdening the other turbines, and while still providing

¹We shall not consider reactive power (voltage) control, nor coordinated nacelle yaw control, in the baseline controller.

accurate plant-wide power tracking.

The hierarchical wind power plant controller will be used in TotalControl D4-6, on electrical-mechanical-control interactions. Towards this end, an electrical system model has been constructed in PSCAD, which interfaces with the plant controller. The PSCAD model can be used to study transient and nonlinear phenomena in the electric grid, over and above what can be done with the present version of STAS. Section 7 describes the integrated modelling of the wind plant and electric grid, with an outlook on some of the case studies that will be done for D4-6. Our approach is significantly novel, both in terms of the plant controller, and the way in which we are linking high-resolution models of the wind turbines and electric grid. We expect that the outcome will give insight into the interactions between the electrical and mechanical components in renewables-dominated grids, as well as the influence of control on grid stability; and lead to improved standards for the modelling and analysis of wind plants.

2 Control objectives, requirements, and design philosophy

The primary objective of the wind plant controller is to deliver the demanded active power \hat{P}_{PCC} at the point of common coupling (PCC) with the electric grid. The actual power must accurately track the command on a timescale of seconds, in order that the wind plant can provide primary frequency support. The secondary objective of the wind plant controller – to be considered only after the primary objective is satisfied – is to reduce the levels of loading and fatigue in the turbines.

The design of the baseline plant controller is subject to the following requirements:

1. There is a standard interface, described in Section 2.1, between the plant controller and its surroundings. Only standard SCADA sensor measurements are available: no direct measurements of strains or internal loads, no lidars measuring atmospheric flow, no ocean wave radar.
2. The plant controller does not include an embedded model of atmospheric flow or turbine wakes. Such models would increase the complexity of the controller beyond what is reasonable to expect for a “baseline” design.
3. The control hierarchy must be obeyed. The turbine controller has authority over the turbine’s actuators, and the turbine controller is a “black box” from the perspective of the plant controller: the plant controller sees only the interface variables, not the internal architecture and states of the turbine controller. Because of this, it might be more appropriate to think of the *command*, sent from the plant controller to a turbine, as a *request*: it is not necessarily delivered, and the plant controller must be able to adapt.
4. The plant controller must be of a type that the small-signal stability properties can be computed explicitly. The implication is that a state-space model of the plant controller can be linearized about a given steady-state operating point. This rules out certain features like hysteretic frequency-exclusion bands. The motivation for this restriction is to be able to study the dynamics and stability of large electric grids, to which the wind plant is connected, using eigenmodes and frequency-domain analysis.

2.1 Interface between the plant controller and other subsystems

A sketch of the closed-loop system, encompassing the wind plant and its surroundings, is shown in Fig. 3. What is important here is the definition of the interface variables passed between the blocks; Table I. In the view presented in Fig. 3, which is one of a number of possible ways to categorize the system,

- the wind plant consists of the wind turbines, local vortex wakes with induced velocity, collection grid, and electrical transmission system to the point of common coupling with the main electric grid;




Figure 3: A high-level block diagram of a wind power plant and its surroundings, showing the interface variables.

- the atmosphere consists of the surrounding flow that mixes with the turbine wakes upon breakup of their vortical structure; and,
- the electric grid here refers to the regional grid to which the wind plant is connected.

Then, the wind plant passes to the atmosphere the induced velocity \mathbf{V}_i caused by the vortical wake of each turbine, and receives in return the effective wind velocity \mathbf{V} approaching each wind turbine.² The wind plant delivers a three-phase electrical current waveform \mathbf{i}_{PCC}^θ (in the d - q frame labelled θ) to the PCC bus, while the grid provides the bus voltage \mathbf{v}_{PCC}^θ .

Consistent with the objectives of the present wind plant controller, it is assumed that only standard sensor measurements are available. From the electric grid, the plant controller receives the measured current $\bar{\mathbf{i}}_{PCC}^\theta$ and voltage $\bar{\mathbf{v}}_{PCC}^\theta$. Although wind field measurements $\bar{\mathbf{V}}$ are indicated in Fig. 3, these would require non-standard sensors and are not presently implemented. Rather, an estimate of the local wind speed at each turbine is available as part of the state observer of Section 4. The plant controller receives a large set of data from the wind plant, namely the sensors associated with the SCADA system mounted on each wind turbine. These provide measurements of the current and voltage waveforms, $\bar{\mathbf{i}}_s^\theta$ and $\bar{\mathbf{v}}_s^\theta$, at the transformer high-voltage terminals; the rotor speed $\bar{\Omega}$; the collective blade pitch $\bar{\beta}$; the nacelle yaw angle $\bar{\chi}$; the absolute nacelle velocity $\bar{\mathbf{v}}_n^g$, measured by an inertial measurement unit; the anemometer wind speed \bar{V}_a ; and the anemometer wind angle $\bar{\theta}_a$.

The wind plant controller provides active and reactive power (\hat{P} and \hat{Q}) and yaw angle ($\hat{\chi}$) set-point commands to each wind turbine. The wind turbine controller implements the power commands, if possible, through coordinated control of its generator (typically via a power converter) and blade pitch. This delegation of responsibility respects the hierarchy of control and safety functions in the plant. In addition to power and yaw commands, a wild-card variable labelled λ is shown in Fig. 3; the thought is that this may contain flags to set the turbines' operating mode, if the turbine controller allows for this.

²This view is most natural if the induced velocity is computed using a blade element momentum type method. If CFD is used to resolve the wake, local induction, and aerodynamic forces, then the interface would be the motions of the blades. There also exist lower-resolution CFD approaches landing somewhere in between these two bounding cases: the interface could also be drawn at the local aerodynamic forces, for instance.

Inputs	
\hat{P}_{PCC}	Active power demand at PCC
\hat{Q}_{PCC}	Reactive power demand at PCC
E	A metric representing the present electricity price
$\bar{i}_{PCC}^{-\theta}$	Measured current at PCC
$\bar{v}_{PCC}^{-\theta}$	Measured voltage at PCC
$\bar{i}_s^{-\theta}$	Measured d - q current at the turbine transformer terminals
$\bar{v}_s^{-\theta}$	Measured d - q voltage at the turbine transformer terminals
$\bar{\Omega}$	Measured rotor speed
$\bar{\beta}$	Measured collective blade pitch
$\bar{\chi}$	Measured yaw angle
$\bar{\mathbf{v}}_n^g$	Measured nacelle velocity in the global frame
\bar{V}	Anemometer wind speed
$\bar{\theta}_a$	Anemometer wind direction
Outputs	
\hat{P}	Active power command
\hat{Q}	Reactive power command
$\hat{\chi}$	Nacelle yaw command
λ	Flags setting the operating mode

Table I: General interface variables available to the control law. The second block of inputs, as well as the outputs, are repeated for each wind turbine in the plant.

2.2 Active load control and its limitations

There are different types of active load control, and the plant controller cannot provide all of them. It is useful to classify the types of active load control along the lines of Merz and Pedersen (2018, pp 148-149):

The wind turbine responds to the applied environmental loads. We can distinguish between three types of responses. The first is the response to the applied loads that would be encountered during steady-state operation in a uniform, constant wind, and ocean currents. These loads include the thrust and torque associated with extracting energy from the wind, and the unavoidable gravity cycles in rotating components. The second category consists of the response to fluctuations and nonuniformities in the environment – turbulence, wind shear, and ocean waves – at frequencies which are not near resonance. The dynamic amplification is low, and the response is concentrated at frequencies associated with peaks in the applied loads. The third type of response is resonance. The dynamic amplification is high at particular resonant frequencies. A small amount of energy in the environmental loads at these frequencies may trigger a large response...

We define the control objectives of *load reduction* as a strategy which reduces the steady-state loads on the turbine; *load rejection* as counteracting fluctuations in the applied loads, usually (though not necessarily) at frequencies away from resonance; and *active damping* as counteracting the response at poorly-damped resonance frequencies.

Through intelligently tailored power commands, a plant controller can contribute to load reduction, and the rejection of low-frequency aerodynamic loads. Consider a spectrum of internal bending moments in the tower, like that shown in Fig. 4. This shows two instances of load rejection, (a) rejecting turbulent wind and (b) ocean waves; and one instance of active damping, at the first tower resonant frequency. The plant controller could conceivably contribute to (a), but not (b) or (c).




Figure 4: Active load control, illustrated on a spectrum of tower bending moments. The dashed lines indicate the original spectrum, with no active load control. (a): rejection of low-frequency aerodynamic loads. (b): rejection of wave loads. (c): active damping of resonance.




Figure 5: Sketches of the desired response characteristics of the commanded power \hat{P} under different combinations of power-tracking $\lambda\epsilon_P$ and thrust-tracking $P(\hat{F}_T)$ errors.

2.3 Desired response characteristics

Figure 5 shows the desired trends in the response under different combinations of power-tracking and load-tracking errors; here it is assumed that the load is the rotor thrust F_T . In essence, the nominal power level is set by the thrust-tracking function. Then, a correction with integral action is added based on the turbine’s share of the plant-wide power error, $\lambda\epsilon_P$. Note in particular how power-tracking takes precedence; that is, under a sustained error in the target thrust, the thrust-tracking controller must not overwhelm the offset dictated by the power-tracking controller.

2.4 Design philosophy

Before proceeding, it might be worth adding a few words about the philosophy used in designing the control architecture. In addition to the objectives, requirements, and limitations mentioned above, we wanted the control law to be simple. This makes it understandable and accessible: useful as a baseline for comparison against more advanced algorithms. Being simple makes the control law useful for other types of studies, such as the grid integration of wind power plants, or wind turbine loads

when operating in coordination with other turbines: studies, in other words, which require a good wind plant controller, but whose emphasis is not the controller itself.

Drawing inspiration from Madjidian (2016), we came upon the idea of parallel control loops: one loop provides the “ideal” active load control for each turbine, and the other provides a correction to the power set-point, ensuring that the overall plant power equals that commanded by the operator, to the extent that wind conditions allow. Note the implication: the problem is not how to dispatch power set-points to the turbines in the plant; rather, the problem becomes how to dispatch power *corrections* to the turbines in the plant. There is little correlation in turbulent wind speed fluctuations from one turbine to the next, and when averaged over a large wind plant, the corrections will tend to be small. This is nice, because we can then consider the active load control function independently of the plant power command tracking function.

We chose to pursue a model-based controller, where a model of the system is used to provide meaningful quantities of interest, on which a set of straightforward control laws can be based. Thus, although our controller is model-based, we did not employ “optimal” control synthesis, in either its linear LQG variant, or a more difficult nonlinear Lyapunov/Pontryagin approach. We expect, based on the results of Madjidian, that the potential for improvement is small; although, an optimal control approach would be useful for simultaneously implementing different kinds of active load control, rather than pursuing only single goals like “steady thrust” or “steady torque”.

Adaptive control was considered, but we decided not to go this route. In the context of wind plant control, the performance of model-free adaptive controllers suffers due to the continually fluctuating wind conditions; extremum-seeking controllers will have a hard time finding a steady “optimum” to which to converge. Parameter adaptation of a fixed-architecture model-based controller could be useful, but we leave this to a later iteration of the control design. Algorithms that continually adapt the control architecture fall into the category of machine learning, which we find too advanced – and, in practice, uncertifiable – to be considered as a baseline.³

While the control law (Section 3) is simple, our state observer (Section 4) ended up being complex, due to an effort to obtain realistic estimates of fatigue from limited sensor measurements. If it is not desired to implement such a complicated observer, it is possible to retain the spirit of the control algorithm while using a simpler metric for the severity of loading, like the standard deviation of nacelle displacements.

3 Control architecture

With reference to Fig. 3, let us zoom in for a closer look at the plant control block. This is shown in Fig. 6. The plant controller consists of a state observer – an embedded model of the system – and a control law.

The role of the state observer is to take the SCADA data feed from each wind turbine and estimate quantities used by the control law. Although the state observer estimates the state variables \mathbf{x}^* , these are not used directly in the control law. Following the selected design philosophy, the observer provides physically meaningful outputs. These are classified into two categories: standard outputs \mathbf{z}^* , which can be expressed in the form $g(\mathbf{x}^*)$ and linearized as $\mathbf{C} \Delta \mathbf{x}^*$; and fatigue outputs $\mathbf{\Pi}^*$, which are derived from spectral analysis and fatigue cycle counts. A more detailed description of the state observer and its outputs is given in Section 4. What matters here is that \mathbf{z}^* includes important loads like the rotor thrust F_T^* and shaft torque T^* , the wind speed V^* , and available power P_a^* ; while $\mathbf{\Pi}^*$ includes the wear (degradation, fatigue) rate $\mathbf{\dot{D}}^*$ for key components like the blades, driveshaft, foundation, and actuators.

Figure 7 shows the proposed control law. This consists of two main pathways: one tracking the operator power command, and the other providing active load control. Both pathways are informed by

³We are optimistic about the future of data-driven algorithms, and automated synthesis of nonlinear, robust, multi-objective controllers.




Figure 6: An expansion of the plant control block, showing how the interface variables feed into the state observer and control law.

the component wear rate $\dot{\mathbf{D}}^*$, and the available power P_a^* based on the current wind conditions. The control law is repeated for each turbine.⁴ This implies that the integration of the error in plant power ε_{PCC} , which ensures tracking, is done separately for each turbine. The reason for doing things this way is so that gain-scheduling via λ (associated with damage rate) can be done *inside* the integrator, as is proper; and so that saturation can be customized to fit the operating conditions of each turbine.

Note that both the thrust-tracking and power-tracking controls act through the same power command. In typical PI control designs, the effectiveness of each control loop is ensured by separating them in frequency; for example, an active tower damping controller may act in parallel with the rotor speed controller, since filters are applied to isolate the action of the former to the vicinity of the tower resonant frequency, and remove this frequency from the latter. This frequency separation is not possible here: the specifications call for a controller that can adjust the power so as to track two targets – rotor thrust and plant power – at once, in the same frequency band. For an individual turbine, this would be impossible. But the plant power is the sum of N turbines; hence the hope is that some of the turbines are in a position to compensate for imperfect tracking by other turbines.

3.1 Duelling controllers

A viable architecture is provided by giving the pathways different response characteristics: the power-tracking branch is a PI controller, with integral action, while the thrust-compensation branch is a proportional gain acting against a low-pass filter.⁵ They can then duel, but after some time the power-tracking function will inevitably win. The thought behind the “duelling-controllers” architecture can be stated like this: When the rotor thrust and plant power targets agree, then they reinforce each other, and the turbine will move towards the common set-point quickly. When the rotor thrust and plant power targets disagree, then they partially cancel each other, and the turbine will move towards

⁴The plant control law is implemented on central control hardware, not at each turbine; but on the central hardware there is a block (Fig. 7) associated with each turbine.

⁵...that is, an *additional* low-pass filter, since all the inputs are already low-pass and notch filtered. A proportional-gain/low-pass filter arrangement is common in power systems, this being the effective response of a governor with speed-droop. (Kundur 1994, p. 589)




Figure 7: Architecture of the wind plant control law. The starred block provides anti-windup features; it is shown in Fig. 8.




Figure 8: The block providing anti-windup features.




Figure 9: Output of the anti-windup block for a case in which the valid range of u is between 0 and 1. The smoothed transition between $\eta = 0$ and $\eta = 1$ is convenient for numerical solution of the nonlinear state equations.

the power set-point slowly at first, but then more quickly as the thrust error reaches its equilibrium. Therefore, turbines for whom a given power correction is favorable will end up taking more, and those for whom it is unfavorable less.

When a turbine reaches the maximum power allowed by the local wind conditions, it may never be able to reach the commanded power set-point. The anti-windup strategy is represented by the starred block in Fig. 7, expanded in Fig. 8. This represents a “soft” anti-windup algorithm, with a smoothed transition between zero and full saturation, with an example shown in Fig. 9. This smoothing is perhaps not necessary for the operation of the controller, but it is implemented for consistency with the rest of the STAS program, where smoothed transitions are needed for numerical stability during Newton-Raphson solutions and gradient-based optimization. Soft anti-windup also has the benefit of allowing the integrator to track a bit beyond the available power, whose estimate may not always be precise.

The turbine’s nominal share of the plant power command \hat{P}_{PCC} – what \hat{P} would be in the absence of thrust and power-tracking corrections $\delta\hat{P}_T$ and $\delta\hat{P}_P$ – is provided as a feed-through pathway, with no control dynamics except those provided by the filtering of the input signals. Power-tracking and thrust compensation are provided as corrections to this nominal output.

The target thrust \hat{F}_T is set according to the wind turbine’s typical $C_T(\text{TSR}, \beta)$ table, where $\text{TSR} = R\Omega/V_\infty$ is the tip-speed ratio. However, there is a major catch, because it is not desired that the target thrust tracks the actual blade pitch angle, nor local turbulent fluctuations in the wind speed; it is, after all, the effect of these that we are trying to reject! For the wind speed, we have the cluster estimate μ_V of Section (4.5.2). The blade pitch is trickier. Starting with the nominal power,

$\lambda_1 \hat{P}_{PCC}$, it is possible to convert this to a power coefficient C_P , and then solve the $C_P(\lambda, \beta)$ table for a nominal value of the blade pitch, $\hat{\beta}$. It is this value, rather than the actual blade pitch, that is then used in the C_T table – that is, $C_T(R\bar{\Omega}/\mu_V, \hat{\beta})$ – to get the target thrust. It will be seen in Section 4 that the value of \dot{D}^* changes slowly, on a timescale of several minutes, and so the variability of λ_1 is not a problem for purposes of setting \hat{F}_T .

3.2 Scheduling as a function of damage rate

The nominal power in the feed-through pathway and the gains on the power-tracking controller, are scheduled as a function of the estimated damage rate \dot{D}^* . The scheduling is determined by two functions $\alpha_1(\dot{D}^*)$ and $\alpha_2(\dot{D}^*)$. In the case study of Section 5.4 a simple linear relationship is implemented, but the functions may instead be based on a more thorough consideration of whatever “damage rate” is intended to represent in practice. In any case, the α ’s are a turbine’s way of saying, “I would like to reduce my static load level and my share of power-command tracking.” This information is shared with the other turbines: everyone can see everyone else’s α ’s.⁶ The j^{th} turbine then gets its fair share of the burden by weighting according to

$$\lambda_{1,j} = \frac{\alpha_{1,j}}{\alpha_{1,j} + \sum_{k \neq j} \alpha_{1,k}}, \quad (1)$$

and similarly for $\lambda_{2,j}$. The sum in the denominator is taken over all N turbines in the plant, except for the j^{th} one, whose control block we are in.

4 Estimation of wind, loading, and fatigue

In order to effectively fulfil its control objectives, a wind plant controller needs to know something about the operating state of each wind turbine. An estimate of the operating state is obtained through sensor measurements. Combining these sensor measurements with a model of the system – an observer – provides additional knowledge and insight, which can be used to aid decision-making, finding the best control actions.

4.1 Design philosophy

The observer is a model of the system, and there are many possible ways to construct such a system model. One aspect is the fidelity of the model, how much detail is used in modelling the physical processes. A model embedded in a real-time controller must be at the low-fidelity end of the range, but the model used for control *design* (including synthesis) does not need to be.⁷ Thus the upper constraint on fidelity is mostly due to the analysis tools at our disposal, and if there is a need or desire for human-understandability in the design process.

Another key choice regards the synthesis of the control architecture, and accordingly the outputs needed from the observer. The control architecture may be determined automatically using state-feedback methods and the separation principle, in which case the designer’s role is to specify a suitable performance metric. Often the performance metric is chosen as a quadratic function of the state and control variables, and then the designer iteratively adjusts the weights until a desirable performance is obtained. The observer’s role is to output estimates of a set of state variables that can be assumed to represent the “complete” system dynamics, and, by the separation principle, the observer is designed and tuned independently of the control law and its performance metric. Alternatively the control

⁶Again, this is all done on the central control hardware, so it isn’t actually necessary to send the information out to the individual turbines, just their plant-control blocks.

⁷Zhou *et al.* (1996), for instance, recommend synthesizing a controller based on a comparatively high-fidelity model, and then reducing the controller in a subsequent step.

Table II: Definition of observer outputs. Above the line: \mathbf{z}^* . Below the line: $\mathbf{\Pi}^*$.

\mathbf{x}_t^*	Wind turbine state variables
V^*	Effective wind speed
θ^*	Effective wind direction
a^*	Effective axial induction factor
F_T^*	Rotor thrust
F_w^*	Ocean wave force
θ_w^*	Ocean wave direction
P_a^*	Available active power
Q_a^*	Available reactive power
S_V^*	Spectrum of effective wind speed
\mathbf{S}_w^*	Wave force cross-spectra
\mathbf{S}_σ^*	Cross-spectra of structural stresses
S_β^*	Spectrum of pitch actuator rate
\mathbf{D}^*	Component wear (degradation, fatigue) rate
\mathbf{D}_a^*	Estimated remaining life

architecture may be designed manually, and in this case the observer should output quantities that are most directly meaningful to the designer. Here the performance metrics – things to be maximized, or minimized, or driven to zero – become part of the observer, and the control law is built around them; the metrics are not just used in the synthesis step, they are retained in real-time operation. There is no requirement that the observer model represents the complete system; selected processes or subsystems can be modelled, and others omitted.

The most general definition of an observer would be one that models a full set of state variables, and also outputs performance-related variables. We pursue such an observer, to allow flexibility in the approach to control design.

4.2 Interface variables

The observed variables may be organized as follows:

$$\mathbf{z}^* = \left\{ \mathbf{x}_t^*, \begin{bmatrix} V^* \\ \theta^* \\ a^* \\ F_T^* \end{bmatrix}, \begin{bmatrix} F_w^* \\ \theta_w^* \end{bmatrix}, \begin{bmatrix} P_a^* \\ Q_a^* \end{bmatrix} \right\}, \quad \mathbf{\Pi}^* = \left\{ \begin{bmatrix} S_V^* \\ \mathbf{S}_w^* \\ \mathbf{S}_\sigma^* \\ S_\beta^* \end{bmatrix}, \begin{bmatrix} \mathbf{D}^* \\ \mathbf{D}_a^* \end{bmatrix} \right\}. \quad (2)$$

Table II defines the variables in (2). It is useful if we divide the variables into two types, denoted respectively \mathbf{z}^* and $\mathbf{\Pi}^*$. The \mathbf{z}^* variables listed in the upper part of Table II are estimates of things that are happening “now”, at the present time. By contrast, the spectral variables in $\mathbf{\Pi}^*$ derive from correlations over different lengths of time. Also, component wear is a fundamentally nonlinear function of material stress cycles – that is to say, it depends on a count and categorization of reversals in the time series of material stress. Though straightforward to compute in either the time or frequency domain, the $\mathbf{\Pi}^*$ variables require special consideration, especially as regards a state-space implementation.

The state observer was designed under the assumption that the system is linear, operating in the vicinity of a set of equilibrium points. Working with linear system models makes the dynamics human-understandable through eigenmodes; the principle of superposition greatly simplifies analysis; and frequency-domain transfer functions with associated load and response spectra enable rapid calculation of fatigue (component wear rates), as well as rough estimates of the extreme response. Synthesis of a final, nonlinear controller is done in a subsequent step by scheduling over the family of linear models. This may require special logic to handle strong nonlinearities, such as the transition between maximum-power-point tracking and power set-point control modes. The scope of this report is limited




Figure 10: The architecture of the state observer, showing the embedded wind turbine model, a module that distinguishes local turbulence from broader weather patterns, and an analytical spectral model that estimates the severity of loading and fatigue.

to the first phase, designing state observers for selected operating points, with the use of linear systems theory.

4.3 Architecture

The architecture of the state observer is shown in Fig. 10. The observer consists of an embedded model of the wind turbine, which, together with its gain matrix \mathbf{G} , are set up as a typical linear optimal filter (Sections 4.3 and 4.4). The outputs are then augmented by a set of calculations that culminate in an estimate of the wear rate of selected structural components (Sections 4.5 through 4.7).

The state observer is based on an *open-loop* model of the wind turbine, including the local wind and wave environments. Figure 11 shows the architecture of the turbine model. Inputs are the signals from the plant controller, and a collection of white noise processes exciting the wind (turbulence) and ocean waves. The turbulent wind and ocean waves are modelled as filtered white noise: this is an assumption, a rough approximation of the true behavior, for purposes of state observation. The wind turbine consists of the local rotor aerodynamics, structures, actuators, electrical components, and the turbine’s own controller. Together, the states of these subsystems are collected into \mathbf{x}_t . Some of the components are equipped with sensors, which will eventually be used to form \mathbf{y}^* ; and the quantities needed to form the outputs \mathbf{z}^* and $\boldsymbol{\Pi}^*$ are obtained from various parts of the model.

The equation of a (linear) state observer has the form

$$\frac{d\Delta\mathbf{x}^*}{dt} = \mathbf{A} \Delta\mathbf{x}^* + \mathbf{B} \Delta\mathbf{u} + \mathbf{G}(\Delta\bar{\mathbf{y}} - \Delta\mathbf{y}^*) \quad (3)$$

with

$$\Delta\mathbf{y}^* = \mathbf{C}_y \Delta\mathbf{x}^* \quad (4)$$

in continuous time. Here \mathbf{x}^* is a vector of states; \mathbf{u} is a vector of control inputs, fed back from the controller; $\bar{\mathbf{y}}$ is the set of sensor measurements that are passed to the observer; and \mathbf{y}^* is the




Figure 11: The architecture of the wind turbine model around which the state observer is built. The arrow behind the turbine subsystem blocks is to indicate that many internal variables, not resolved in the figure, are passed back and forth between these blocks.

prediction of the same sensor measurements, derived from the state vector. The Δ 's indicate that the equations have been linearized, such that the dynamic variables represent perturbations with respect to an operating point, usually an equilibrium point. The discrete-time equation is similar,

$$\Delta \mathbf{x}^{*,k+1} = \mathbf{A} \Delta \mathbf{x}^{*,k} + \mathbf{B} \Delta \mathbf{u}^{k+1} + \mathbf{G}(\Delta \bar{\mathbf{y}}^{k+1} - \Delta \mathbf{y}^{*,k+1}), \quad (5)$$

the only “catch” being that the estimated sensor measurements

$$\Delta \mathbf{y}^{*,k+1} := \mathbf{C}_y(\mathbf{A} \Delta \mathbf{x}^{*,k} + \mathbf{B} \Delta \mathbf{u}^{k+1}) \quad (6)$$

are computed using a *prediction* $\mathbf{A} \Delta \mathbf{x}^{*,k} + \mathbf{B} \Delta \mathbf{u}^{k+1}$ of the states at timestep $k + 1$. We shall formulate the state equations in continuous time, switching to discrete time when this allows us to avoid mathematical difficulties with white noise: dealing with continuous-time white noise is an unnecessary complication, since the equations will ultimately be solved numerically in discrete time. A second-order Runge-Kutta scheme is sufficient: if a continuous-time state equation is $d\mathbf{x}/dt = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}$, then in discrete time

$$\mathbf{x}^{k+1} \approx \left(\mathbf{I} + \mathbf{A} \Delta t + \frac{1}{2} \mathbf{A}^2 \Delta t^2 \right) \mathbf{x}^k + \left(\mathbf{B} \Delta t + \frac{1}{2} \mathbf{A} \mathbf{B} \Delta t^2 \right) \mathbf{u}^{k+1}. \quad (7)$$

Alternatively, to guarantee numerical stability with a comparatively large timestep, an implicit integration scheme may be used, such as the second-order trapezoidal method,

$$\mathbf{x}^{k+1} \approx \left(\mathbf{I} - \frac{1}{2} \mathbf{A} \Delta t \right)^{-1} \left(\mathbf{I} + \frac{1}{2} \mathbf{A} \Delta t \right) \mathbf{x}^k + \left(\mathbf{I} - \frac{1}{2} \mathbf{A} \Delta t \right)^{-1} \mathbf{B} \Delta t \mathbf{u}^{k+1}. \quad (8)$$

The outputs \mathbf{z}^* and $\mathbf{\Pi}^*$ are functions of the state variables \mathbf{x}^* ; in fact, many of the \mathbf{z}^* are state variables. Let us go through the outputs in turn and look at how each is estimated by the observer.

4.3.1 Turbulent wind

The observer’s internal representation of the wind starts with an effective wind speed V and wind direction θ . This can be thought of as the incoming wind, spatially averaged over the rotor swept area; it is the effective “upstream” wind that does not include the induction due to the vortex wake. The turbine’s standard sensors – which, it is assumed, do not include strain measurements in the blades – are not sufficient to resolve features in the turbulence with a length scale smaller than the rotor diameter. Rather, small-scale, high-frequency, and rotationally-sampled features in the turbulence are filled in by a stochastic model, calibrated to the observed large-scale properties.⁸ This is necessary in order to predict the severity of the dynamic loading and response in various components.

There are three phases to the calculation of the turbulent wind and turbine response. The first phase (Section 4.4) is the estimation of the effective wind speed and direction at each wind turbine, using the turbine’s sensors. The second phase (Section 4.5) is to separate the “weather” from the “turbulence” by establishing a consensus, among a cluster of adjacent turbines, of the large-scale flow properties. The third phase (Section 4.6) is to fit a spectral model to the estimates of effective wind speed, and use this to estimate the stochastic properties of the turbine response.

At present, suffice it to say that the internal representation of the effective wind speed and direction, for purposes of state estimation, is low-pass filtered white noise. There is also significant energy at the $3P$ blade-passing frequency; this can be represented as a band pass filter. The combined filters have the functional form

$$\frac{d}{dt} \begin{bmatrix} V \\ \theta_V \\ \Psi_V \\ V_{3P} \end{bmatrix} = \begin{bmatrix} -\alpha_V & 0 & 0 & 0 \\ 0 & -\alpha_\theta & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -\alpha_{3P}^2 & -2\zeta_{3P}\alpha_{3P} \end{bmatrix} \begin{bmatrix} V \\ \theta_V \\ \Psi_V \\ V_{3P} \end{bmatrix} + \begin{bmatrix} \alpha_V & 0 \\ 0 & \alpha_\theta \\ 0 & 0 \\ 2\zeta_{3P}\alpha_{3P} & 0 \end{bmatrix} \begin{bmatrix} n_V \\ n_\theta \end{bmatrix}. \quad (9)$$

⁸The anemometer measures small-scale and high-frequency turbulence at a single point near the nacelle, but this cannot be used to resolve the turbulent flow over the rest of the rotor plane.

Here α is the corner frequency in rad/s, ζ is a damping ratio, and n represents either white noise or an input time series. The parameters are calibrated based on the measured wind speed and direction; refer to Section 4.5 for the details.

4.3.2 Rotor aerodynamics

Rotor aerodynamic outputs include the rotor-average induction factor and the rotor thrust. The induction factor a is related to the strength of the vortex wake, and hence the severity of downstream wake deficits and turbulence. The rotor thrust is a useful metric of the aerodynamic loading; for instance, a load-reducing control law can track a constant rotor thrust and thereby eliminate low-frequency aerodynamic load fluctuations.

The rotor-average induction factor is defined as the ratio of the induced velocity V_i and effective wind speed V ,

$$a := -\frac{V_i}{V}. \quad (10)$$

The negative sign makes a a positive quantity when V_i opposes V , as is always the case for a wind turbine. The induced velocity is one of the aerodynamic states associated with each blade element; that is to say, it is part of the state vector \mathbf{x}_t and is therefore readily available. The elements' induced velocities are averaged over the rotor swept area to obtain a scalar V_i . Linearizing (10) gives

$$\Delta a := -\frac{1}{V_0} \Delta V_i^* + \frac{V_{i0}}{V_0^2} \Delta V^*, \quad (11)$$

where V^* is the estimated wind speed.

The rotor thrust is the summation of the axial component of aerodynamic force along the blades. The element forces are output from the STAS aeroelastic module. However, these are expressed in multi-blade coordinates and with respect to the blade body (pitch) coordinate system. They need to be transformed into the rotor coordinate system before extracting the axial component and summing. The full transformation takes the form

$$\mathbf{F}^{r,\psi} = \mathbf{T}_B^\psi \mathbf{T}_r^r \mathbf{T}_y^y \mathbf{T}_{y0}^{y0} \mathbf{T}_g^{g0} \mathbf{T}_{p0}^{p0} \mathbf{T}_p^p \mathbf{F}^{p,\psi}. \quad (12)$$

Linearized, this is

$$\Delta \mathbf{F}^{r,\psi} = \mathbf{T}_B^\psi \mathbf{T}_p^p \Big|_0 \mathbf{T}_\psi^B \Delta \mathbf{F}^{p,\psi} + \mathbf{T}_B^\psi \left(\mathbf{T}_y^y \frac{\partial \mathbf{T}_{y0}^{y0}}{\partial \mathbf{q}^B} \Big|_0 \mathbf{T}_p^p \Big|_0 + \mathbf{T}_{p0}^{p0} \Big|_0 \frac{\partial \mathbf{T}_p^p}{\partial \mathbf{q}^B} \Big|_0 \right) \mathbf{T}_\psi^B \mathbf{F}_0^{p,\psi} \mathbf{T}_\psi^B \Delta \mathbf{q}^\psi. \quad (13)$$

The equations describing the transforms are documented in Merz (2018). The z component of \mathbf{F}^r is the thrust-direction force at the given blade element, and the sum over all blade elements gives the total rotor thrust.

The available power P_a is defined based on the observed wind speed V^* . Let the maximum power coefficient of the wind turbine be C_P^* . The available power is then

$$P_a = C_P^* \frac{1}{2} \bar{\rho} A (V^*)^3, \quad (14)$$

where $\bar{\rho}$ is the measured air density and A the rotor swept area. Linearized, this is

$$P_a = C_P^* \frac{3}{2} \bar{\rho} A (V_0^*)^2 \Delta V^*. \quad (15)$$

4.3.3 Ocean wave forces

Ocean waves are represented as an oscillating wave force on the support structure, and a direction from which the dominant waves approach. The force amplitude can be represented by a second-order band-pass filter⁹

$$\frac{d}{dt} \begin{bmatrix} \Psi_w \\ F_w \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -\omega_w^2 & -2\zeta_w\omega_w \end{bmatrix} \begin{bmatrix} \Psi_w \\ F_w \end{bmatrix} + \begin{bmatrix} 0 \\ 2\zeta_w\omega_w \end{bmatrix} n_w, \quad (16)$$

with $\zeta_w = 0.1$ and $\omega_w = 2\pi/T_p$. It is assumed that the wave direction θ_w is known from either forecasts or local measurements at the site.

4.3.4 Available power

Under a given steady wind condition a wind turbine has a maximum steady-state active and reactive power that can be provided. Active and reactive power capabilities are completely different, but not always independent. The active power capacity is determined by the wind, or in above-rated conditions by the power rating of the generator and other electrical components. Active power draws energy from the rotating driveshaft, and therefore impacts the mechanical components. Reactive power, by contrast, is a purely electrical quantity. In some wind turbine designs, the full reactive power capacity can be provided under any operating conditions. In other designs, there is a relationship between the active power and the reactive power capacity. At present the former type is assumed, such that the reactive power capacity is constant, provided as an input at the start of the analysis.

As for the active power capacity, this is computed from the measured rotor speed. The power-speed curve for maximum power tracking is known to the controller, and for a given rotor speed the maximum power can be read directly from this curve. When tracking a curtailed power command, the present turbine controller¹⁰ sets the target rotor speed based on the estimated wind speed, and so there is a direct relationship between the rotor speed and the available power – or at least, a best estimate of the available power.

4.4 Observer gain matrix

A Kalman filter (least-squares optimal filter) is used to estimate turbine states and the derived outputs of Section 3. Errors between predicted and measured sensor outputs are linked with estimates of the state variables through a gain matrix. This gain matrix is determined automatically as part of the Kalman filter. It can be said to be “optimal”, but this can be misleading, since it is usually necessary to tune a set of weights in the performance metric in order to obtain a good result. That said, it is quite convenient to work with a small set of tunable parameters, rather than manually specifying the structure and values of the gain matrix.

The architecture of the system is sketched in Fig. 12. This is based on a discrete-time representation of the dynamics. The system can then be “unrolled” in time, in essence converting the dynamic system into a static network along which the variables, and their probability distributions, propagate. The system model is

$$\mathbf{x}^k = \mathbf{A}\mathbf{x}^{k-1} + \mathbf{B}\mathbf{w}^k, \quad \bar{\mathbf{y}}^k = \mathbf{C}\mathbf{x}^k + \mathbf{n}^k \quad (17)$$

and the corresponding observer model is

$$\mathbf{x}^{*,k} = \mathbf{A}\mathbf{x}^{*,k-1} + \mathbf{G}(\bar{\mathbf{y}}^k - \mathbf{y}^{*,k}), \quad \mathbf{y}^{*,k} = \mathbf{C}\mathbf{A}\mathbf{x}^{*,k-1}. \quad (18)$$

Here \mathbf{w} and \mathbf{n} are disturbances, \mathbf{w} on the system and \mathbf{n} on the sensor measurements. They are intended to be white-noise processes, although we shall avoid mathematical treatment of continuous-time white noise. Other inputs \mathbf{u} , like control signals, which cannot be considered disturbances, are neglected.

⁹Fossen (1994), Smilden (2019)

¹⁰Merz *et al.* (2019)




Figure 12: The discrete-time Kalman filter “unrolled” in time, as a network.

To be concise we have dropped the Δ 's on the dynamic variables: it is implicit that we are wholly in the world of linear dynamic systems.

The crux of the problem is to find the gain matrix \mathbf{G} . This is developed in two phases, first as a formal optimization problem, which provides some guidance; then the approach is modified to be more practical.

In the formal version, we seek to minimize the mean-squared error between the actual and predicted state variables.¹¹ Defining the error $\boldsymbol{\xi} = \mathbf{x} - \mathbf{x}^*$,

$$\Pi = E[\boldsymbol{\xi}^T \boldsymbol{\xi}] = \text{Tr}\{E[\boldsymbol{\xi} \boldsymbol{\xi}^T]\} = \text{Tr}\{\boldsymbol{\Sigma}\}, \quad (19)$$

where $\boldsymbol{\Sigma}$ is the covariance. The dynamics of the error follow from (17) and (18),

$$\boldsymbol{\xi}^k = (\mathbf{I} - \mathbf{G}\mathbf{C})\mathbf{A}\boldsymbol{\xi}^{k-1} + (\mathbf{I} - \mathbf{G}\mathbf{C})\mathbf{B}\mathbf{w}^k - \mathbf{G}\mathbf{n}^k. \quad (20)$$

Write this as

$$\boldsymbol{\xi}^k = \boldsymbol{\Phi}\boldsymbol{\xi}^{k-1} + \boldsymbol{\Psi}\mathbf{w}^k - \mathbf{G}\mathbf{n}^k. \quad (21)$$

From the definition of covariance,

$$E[\boldsymbol{\xi}^k (\boldsymbol{\xi}^k)^T] = E[(\boldsymbol{\Phi}\boldsymbol{\xi}^{k-1} + \boldsymbol{\Psi}\mathbf{w}^k - \mathbf{G}\mathbf{n}^k)(\boldsymbol{\Phi}\boldsymbol{\xi}^{k-1} + \boldsymbol{\Psi}\mathbf{w}^k - \mathbf{G}\mathbf{n}^k)^T], \quad (22)$$

such that

$$\boldsymbol{\Sigma}^k = \boldsymbol{\Phi}\boldsymbol{\Sigma}^{k-1}\boldsymbol{\Phi}^T + \boldsymbol{\Psi}\mathbf{Q}\boldsymbol{\Psi}^T + \mathbf{G}\mathbf{R}\mathbf{G}^T, \quad (23)$$

with $\mathbf{Q} = E[\mathbf{w}\mathbf{w}^T]$ and $\mathbf{R} = E[\mathbf{n}\mathbf{n}^T]$ being the covariance of \mathbf{w} and \mathbf{n} , respectively. Since \mathbf{w} and \mathbf{n} are assumed to be white noise,¹² all the cross-covariance terms like $E[\boldsymbol{\xi}\mathbf{w}^T]$, $E[\mathbf{w}\mathbf{n}^T]$, and so on, are zero.

¹¹Note that for this measure to be meaningful it is critical to scale the variables to the same order of magnitude. If needed this can be accomplished with a simple transformation $\mathbf{z} = \mathbf{S}\mathbf{x}$, applied to the state equations.

¹²There is really no loss of generality with this assumption, since correlated signals can be generated through augmentation of the state space.

We can see in (20) and (23) the tension between the gain matrix \mathbf{G} and the error $\boldsymbol{\xi}$. Think in terms of the scalar case,

$$\xi^k = (1 - gc)a\xi^{k-1} + (1 - gc)bw^k - gn^k, \quad \sigma^k = \varphi^2\sigma^{k-1} + \psi^2q + g^2r. \quad (24)$$

Both disturbances are always, at every step, acting to increase the covariance. We can assume that the system is stable, so φ has the opposite effect, tending to decay the error and its covariance. On the one hand, we want to put the gain high, so that the error decays rapidly: say, the upper limit of $g = 1/c$, giving $\varphi = 0$ and $\psi = 0$. On the other hand, so high a gain will amplify the measurement noise n , and this is ultimately counterproductive. The optimal gain minimizes the covariance σ^k , such that

$$\frac{\partial \sigma^k}{\partial g} = -2(1 - gc)ca^2\sigma^{k-1} - 2(1 - gc)cb^2q + 2gr = 0 \quad (25)$$

giving

$$g = \frac{ca^2\sigma^{k-1} + cb^2q}{c(ca^2\sigma^{k-1} + cb^2q) + r}. \quad (26)$$

Let us now return to the matrix case. Here the gain is chosen so as to minimize (19), the trace of $\boldsymbol{\Sigma}^k$. That is, each element of the gain matrix is chosen such that

$$\frac{\partial \text{Tr}\{\boldsymbol{\Sigma}^k\}}{\partial g_{ij}} = \frac{\partial}{\partial g_{ij}} \text{Tr}\{\boldsymbol{\Phi}\boldsymbol{\Sigma}^{k-1}\boldsymbol{\Phi}^T + \boldsymbol{\Psi}\mathbf{Q}\boldsymbol{\Psi}^T + \mathbf{G}\mathbf{R}\mathbf{G}^T\} = 0. \quad (27)$$

Now,

$$\frac{\partial}{\partial a_{ij}} \text{Tr}\{\mathbf{A}\mathbf{B}\mathbf{A}^T\} = b_{jp}a_{pi}^T + a_{ip}b_{pj}, \quad (28)$$

so if \mathbf{B} is symmetric, as is the case in our problem,

$$\frac{\partial}{\partial a_{ij}} \text{Tr}\{\mathbf{A}\mathbf{B}\mathbf{A}^T\} = 2a_{ip}b_{pj}. \quad (29)$$

Also,

$$\frac{\partial}{\partial a_{ij}} \text{Tr}\{\mathbf{A}\mathbf{B}\} = b_{ji}, \quad \frac{\partial}{\partial a_{ij}} \text{Tr}\{\mathbf{B}\mathbf{A}^T\} = b_{ij} \quad (30)$$

We proceed term-by-term. The first two terms are similar,

$$\frac{\partial}{\partial g_{ij}} \text{Tr}\{\boldsymbol{\Phi}\boldsymbol{\Sigma}^{k-1}\boldsymbol{\Phi}^T\} = -2(\mathbf{A}\boldsymbol{\Sigma}^{k-1}\mathbf{A}^T\mathbf{C}^T)_{ij} + 2(\mathbf{G}\mathbf{C}\mathbf{A}\boldsymbol{\Sigma}^{k-1}\mathbf{A}^T\mathbf{C}^T)_{ij} \quad (31)$$

and

$$\frac{\partial}{\partial g_{ij}} \text{Tr}\{\boldsymbol{\Psi}\mathbf{Q}\boldsymbol{\Psi}^T\} = -2(\mathbf{B}\mathbf{Q}\mathbf{B}^T\mathbf{C}^T)_{ij} + 2(\mathbf{G}\mathbf{C}\mathbf{B}\mathbf{Q}\mathbf{B}^T\mathbf{C}^T)_{ij}. \quad (32)$$

The third term becomes

$$\frac{\partial}{\partial g_{ij}} \text{Tr}\{\mathbf{G}\mathbf{R}\mathbf{G}^T\} = 2(\mathbf{G}\mathbf{R})_{ij}. \quad (33)$$

Together,

$$\mathbf{G} = (\mathbf{A}\boldsymbol{\Sigma}^{k-1}\mathbf{A}^T + \mathbf{B}\mathbf{Q}\mathbf{B}^T)\mathbf{C}^T [\mathbf{C}(\mathbf{A}\boldsymbol{\Sigma}^{k-1}\mathbf{A}^T + \mathbf{B}\mathbf{Q}\mathbf{B}^T)\mathbf{C}^T + \mathbf{R}]^{-1} \quad (34)$$

and it is evident how (34) preserves the form of the scalar equation (26).

The filter equations (18), (23), and (34) can be implemented in real-time state observation, starting with estimates of $\mathbf{x}^{*,0}$ and $\boldsymbol{\Sigma}^0$. Our use of these equations is *design*: we want to find the gain matrix \mathbf{G} – the structure of the state observer – at a set of specified operating points. The conditions at each

operating point are assumed to be stochastically stationary. Rather than tracing the time evolution of Σ , this is prescribed to be constant. As a result, (23) becomes

$$\Phi \Sigma \Phi^T - \Sigma + \Psi Q \Psi^T + \mathbf{G} \mathbf{R} \mathbf{G}^T = \mathbf{0}, \quad (35)$$

which is to be solved together with (34) for Σ and \mathbf{G} .

There are two numerical algorithms implemented in STAS for the solution of (35).¹³ A slow-but-reliable method is to treat the stationary problem as a transient one, making an initial guess for Σ^0 and solving (23) and (34) in time. Another possibility, with faster convergence, is to iterate between the gain and covariance equations:

1. Guess an initial gain matrix \mathbf{G}^0 . For the method to work, \mathbf{G}^0 must be such that the system is stable; that is, all the eigenvalues of $\Phi^0 = (\mathbf{I} - \mathbf{G}^0 \mathbf{C}) \mathbf{A}$ have a magnitude less than 1. Set $k = 0$ and $\mathbf{G}^k = \mathbf{G}^0$.
2. Solve the discrete Lyapunov equation (35) for Σ^k , using the present gain matrix \mathbf{G}^k .
3. Solve the gain equation (34) for \mathbf{G}^{k+1} , using Σ^k .
4. Check for convergence, which can be based on the change in the gain and/or covariance matrices from one iteration to the next. Increment k and return to Step 2.

If we were rigorously following Kalman, which we are not, then \mathbf{Q} is the covariance of the external disturbances acting on the system; \mathbf{R} is the covariance of the disturbances or noise on the sensor measurements; and \mathbf{G} is the optimal gain that minimizes the covariance of the error in the state estimate. One problem with this approach is that the atmospheric turbulence and ocean waves are not white noise. It is possible to shape the frequency response in a rudimentary way by augmenting the state space like (9) and (16) – in essence, telling the observer that we expect the disturbances to behave in this way. The result is a rough approximation, not precise, and it is required to tune the damping parameters to provide an appropriate weight on peak frequencies.

The main downside to the formal Kalman filter is the characterization of measurement noise \mathbf{R} , which is poorly defined in general. Even if one rigorously calibrated a noise model to trial measurements, there are a range of factors, including more obscure forms of uncertainty, that call for higher level of “noise” than that which would be associated with high-quality sensors. The gains result from a tradeoff between process and measurement noise; but the measurement noise is, to some extent, a set of arbitrary tunable parameters. A practical approach is therefore to treat the weights in \mathbf{R} as arbitrary tunable parameters, and iterate until the observer performance, and robustness, is satisfactory. It is usually acceptable to define \mathbf{R} as a diagonal matrix, so there are as many tunable parameters as sensors.

On initial trials, it was found that the “optimal” filter derived from the above procedure had a potential flaw: the error $\bar{\mathbf{y}} - \mathbf{y}^*$ was not driven to zero; that is, given a steady-state input, there appeared a steady-state offset between the sensor measurements and estimates. This is to be expected: say that a step function is given as input for \mathbf{w} ; then, an expression for the steady-state sensor error, based on (17) and (18), is

$$\bar{\mathbf{y}} - \mathbf{y}^* = \mathbf{C} [\mathbf{I} - \mathbf{A}(\mathbf{I} - \mathbf{A} + \mathbf{G} \mathbf{C} \mathbf{A})^{-1} \mathbf{G} \mathbf{C}] (\mathbf{I} - \mathbf{A})^{-1} \mathbf{B} \mathbf{w}. \quad (36)$$

The error approaches zero only in the limit, as the gains \mathbf{G} become large. Nonzero steady-state error is an annoyance, since it is clearly incorrect, and non-physical in the case of variables like nacelle

¹³Stengel (1994). The original references for the iterative method are two letters to the IEEE Transactions on Automatic Control: Kleinman (1968) and Hwer (1971).

velocity.¹⁴ To fix this, an integral term was added to each of the sensor errors; essentially, the sensor error $\bar{\mathbf{y}} - \mathbf{y}^*$ is passed through a PI controller, with $\mathbf{K}_P = \mathbf{I}$. The observer (18) becomes

$$\begin{bmatrix} \mathbf{x}^{*,k} \\ \mathbf{x}_I^k \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{G} \\ \mathbf{I} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{x}^{*,k-1} \\ \mathbf{x}_I^{k-1} \end{bmatrix} + \begin{bmatrix} \mathbf{G} \\ \mathbf{K}_I \end{bmatrix} (\bar{\mathbf{y}}^k - \mathbf{y}^{*,k}), \quad (37)$$

where \mathbf{K}_I is a diagonal matrix of integral gains. The integral gains are small, so that they do not interfere noticeably with the “optimal” state estimates, but rather slowly eliminate residual steady-state errors.

4.5 Estimating the weather and turbulence

In order to be consistent with the spectral analysis of Section 4.6, turbulence is defined as a fluctuation with respect to a mean value. This implies a *separation of scales*. There is a background flow associated with what we might call the “weather”, upon which is superposed fluctuating eddies associated with the atmospheric boundary layer, wakes from other turbines, thermal instabilities, and so on. This view of atmospheric flow is useful, but has its flaws. In reality, the weather is always shifting: often gradually – many locations experience a diurnal cycle in wind conditions, for instance – but sometimes abruptly, during convective conditions and the passage of fronts. The questions are then, what is the “mean wind”, what is the “turbulent” component of V (the effective wind speed over the rotor), and what is fine-scale turbulence versus measurement noise?

Consider the information at our disposal. At each turbine, the effective wind speed $V(t)$ is estimated from two sources: the rotor dynamics, which are an integral part of the turbine state model; and the anemometer wind speed measurement on the nacelle housing. The former is reliable under normal operating conditions, in particular when there is little yaw misalignment and when the rotor aerodynamic model is well-calibrated to the actual turbine performance.¹⁵ The anemometer measurement is sensitive to local wind conditions at the nacelle, and so it is noisy.

The weather consists of features in the atmospheric flow that are on the scale of the wind plant, or larger. The weather can therefore be identified by comparing the winds measured over clusters of wind turbines. In fact, this can be used as a definition, for purposes of wind plant control: *the weather is that component of the measured wind speed and direction that is consistent across a cluster of several adjacent wind turbines*. Note that this includes

1. what we would normally consider the “weather”;
2. the flow on each side of a passing front, but not the frontal boundary itself; and,
3. persistent gradients in the wind speed associated with the extraction of energy from the atmospheric boundary layer, but not local, meandering wake effects.

From this definition of weather, it follows that the “turbulence” is that component of the measured wind that is *not* consistent from one turbine to the next in the cluster.

So the key to identifying μ_V , and by extension σ_V , I , L_u , and S_V , is to establish consensus among the turbines in a cluster. Here it pays to take a Bayesian approach; this can handle things like shut-down turbines, sensor miscalibration, and other faults that would throw off a simple averaging of the signals.

¹⁴There may be drift in the input sensor measurements, and it is assumed that this has been removed; here we are talking about nonzero errors that are an artifact of the observer model.

¹⁵It is usually possible to arrive at a well-calibrated model of the rotor aerodynamics, which can be evaluated by comparing the computed and measured electrical power over long periods of operation. The performance can change over time, in particular due to leading-edge fouling and erosion. The ambient air density must also be taken into account.




Figure 13: Examples of two turbine clusters, used to estimate the “mean” wind speed and direction at the starred turbines.

4.5.1 Clusters of turbines

It is assumed that the wind turbines are arranged in an array of approximately uniform spacing. The layout may be regular or irregular. For each turbine, a cluster is defined as the turbine plus its nearest eight neighbors. Figure 13 shows an example of two such clusters on a regular, staggered layout.

The choice of eight neighboring turbines is somewhat arbitrary. It is a reasonable choice, though, for the following reasons. The cluster spans three rows and columns in a regular gridded layout. It is evident in Fig. 13 that for a plant of a moderate size such a cluster achieves a balance between being local and encompassing a significant sector of the plant. For instance, if there was a steady gradient in the wind speed across the plant, then the values of μ_V at the blue and red starred turbines would correctly differ. Nine turbines (one plus eight neighbors) provide multiple redundancy, with the possibility of detecting and eliminating faulty or miscalibrated sensors; it is unlikely that more than four out of nine sensors would be severely miscalibrated. For large commercial wind turbines, and typical plant spacings of $s = 5D$ to $8D$, the convection time $2s/\mu_V$ across the cluster is on the order of several minutes. This is a reasonable timescale at which to distinguish between random, local turbulence and more persistent features in the flow. The tradeoff is essentially one of consistency during steady weather conditions, and achieving a faster response to dynamic changes in the weather.¹⁶

4.5.2 Estimation of the cluster wind speed and direction

A Bayesian approach works explicitly with probability distributions over the possible values of the cluster wind speed μ_V and direction μ_θ . Fluctuations in the speed and direction are assumed to be

¹⁶Annoni *et al.* (2019) applied a clustering algorithm to an onshore wind plant with an irregular layout, and found that a 15-turbine cluster produced minimum errors in the estimate of wind direction. Nine-turbine clusters produced a slightly higher error: 2.8° average error for 9 turbines, versus 2.5° for 15 turbines, in the particular case considered.

independent. In reality this assumption is not rigorously true, nor entirely false – but it is practical. The way that the algorithm incorporates sensor measurements compensates for imprecision in the stochastic model, which in any case is highly simplified.

The procedure consists of propagating the probability distribution in time, incorporating the wind speed estimates provided by the turbine state observers. The probability distribution is low-dimensional, due to some fortunate features of the problem. Specifically, that we are to estimate a single μ_V and μ_θ for the cluster, and that these are independent, means that we are dealing with two separate scalar (one-dimensional) probability distributions. In addition, the turbulent fluctuations – the deviations of the individual turbine signals with respect to the cluster μ_V or μ_θ – are expected to be independent and uncorrelated. This means that it is not necessary to formulate a high-dimensional joint probability distribution over the measurements.

There are a variety of algorithms we could use to solve for the probability distribution at each timestep. The most common algorithms would work with a linear system model and Gaussian probabilities, and propagate an expected value and variance. There are also advanced Monte-Carlo type algorithms that approximate the stochastic properties through repeated realizations. However, because the problem deals with low-dimensional probability distributions, it is feasible to apply the cell-to-cell mapping technique (CCMT).¹⁷ This is a general solution to stochastic problems, both linear and nonlinear, whose numerical implementation consists of matrix operations. CCMT suffers from the “curse of dimensionality”, but for low-dimensional problems it is efficient enough for real-time applications; and it is easy to program.

Let us walk through the procedure for a single timestep; this is then repeated to advance in time. The procedure is the same for μ_V and μ_θ ; we will consider μ_V as an illustration. The input available at the start of the k^{th} timestep includes the probability density $\varphi(\mu_V^{k-1})$ from the previous timestep, and estimates \bar{V}_i^k , obtained from the state observers associated with the $i = \{1, 2, \dots, n\}$ turbines in the cluster.

The first step is to update $\varphi(\mu_V)$ from timestep $k - 1$ to k , prior to incorporating the new measurements. We are free to choose the update equation. The choice is constrained by the information available as input to the timestep. If we had a lot of information – say, lidar measurements of the wind field ahead of each turbine – then the update equation could be based on a quite elaborate stochastic or physical flow model. For the present controller design, where simplicity is desirable, and only basic sensor measurements are available, we represent the drift in μ_V as discrete, binary white noise. The formula is

$$\hat{\mu}_V^k = \mu_V^{k-1} + \gamma \Delta t \delta^k, \quad (38)$$

where $\delta^k \in \{1, -1\}$ with probability $p(\delta) \in \{0.5, 0.5\}$ is binary white noise, and γ is a tunable parameter. In effect, (38) says that “ $\varphi(\mu_V)$ diffuses”; and γ controls the strength of diffusion.

According to CCMT, we divide the possible values of $\hat{\mu}_V$ into discrete bins, or cells. There is a probability density $\varphi(\hat{\mu}_{V,m})$ or total probability $p(\hat{\mu}_{V,m}) = \varphi(\hat{\mu}_{V,m}) \Delta\mu_V$ associated with each cell. Applying (38) to each cell in turn, considering the possible input values δ with their probabilities $p(\delta)$, a matrix equation can be written,

$$\boldsymbol{\varphi}(\hat{\mu}_V^k) = \mathbf{A}\boldsymbol{\varphi}(\mu_V^{k-1}). \quad (39)$$

Here $\boldsymbol{\varphi}$ collects the $\varphi(\hat{\mu}_{V,m})$ into a column vector. Thus the update takes the form of a sparse matrix-vector multiplication.

The second step is to incorporate the information contained in the measurements. This is done using Bayesian inference:

$$\varphi(\mu_V^k | \bar{V}_1^k, \bar{V}_2^k, \dots, \bar{V}_n^k) = \frac{\varphi(\bar{V}_1^k, \bar{V}_2^k, \dots, \bar{V}_n^k | \hat{\mu}_V^k) \varphi(\hat{\mu}_V^k)}{\sum_m \varphi(\bar{V}_1^k, \bar{V}_2^k, \dots, \bar{V}_n^k | \hat{\mu}_{V,m}^k) \varphi(\hat{\mu}_{V,m}^k)}. \quad (40)$$

¹⁷Hsu CS (1987)




Figure 14: A two-parameter distribution representing the uncertainty in measurements at a single turbine, given a value of the cluster average.

But here the assumptions of independence some into play, considerably simplifying things. We can write (40) as

$$\varphi(\mu_V^k | \bar{V}_1^k, \bar{V}_2^k, \dots, \bar{V}_n^k) = \frac{\varphi(\bar{V}_1^k | \hat{\mu}_V^k) \varphi(\bar{V}_2^k | \hat{\mu}_V^k) \dots \varphi(\bar{V}_n^k | \hat{\mu}_V^k) \varphi(\hat{\mu}_V^k)}{\sum_m \varphi(\bar{V}_1^k | \hat{\mu}_{V,m}^k) \varphi(\bar{V}_2^k | \hat{\mu}_{V,m}^k) \dots \varphi(\bar{V}_n^k | \hat{\mu}_{V,m}^k) \varphi(\hat{\mu}_{V,m}^k)}. \quad (41)$$

Evidently, solving (41) requires a model for $\varphi(\bar{V} | \mu_V)$: the probability of measuring a particular \bar{V} at one of the turbines, if the “true” cluster value were μ_V . The difference between the cluster-wide wind speed and the value observed at a single turbine may be due to the real effect of local turbulence, or else errors in the measurement: noise, bias, or sensor failure.

The arguments leading up to (38) apply here as well: we get to choose the model for $\varphi(\bar{V} | \mu_V)$, simple or elaborate. Using a simple model is practical, since the measurements continually serve to correct the modelling errors; though it could be expected that a well-designed elaborate model would increase the precision. We elect to go simple, defining an empirical two-parameter distribution shown in Fig. 14. This consists of a normal distribution and a constant “background” probability. The normal distribution, scaling with the parameter σ , is centered about the true (or given) value of μ_V , and accounts for turbulence as well as measurement errors. The background probability, a constant β , spans the possible range of $\bar{V} - \mu_V$. It accounts for the possibility of sensor failure or other anomalies.

The range of \bar{V} is fixed at ± 50 m/s; this is arbitrary, in that the values of the other parameters can be adjusted to achieve any desired relative weighting between the background and normal distributions. Formally, since $\varphi(\bar{V} | \mu_V)$ is a probability distribution, the integral under the curve is equal to 1. However, (41) is a ratio between like quantities, so $\varphi(\bar{V} | \mu_V)$ can be scaled by any uniform factor without affecting the result.

It is well to keep in mind that the cluster values μ_V and μ_θ are computed in turn for each turbine in the plant. In general, there is a unique cluster associated with each turbine. Even though one turbine and its neighbor are part of each others’ clusters, it is not so that their values of μ_V and μ_θ are the same.

4.5.3 Time series of turbulence

Once the cluster mean wind speed and direction is found using the methods of Section 4.5.2, a time series of turbulence at each turbine can be generated by computing

$$v_i = \bar{V}_i - \mu_V \quad \text{and} \quad \theta_i = \bar{\theta}_i - \mu_\theta. \quad (42)$$

4.6 Turbulence and response spectra

Turbulence and response spectra are computed based on the time series (42), together with a numerical model that fills in the higher frequencies and smaller spatial scales, which cannot be resolved with the basic sensor measurements. The time series is logged over N discrete timesteps, the most recent measurement replacing the last in the sequence. This can be represented as a state space where

$$x_{i+1}^k = x_i^{k-1}, \quad (43)$$

x representing either v or θ . N should be a power of 2, in order to facilitate the numerical Fourier transform. The timestep is Δt , and the length of time over which data is logged is $T = N \Delta t$.

4.6.1 Measured turbulence spectrum

Given a time series $x(k \Delta t)$, the spectrum is computed by

1. mirroring the time series to make it periodic;
2. computing the Fourier coefficients $\kappa(k \Delta f)$, where $\Delta f = 1/(N \Delta t)$; and,
3. computing the spectrum

$$S(k \Delta f) = \frac{1}{\Delta f} \kappa \kappa^*, \quad (44)$$

where κ^* is the complex conjugate of κ .

4.6.2 Theoretical turbulence spectrum and parameters

The stochastic model for the turbulent wind field requires two parameters, the turbulence intensity

$$I := \frac{\sigma_V}{\mu_V}, \quad (45)$$

that is, the standard deviation divided by the mean; and a length scale L_u , representing in some sense the size of the energy-carrying turbulent eddies. We assume an isotropic Von Karman spectrum,¹⁸

$$S_V^o(f) = \sigma_V^2 \frac{4L_u}{\mu_V \left[1 + 70.8 \left(\frac{L_u f}{\mu_V} \right)^2 \right]^{5/6}}. \quad (46)$$

This is a convenient spectrum to use, since there are analytical formulas for the covariance between the three velocity components at any two points in space.¹⁹ Figure 15 shows an example of the Von Karman spectrum.

¹⁸The superscript “o” on S_V^o emphasizes that this is the one-sided spectrum ($0 \leq f < \infty$), twice the amplitude of the two-sided spectrum ($-\infty < f < \infty$).

¹⁹Merz (2015a), Kristensen and Frandsen (1982), Connell (1982)




Figure 15: An example of the Von Karman spectrum, $\mu_V = 10 \text{ m/s}$, $I = 0.183$, $L_u = 180 \text{ m}$.

In operation, the turbulence intensity and length scale are calibrated such that the zero'th and second spectral moments

$$m_0 = \int_0^\infty S_V^o df = \sigma_V^2 = (IV_\infty)^2 \quad \text{and} \quad m_2 = \int_0^\infty f^2 S_V^o df \quad (47)$$

match those observed at each wind turbine.²⁰ There is a catch, though: the observations are made on the *effective rotor-average* wind speed, whereas the theoretical spectrum needs to be that at a single point. It is possible, since the theoretical covariance in space and time is known, to compute the theoretical spectrum of rotor-average wind speed, and the resulting m_0 and m_2 . These can then be related to the single-point parameters I and L_u . For the DTU 10 MW wind turbine ($D = 178.3 \text{ m}$) used in the present analysis, the following relations hold: first,

$$y = 1.3057646 + 0.3628023x + 0.2079409x^2, \quad -4 \leq x \leq -1.8, \quad (48)$$

with

$$y = \log_{10} L_u \quad \text{and} \quad x = \log_{10} \left(\frac{m_2}{m_0} \right), \quad (49)$$

which can be solved for L_u ; and next,

$$y = 0.1852021x - 0.1479345x^2, \quad 1.3 \leq x \leq 3.3, \quad (50)$$

with

$$y = \log_{10} \left(1 - \frac{m_0}{\sigma_V^2} \right) \quad \text{and} \quad x = \log_{10} L_u, \quad (51)$$

which can be solved for I , via σ_V .

4.6.3 Rotationally-sampled turbulence spectra

After the nominal Von Karman spectrum is determined, with its characteristic turbulence intensity and length scale, the spectrum of turbulence seen by the rotating blades can be computed analytically.

²⁰The frequency $\sqrt{f_2/f_0}$ is the zero-crossing frequency of a narrow-banded Gaussian stochastic process. Now, turbulence is wide-banded, so $\sqrt{f_2/f_0}$ is not its zero-crossing frequency; but it gives some motivation for the choice of spectral moments in the calibration.

The method is documented in TotalControl D3-3.²¹ The key assumption is that the turbulence is isotropic and convects at the effective wind speed.

In addition to the stochastic turbulence, periodic effects like wind shear and tower shadow are also modelled. The periodic effects are not observable from the basic sensor measurements. Tower shadow can be estimated theoretically by assuming a potential flow pattern upstream of the tower, together with undeformed blades. For wind shear, it is necessary either to use a nominal theoretical value for the wind speed gradient, or to consider the wind shear as an “external” input provided by the plant operator, who, it is presumed, has access to additional sensors like a meteorological mast.

To reiterate the key concept, it is that the properties of a nominal, theoretical turbulence spectrum are calibrated to the observed wind speed measurements; and then the details of the stochastic and periodic flow fields, as seen by the rotating blades, are filled in using theory.

The result is a pair of spectral matrices, with components associated with the flow at each blade element: $\mathbf{S}_V(f)$ accounting for the stochastic turbulence, and $\tilde{\mathbf{S}}_V(f)$ the periodic fluctuations. These are kept separate, for reasons that will be made clear in later sections.

4.6.4 Scaling of a reference spectrum

The rotationally-sampled turbulence spectrum will scale with the severity of turbulence and with the rotor speed. The amplitude of the spectrum scales as the square of the turbulence intensity $I = \sigma_V/V_\infty$. A special scaling technique is required in order to get the proper scaling with rotor speed.²² Under combined perturbations in turbulence intensity and rotor speed, the amplitude scales as

$$S_v(\Omega + \Delta\Omega, f) = \kappa \left(\gamma^2 S_v(\Omega, \kappa f) - \frac{\partial S_v(\Omega, \kappa f)}{\partial v} \frac{V_\infty}{\Omega} \Delta\Omega \right), \quad (52)$$

where

$$\kappa = \frac{\Omega}{\Omega + \Delta\Omega} \quad (53)$$

and

$$\gamma = \frac{I}{I_0}. \quad (54)$$

4.6.5 Ocean wave force spectra

The ocean wave force spectra \mathbf{S}_w are computed numerically from a time series of the observed values. The procedure of Section 4.6.1 is applied here as well. The ocean wave direction changes slowly, and in the model it is assumed to be constant. It is assumed that the dominant wave direction is available as an external input.

4.6.6 Transfer functions and response spectra

Defining a transfer function between the environmental inputs \mathbf{v} and some outputs \mathbf{z} ,

$$\mathbf{z} = \sum_{k=-N/2}^{N/2-1} \mathbf{H}_k \mathbf{v}_k \exp(-i2\pi kp/N), \quad (55)$$

the spectral matrix transforms as

$$\mathbf{S}_z(k \Delta f) = \mathbf{H}_k \mathbf{S}_v(k \Delta f) \mathbf{H}_k^{*T}. \quad (56)$$

²¹Merz (2020)

²²Merz *et al.* (2019)

4.6.7 Response spectra in fixed and rotating coordinates

The linearized equations of motion of a three-bladed wind turbine can be made time-invariant with a multi-blade coordinate transform. This expresses the rotor dynamics in a fixed (non-rotating) coordinate frame. The state equations (3) are implemented in multi-blade coordinates. Sensor inputs associated with each blade – in the present state observer, this is only the blade pitch angle – are transformed to multi-blade coordinates before being fed to the state observer. Solving for response spectra in the fixed coordinate frame is then straightforward. In order to compute the material stresses required for fatigue analysis (Section 4.7), driveshaft and blade load spectra need to be transformed into the rotating frame.

Consider a trio of variables $\mathbf{r}^B = [r_1, r_2, r_3]$, with each variable associated with the blade indicated by the subscript.²³ The multi-blade coordinate transform is $\mathbf{r}^\psi = \mathbf{T}_B^\psi \mathbf{r}^B$, where

$$\mathbf{T}_B^\psi = \frac{1}{3} \begin{bmatrix} 1 & 1 & 1 \\ 2 \cos \Psi_1 & 2 \cos \Psi_2 & 2 \cos \Psi_3 \\ 2 \sin \Psi_1 & 2 \sin \Psi_2 & 2 \sin \Psi_3 \end{bmatrix} \quad (57)$$

with inverse

$$\mathbf{T}_\psi^B = \begin{bmatrix} 1 & \cos \Psi_1 & \sin \Psi_1 \\ 1 & \cos \Psi_2 & \sin \Psi_2 \\ 1 & \cos \Psi_3 & \sin \Psi_3 \end{bmatrix}. \quad (58)$$

The blade azimuth angles are

$$\Psi_k = \Psi_0 + \Omega t + \frac{2\pi}{3}(k-1). \quad (59)$$

It will be convenient to write (58) in complex exponential form,

$$\mathbf{T}_\psi^B = \mathbf{E}_0 + \mathbf{E} \exp(i\Omega t) + \mathbf{E}^* \exp(-i\Omega t), \quad (60)$$

with

$$\mathbf{E}_0 = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{E} = \frac{1}{2} \begin{bmatrix} 0 & 1 & -i \\ 0 & \exp(i2\pi/3) & -i \exp(i2\pi/3) \\ 0 & \exp(i4\pi/3) & -i \exp(i4\pi/3) \end{bmatrix} \exp(i\Psi_0). \quad (61)$$

Let us define the variable \mathbf{r}^ψ in the following way:

$$\mathbf{r}^\psi = \overline{\mathbf{r}^\psi} + \Delta \mathbf{r}^\psi + \tilde{\Delta} \mathbf{r}^\psi. \quad (62)$$

Here $\overline{\mathbf{r}^\psi}$ contains the steady-state values, the components of $\Delta \mathbf{r}^\psi$ are zero-mean Gaussian processes, and $\tilde{\Delta} \mathbf{r}^\psi$ is periodic. We are given a spectral matrix of $\Delta \mathbf{r}^\psi$, $\mathbf{S}^\psi(f)$, or in discrete form $\mathbf{S}^\psi(k \Delta f) = \mathbf{S}^\psi(k)$. We are also given the steady-state $\overline{\mathbf{r}^\psi}$, as well as the Fourier coefficients \mathbf{c}_k^ψ describing $\tilde{\Delta} \mathbf{r}^\psi$. The goal is to transform these quantities to body (rotating) coordinates. That is, we seek an equivalent form of (62) in body coordinates,

$$\mathbf{r}^B = \overline{\mathbf{r}^B} + \Delta \mathbf{r}^B + \tilde{\Delta} \mathbf{r}^B, \quad (63)$$

where the zero-mean stochastic term $\Delta \mathbf{r}^B$ is characterized by a spectral matrix $\mathbf{S}^B(k)$, and the zero-mean periodic term $\tilde{\Delta} \mathbf{r}^B$ by Fourier coefficients \mathbf{c}_k^B .

To develop the appropriate transformations, let's start in the time domain. The spectra are each the individual Fourier transform of a covariance function; the time-domain counterpart to \mathbf{S}^ψ is a covariance function \mathbf{Q}^ψ . By definition,

$$\mathbf{Q}_{jk}^\psi(t, \tau) = E \left[\left(r_i^\psi(t) - \overline{r_i^\psi} \right) \left(r_j^\psi(t + \tau) - \overline{r_j^\psi} \right) \right], \quad (64)$$

²³These could be any of the dynamic variables: displacements, loads, stresses, or whatever.

From (62),

$$Q_{jk}^{\psi}(t, \tau) = E \left[\Delta r_i^{\psi}(t) \Delta r_j^{\psi}(t + \tau) \right] + E \left[\Delta r_i^{\psi}(t) \tilde{\Delta} r_j^{\psi}(t + \tau) \right] \\ + E \left[\tilde{\Delta} r_i^{\psi}(t) \Delta r_j^{\psi}(t + \tau) \right] + E \left[\tilde{\Delta} r_i^{\psi}(t) \tilde{\Delta} r_j^{\psi}(t + \tau) \right]. \quad (65)$$

It is necessary to make a brief digression to specify what is meant by the expected value. In the context of random variables, the fundamental definition of an expected value is

$$E[x] = \int_{-\infty}^{\infty} x \varphi(x) dx, \quad (66)$$

where $\varphi(x)$ is the probability density over x . It is useful to picture (66) in discrete space – that is, discretized values of x – in which case

$$E[x] = \sum_{k=-n/2}^{n/2-1} x_k \varphi(x_k) \Delta x = \sum_{k=-n/2}^{n/2-1} x_k p(x_k), \quad (67)$$

with $p(\cdot)$ the cell probability. Blind application of (66) to a deterministic function simply returns the value of the function: making use of the Law of the Unconscious Statistician,

$$E[f(x_p)] = \sum_{k=-n/2}^{n/2-1} f(x_k) p(x_k) = \begin{cases} 0, & k \neq p \\ f(x_p) & k = p \end{cases} \quad (68)$$

But this is not really what we want. Rather, given that we are working with a stationary stochastic process and repeating periodic signal, the appropriate definition of the expected value is the *ensemble average over a set of arbitrary times*. That is to say, we consider t – the starting time, in the case of the covariance (65) – to be a uniformly-distributed random variable, and then we evaluate the expected value according to (66). In this case, *the expected value becomes the time-average*. To see this, write

$$E[f(t)] = \lim_{t \rightarrow \infty} \int_{-T/2}^{T/2} f(t) \varphi(t) dt = \lim_{t \rightarrow \infty} \frac{1}{T} \int_{-T/2}^{T/2} f(t) dt. \quad (69)$$

By the ergodic theorem, (66) and (69) are the same, for a stationary stochastic process. This means that for both stochastic and periodic quantities, we may define the expected value as the average value over time.

Picking up Eq. (65) where we left off, it is clear that the two terms in the middle are zero. They are independent, so

$$E \left[\Delta r_i^{\psi}(t) \tilde{\Delta} r_j^{\psi}(t + \tau) \right] = E \left[\Delta r_i^{\psi}(t) \right] E \left[\tilde{\Delta} r_j^{\psi}(t + \tau) \right] = 0, \quad (70)$$

and similarly for the other term. Then,

$$Q_{jk}^{\psi}(t, \tau) = E \left[\Delta r_i^{\psi}(t) \Delta r_j^{\psi}(t + \tau) \right] + E \left[\tilde{\Delta} r_i^{\psi}(t) \tilde{\Delta} r_j^{\psi}(t + \tau) \right]. \quad (71)$$

We are therefore free to consider the stochastic and periodic processes separately. Should we later wish to form the combined covariance or spectra, these can be obtained by superposition.

Let us simplify the terminology by making the time t implicit. So, for instance, we will write (71) as

$$Q_{jk}^{\psi}(\tau) = E \left[\Delta r_i^{\psi}(0) \Delta r_j^{\psi}(\tau) \right] + E \left[\tilde{\Delta} r_i^{\psi}(0) \tilde{\Delta} r_j^{\psi}(\tau) \right]. \quad (72)$$

The development that follows is best accomplished in discrete time, and discrete frequencies, since this is how it is programmed. Working with discrete increments also allows us to avoid integrals,

limits, and delta functions, and the operations are easier to visualize. Let $\tau = p \Delta\tau$ and $f = k \Delta f$. The discrete form of the complex exponential is then

$$\exp(i\omega\tau) = \exp(i2\pi k \Delta f p \Delta\tau) = \exp\left(i2\pi \frac{k \Delta f p}{N \Delta f}\right) = \exp\left(i2\pi \frac{kp}{N}\right). \quad (73)$$

In a similar spirit, let's define $\Omega = w \Delta\tau$, and

$$\exp(i\Omega\tau) = \exp\left(i2\pi \frac{wp}{N}\right), \quad (74)$$

where w does not necessarily need to be an integer.

Now, we want to arrive at the Fourier representation of $\tilde{\Delta}\mathbf{r}^B$,

$$\tilde{\Delta}\mathbf{r}^B(p) = \sum_{k=-N/2}^{N/2-1} \mathbf{c}_k^B \exp\left(i2\pi \frac{kp}{N}\right) \quad (75)$$

along with the covariance in body coordinates,

$$\mathbf{Q}^B(p) = E \left[\left(\mathbf{r}^B(0) - \overline{\mathbf{r}^B} \right) \left(\mathbf{r}^B(p \Delta\tau) - \overline{\mathbf{r}^B} \right)^T \right]. \quad (76)$$

To do so, we need to first find the steady-state, or mean, $\overline{\mathbf{r}^B}$. Since $\mathbf{r}^B(t) = \mathbf{T}_\psi^B(\Omega, t) \mathbf{r}^\psi(t)$, we have

$$\overline{\mathbf{r}^B} = E \left[\mathbf{T}_\psi^B(\Omega, t) \mathbf{r}^\psi(t) \right]. \quad (77)$$

Using (60) for the transform and (62) for \mathbf{r}^ψ ,

$$\overline{\mathbf{r}^B} = E \left[\left(\mathbf{E}_0 + \mathbf{E} \exp\left(i2\pi \frac{wp}{N}\right) + \mathbf{E}^* \exp\left(-i2\pi \frac{wp}{N}\right) \right) \left(\overline{\mathbf{r}^\psi} + \Delta\mathbf{r}^\psi(p \Delta\tau) + \tilde{\Delta}\mathbf{r}^\psi(p \Delta\tau) \right) \right]. \quad (78)$$

Consider this term-by-term:

$$E \left[\left(\mathbf{E}_0 \overline{\mathbf{r}^\psi} \right) \right] = \mathbf{E}_0 \overline{\mathbf{r}^\psi}, \quad (79)$$

since both quantities are constant;

$$E \left[\mathbf{E} \overline{\mathbf{r}^\psi} \exp\left(i2\pi \frac{wp}{N}\right) \right] = \mathbf{0} \quad \text{and} \quad E \left[\mathbf{E}^* \overline{\mathbf{r}^\psi} \exp\left(-i2\pi \frac{wp}{N}\right) \right] = \mathbf{0}; \quad (80)$$

and,

$$E \left[\left(\mathbf{E}_0 + \mathbf{E} \exp\left(i2\pi \frac{wp}{N}\right) + \mathbf{E}^* \exp\left(-i2\pi \frac{wp}{N}\right) \right) \Delta\mathbf{r}^\psi(p \Delta\tau) \right] = \mathbf{0}, \quad (81)$$

since, even if there is some energy in the random process at the rotor frequency, the phase is undetermined. As for the periodic part $\tilde{\Delta}\mathbf{r}^\psi$, it can be expanded in a Fourier series akin to (75),

$$\tilde{\Delta}\mathbf{r}^\psi(p) = \sum_{k=-N/2}^{N/2-1} \mathbf{c}_k^\psi \exp\left(i2\pi \frac{kp}{N}\right). \quad (82)$$

Then,

$$E \left[\mathbf{E}_0 \sum_{k=-N/2}^{N/2-1} \mathbf{c}_k^\psi \exp\left(i2\pi \frac{kp}{M}\right) \right] = \mathbf{0} \quad (83)$$

but

$$E \left[\mathbf{E} \exp\left(i2\pi \frac{wp}{N}\right) \sum_{k=-N/2}^{N/2-1} \mathbf{c}_k^\psi \exp\left(i2\pi \frac{kp}{M}\right) \right] = \mathbf{E} \mathbf{c}_{(-w)}^\psi \quad (84)$$

and

$$E \left[\mathbf{E}^* \exp \left(-i2\pi \frac{wp}{N} \right) \sum_{k=-N/2}^{N/2-1} \mathbf{c}_k^\psi \exp \left(i2\pi \frac{kp}{M} \right) \right] = \mathbf{E}^* \mathbf{c}_w^\psi. \quad (85)$$

Putting together the nonzero terms,

$$\overline{\mathbf{r}^B} = \mathbf{E}_0 \overline{\mathbf{r}^\psi} + \mathbf{E} \mathbf{c}_{(-w)}^\psi + \mathbf{E}^* \mathbf{c}_w^\psi. \quad (86)$$

The transform from MBC to body coordinates can be written

$$\overline{\mathbf{r}^B} + \Delta \mathbf{r}^B + \tilde{\Delta} \mathbf{r}^B = \mathbf{T}_\psi^B (\overline{\mathbf{r}^\psi} + \Delta \mathbf{r}^\psi + \tilde{\Delta} \mathbf{r}^\psi). \quad (87)$$

In light of (71) and (86), it is clear that we can separate the stochastic part from the periodic part, like

$$\Delta \mathbf{r}^B = \mathbf{T}_\psi^B \Delta \mathbf{r}^\psi \quad (88)$$

and

$$\overline{\mathbf{r}^B} + \tilde{\Delta} \mathbf{r}^B = \mathbf{T}_\psi^B (\overline{\mathbf{r}^\psi} + \tilde{\Delta} \mathbf{r}^\psi). \quad (89)$$

In other words, the stochastic signal remains stochastic: it influences neither the mean nor the periodic signal.

A Fourier expansion of (89) at time $t = p\Delta\tau$ gives

$$\begin{aligned} \mathbf{E}_0 \overline{\mathbf{r}^\psi} + \mathbf{E} \mathbf{c}_{(-w)}^\psi + \mathbf{E}^* \mathbf{c}_w^\psi + \sum_{j=-N/2}^{N/2-1} \mathbf{c}_j^B \exp \left(i2\pi \frac{jp}{N} \right) = \\ \left(\mathbf{E}_0 + \mathbf{E} \exp \left(i2\pi \frac{wp}{N} \right) + \mathbf{E}^* \exp \left(-i2\pi \frac{wp}{N} \right) \right) \left(\overline{\mathbf{r}^\psi} + \sum_{k=-N/2}^{N/2-1} \mathbf{c}_k^\psi \exp \left(i2\pi \frac{kp}{N} \right) \right). \end{aligned} \quad (90)$$

Isolating each frequency in turn,

$$\mathbf{c}_k^B = \begin{cases} \mathbf{0}, & k = 0 \\ \mathbf{E}_0 \mathbf{c}_k^\psi + \mathbf{E} \mathbf{c}_{k-w}^\psi + \mathbf{E}^* \mathbf{c}_{k+w}^\psi, & k \neq \{\pm w, 0\} \\ \mathbf{E}_0 \mathbf{c}_w^\psi + \mathbf{E} \overline{\mathbf{r}^\psi} + \mathbf{E}^* \mathbf{c}_{2w}^\psi, & k = w \\ \mathbf{E}_0 \mathbf{c}_{-w}^\psi + \mathbf{E} \mathbf{c}_{-2w}^\psi + \mathbf{E}^* \overline{\mathbf{r}^\psi}, & k = -w \end{cases} \quad (91)$$

The spectra can be obtained directly from the Fourier coefficients. Denoting as $\tilde{\mathbf{S}}^B$ the periodic part of the spectral matrix,

$$\tilde{\mathbf{S}}^B = \frac{1}{\Delta f} \mathbf{c}^B (\mathbf{c}^B)^{*T}. \quad (92)$$

As for the stochastic part, the covariance transforms as

$$\mathbf{Q}^B(p) = \mathbf{T}_\psi^B(0) \mathbf{Q}^\psi(p) \mathbf{T}_\psi^{B,T}(p). \quad (93)$$

The spectral matrix $\mathbf{S}^B(k)$ is the (discrete) Fourier transform of (93). This is

$$\mathbf{S}^B(k) = \Delta\tau \sum_{p=-N/2}^{N/2-1} \mathbf{Q}^B(p) \exp \left(-i2\pi \frac{kp}{N} \right) \quad (94)$$

or

$$\mathbf{S}^B(k) = \Delta\tau \mathbf{T}_\psi^B(0) \sum_{p=-N/2}^{N/2-1} \mathbf{Q}^\psi(p) \left[\mathbf{E}_0 + \mathbf{E} \exp \left(i2\pi \frac{wp}{N} \right) + \mathbf{E}^* \exp \left(-i2\pi \frac{wp}{N} \right) \right]^T \exp \left(-i2\pi \frac{kp}{N} \right). \quad (95)$$

This can be written as the sum of three terms,

$$\mathbf{S}^B(k) = \mathbf{S}_1(k) + \mathbf{S}_2(k) + \mathbf{S}_3(k) \quad (96)$$

with

$$\mathbf{S}_1(k) = \Delta\tau \mathbf{T}_\psi^B(0) \sum_{p=-N/2}^{N/2-1} \mathbf{Q}^\psi(p) \exp\left(-i2\pi\frac{kp}{N}\right) \mathbf{E}_0^T = \mathbf{T}_\psi^B(0) \mathbf{S}^\psi(k) \mathbf{E}_0^T, \quad (97)$$

$$\mathbf{S}_2(k) = \Delta\tau \mathbf{T}_\psi^B(0) \sum_{p=-N/2}^{N/2-1} \mathbf{Q}^\psi(p) \exp\left(-i2\pi\frac{(k-w)p}{N}\right) \mathbf{E}^T = \mathbf{T}_\psi^B(0) \mathbf{S}^\psi(k-w) \mathbf{E}^T, \quad (98)$$

and

$$\mathbf{S}_3(k) = \Delta\tau \mathbf{T}_\psi^B(0) \sum_{p=-N/2}^{N/2-1} \mathbf{Q}^\psi(p) \exp\left(-i2\pi\frac{(k+w)p}{N}\right) \mathbf{E}^{*T} = \mathbf{T}_\psi^B(0) \mathbf{S}^\psi(k+w) \mathbf{E}^{*T}. \quad (99)$$

Note how part of the energy at the frequency f in rotating coordinates comes from $f \pm \Omega$ frequencies in multi-blade coordinates.

Obtaining stress spectra at points in the driveshaft requires a similar transformation into the rotating frame. The transform from nacelle to driveshaft coordinates involves a rotation about the $z^n = z^d$ axis by the rotor azimuth angle Ψ , which is identical to Ψ_1 from above. The transform can be written in the same form as (60),

$$\mathbf{T}_n^d = \mathbf{E}_0 + \mathbf{E} \exp(i\Omega t) + \mathbf{E}^* \exp(-i\Omega t), \quad (100)$$

with

$$\mathbf{E}_0 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad \mathbf{E} = \frac{1}{2} \begin{bmatrix} 1 & -i & 0 \\ i & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \exp(i\Psi_0). \quad (101)$$

The remainder of the derivation, (62) through (99), is unchanged.

4.7 Component wear

Every component in a wind plant wears, or degrades, over time. The degradation process depends on some combination of *time*, *environment*, *loading*, and *chance*. The nature of the process is probabilistic, influenced by chance events (i.e. unmodelled processes): minute inclusions or manufacturing flaws in the material, bonding failures or in-service impacts that damage the protective coating or insulation, overload events,²⁴ and other miscellaneous occurrences.

The present observer employs a standard method for fatigue analysis, involving cycle counting and damage summation (Miner's rule). This type of fatigue analysis, intended primarily for design, is not necessarily well-suited to economic decision-making in operation. There are two shortcomings. First, the methods are based on a deterministic value of the material capacity, how many load cycles the material can take. The allowable capacity must be set at a safe level, based on overlapping conservative assumptions regarding the above effects (time, environment, loading, chance). The design value is often set at a lifetime that is one or more orders-of-magnitude below the bulk of the test data.²⁵ It is important to be conservative in the design phase, when the entire lifetime of the component lies ahead; conservatism is less warranted during operation, especially late in the component's scheduled life when the economic consequences of failure are not so severe. The second shortcoming of typical

²⁴Overload isn't simple. In ductile metals, tensile overload can *extend* the lifetime of a cracked part, by stretching and hardening the material at the crack tip. In composites, tensile overload is likely to cause further damage by breaking fibers.

²⁵The design lifetime may be lower than the mean of the data by an order-of-magnitude, but the difference in load level is smaller, due to the power-law relationship between the amplitude of alternating loads and fatigue life.

fatigue design methods is that they do not incorporate operational data – condition monitoring, whether by automated sensors or manual inspections – of the actual state of the component. When the accumulated fatigue damage of the component hits “1.0”, the component is expected to fail; never mind particular circumstances that may lead to accelerated or retarded degradation. In structural components, the alternative of crack-growth methods does provide a framework that can incorporate in-service inspection data. But inspecting for minute cracks, like the commercial aviation industry does, is expensive. More likely, other types of dynamic sensor measurements (vibrations, temperature, oil particle counts), together with superficial inspections (visual inspections of corroded, eroded, or pitted areas), will be available for estimating the condition of the turbines.²⁶

The present state observer does not offer a solution to the above critique. A cycle-counting method is used to estimate the *rate of accumulation* of stress cycles. This estimate is expected to scale properly with inputs like wind speed, turbulence intensity, the plant control signals, and so on. The problem comes not in the observer, but rather in the control law, where the objective is to use the estimate of fatigue rate to trade power generation and load reduction. Beware that the answer is not straightforward.

With these caveats aside, let us look at a set of methods that can be used to obtain a real-time estimate of the wear rates in structural and mechanical components.

4.7.1 Material stress

It is assumed that all structures behave linearly, such that there is a linear relationship between displacement and strain, also strain and stress. This is a sound assumption for normal operation.²⁷ For the design of the present state observer, we can assume that cross-sections do not grossly deform, metals do not yield, and composites do not crack or buckle.

The structure is represented by a finite beam element model. There is a particular relationship between the elastic nodal displacements and rotations $\boldsymbol{\mu}_e$ and the material strain $\boldsymbol{\epsilon}$,

$$\boldsymbol{\epsilon}(x, y, z, t) = \begin{bmatrix} \epsilon_x \\ \gamma_{xy} \\ \gamma_{xz} \end{bmatrix} = \mathbf{D}(y, z) \frac{\partial \mathbf{S}(x)}{\partial x} \boldsymbol{\mu}_e(t). \quad (102)$$

The definitions of \mathbf{S} and \mathbf{D} can be found in Merz (2018). Specifically,

$$\mathbf{D}(y, z) = \begin{bmatrix} 1 & 0 & 0 & 0 & r_z & -r_y \\ 0 & 0 & 0 & -r_z & 0 & 0 \\ 0 & 0 & 0 & r_y & 0 & 0 \end{bmatrix}, \quad (103)$$

where r_y and r_z are the in-plane coordinates of a point in the material relative to some reference line with respect to which the cross-sectional properties are defined. The matrix \mathbf{S} is 6-by-12, and for standard beam elements its nonzero entries are

$$S_{11} = 1 - \frac{x}{L}, \quad S_{17} = \frac{x}{L}, \quad (104)$$

$$S_{22} = 1 - \frac{3x^2}{L^2} + \frac{2x^3}{L^3}, \quad S_{26} = x - \frac{2x^2}{L} + \frac{x^3}{L^2}, \quad S_{28} = \frac{3x^2}{L^2} - \frac{2x^3}{L^3}, \quad S_{2,12} = -\frac{x^2}{L} + \frac{x^3}{L^2}, \quad (105)$$

$$S_{33} = 1 - \frac{3x^2}{L^2} + \frac{2x^3}{L^3}, \quad S_{35} = -x + \frac{2x^2}{L} - \frac{x^3}{L^2}, \quad S_{39} = \frac{3x^2}{L^2} - \frac{2x^3}{L^3}, \quad S_{3,11} = \frac{x^2}{L} - \frac{x^3}{L^2}, \quad (106)$$

$$S_{44} = 1 - \frac{x}{L}, \quad S_{4,10} = \frac{x}{L}, \quad (107)$$

²⁶For some components there may be data from detailed surveys, like ultrasonic inspections of blades.

²⁷There may be certain extreme load cases where nonlinearity is evident: one example would be buckling of blade skin panels under extreme gust loads.

$$S_{53} = \frac{6x}{L^2} - \frac{6x^2}{L^3}, \quad S_{55} = 1 - \frac{4x}{L} + \frac{3x^2}{L^2}, \quad S_{59} = -\frac{6x}{L^2} + \frac{6x^2}{L^3}, \quad S_{5,11} = -\frac{2x}{L} + \frac{3x^2}{L^2}, \quad (108)$$

and

$$S_{62} = -\frac{6x}{L^2} + \frac{6x^2}{L^3}, \quad S_{66} = 1 - \frac{4x}{L} + \frac{3x^2}{L^2}, \quad S_{68} = \frac{6x}{L^2} - \frac{6x^2}{L^3}, \quad S_{6,12} = -\frac{2x}{L} + \frac{3x^2}{L^2}. \quad (109)$$

If x is set to zero – that is, if the internal loads and stresses are to be computed at the “base” of a particular element – then

$$\begin{aligned} \frac{dS_{11}}{dx} &= -\frac{1}{L}, & \frac{dS_{17}}{dx} &= \frac{1}{L}, & \frac{dS_{26}}{dx} &= 1, & \frac{dS_{35}}{dx} &= -1, & \frac{dS_{44}}{dx} &= -\frac{1}{L}, & \frac{dS_{4,10}}{dx} &= \frac{1}{L}, \\ \frac{dS_{53}}{dx} &= \frac{6}{L^2}, & \frac{dS_{55}}{dx} &= -\frac{4}{L}, & \frac{dS_{59}}{dx} &= -\frac{6}{L^2}, & \frac{dS_{5,11}}{dx} &= -\frac{2}{L}, \\ \frac{dS_{62}}{dx} &= -\frac{6}{L^2}, & \frac{dS_{66}}{dx} &= -\frac{4}{L}, & \frac{dS_{68}}{dx} &= \frac{6}{L^2}, & \frac{dS_{6,12}}{dx} &= -\frac{2}{L}. \end{aligned}$$

It is convenient to linearize (102) as

$$\Delta \boldsymbol{\varepsilon} = \mathbf{D} \frac{\partial \mathbf{S}}{\partial x} \frac{\partial \boldsymbol{\mu}_e}{\partial \mathbf{q}^B} \Delta \mathbf{q}^B. \quad (110)$$

Functions for computing $\boldsymbol{\mu}_e$ and $\partial \boldsymbol{\mu}_e / \partial \mathbf{q}^B$ based on \mathbf{q}^B are provided in the STAS Aeroelastic module.

Now, stress is related to strain by Hooke’s law,

$$\boldsymbol{\sigma} = \mathbf{E} \boldsymbol{\varepsilon}. \quad (111)$$

So, if we pick a set of points around a particular cross-section at which we would like to compute the stress, each point has a particular (x, y, z) coordinate on the element, and the stress can be computed by the above equations; or alternatively, a transfer function is obtained,

$$\mathbf{H}_\sigma = \frac{\partial \boldsymbol{\sigma}}{\partial \mathbf{q}^B} = \mathbf{E} \mathbf{D} \frac{\partial \mathbf{S}}{\partial x} \frac{\partial \boldsymbol{\mu}_e}{\partial \mathbf{q}^B}. \quad (112)$$

This transfer function is constant, not a function of frequency.

Yielding and fatigue in ductile metals like steel and aluminum are most accurately predicted with the Von Mises equivalent stress criterion. In thin-walled sections a state of plane stress can be assumed, with the equivalent stress

$$\sigma_{\text{VM}} = \sqrt{\sigma_{xx}^2 + \sigma_{yy}^2 - \sigma_{xx}\sigma_{yy} + 3\tau_{xy}^2}. \quad (113)$$

This is a nonlinear function of the stress components, so the resulting probability distribution is no longer Gaussian. Pitoiset and Preumont (2000) recommend defining a Von Mises stress spectrum as

$$S_{\text{VM}} = S_{xx} + S_{yy} - S_{xy} + 3S_\tau. \quad (114)$$

This approximates the spectrum derived from the nonlinear equation (113), and remains Gaussian.

4.7.2 Material fatigue

Material fatigue is estimated using Miner’s rule, which is based on the presumption of proportional damage accumulation,

$$D_i = \frac{1}{N_i} \quad \text{or} \quad D = \sum_i \frac{1}{N_i} \quad (115)$$

where D_i is the damage done by the i^{th} stress cycle, and N_i is the number of stress cycles of this amplitude that would cause failure in a specimen of the same material. $D \geq 1$ implies failure.

Computations of fatigue damage therefore involve two elements: identification and counting of stress cycles at critical points in the material (structural *details*); and estimating the corresponding number of cycles to failure in a representative specimen.

The rainflow method is used to count cycles. Downing and Socie (1982) provide a useful algorithm that conducts this counting in real time. Given a time series of stress, this algorithm represents the “correct” result against which other methods can be compared. However, it isn’t of direct use in the observer, because we don’t have a complete time series of the stress; remember, we are filling in the higher frequencies based on an analytical, spectral model.

A linear model with Gaussian inputs, like the one we are using, has Gaussian outputs. The probability distribution of the stress at any point in the material is thus Gaussian – and the Von Mises stress is Gaussian, if the approximation (114) is employed. So we are working in the confines of linear, Gaussian theory.

Dirlik’s method²⁸ provides an estimate of rainflow fatigue cycle counts from a frequency-domain spectrum of the signal. It is a semi-empirical method, where rainflow cycle counts were generated in the time domain for a large number of archetypal wide-band Gaussian processes. Based on their observed form, Dirlik chose to represent the probability of rainflow cycle ranges using a sum of three distributions: one exponential and two Rayleigh. These are formulated as functions of moments of the frequency-domain spectra. The parameters in the formula were calibrated to best reproduce, in a least-squares sense, the time-domain cycle counts. Although it may be possible to improve marginally on the formula, especially with more complex and esoteric machine-learning algorithms, Dirlik’s method has repeatedly been found to give accurate results *provided that the spectrum represents a Gaussian process*.²⁹

Dirlik’s method defines the probability density function of rainflow cycle *amplitudes* (half of a *range*) as

$$\varphi(\sigma) = \frac{D_1}{Q\sqrt{m_0}} \exp\left(-\frac{Z}{Q}\right) + \frac{D_2 Z}{R^2\sqrt{m_0}} \exp\left(-\frac{Z^2}{2R^2}\right) + \frac{D_3 Z}{\sqrt{m_0}} \exp\left(-\frac{Z^2}{2}\right) \quad (116)$$

with the spectral moments

$$m_k = \int_{-\infty}^{\infty} |f^k S(f)| df \quad (117)$$

and other parameters

$$\begin{aligned} Z &= \frac{\sigma}{\sqrt{m_0}}, & x_m &= \frac{m_1}{m_0} \sqrt{\frac{m_2}{m_4}}, & \gamma &= \frac{m_2}{\sqrt{m_0 m_4}}, \\ D_1 &= 2 \frac{x_m - \gamma^2}{1 + \gamma^2}, & D_2 &= \frac{1 - \gamma - D_1 + D_1^2}{1 - R}, & D_3 &= 1 - D_1 - D_2, \\ Q &= 1.25 \frac{\gamma - D_3 - D_2 R}{D_1}, & \text{and } R &= \frac{\gamma - x_m - D_1^2}{1 - \gamma - D_1 + D_1^2}. \end{aligned}$$

The number of cycles in the range $\sigma_k - \Delta\sigma/2 \leq \sigma < \sigma_k + \Delta\sigma/2$, over a specified bin width $\Delta\sigma$ and period of time T , is obtained by

$$n(\sigma_k) = \varphi(\sigma_k) \Delta\sigma T \sqrt{\frac{m_4}{m_2}}. \quad (118)$$

For narrow-band processes, $D_1 \rightarrow 0$, $D_2 \rightarrow 0$, and $D_3 \rightarrow 1$, and Dirlik’s method approaches a Rayleigh distribution,

$$\varphi(\sigma) = \frac{\sigma}{m_0} \exp\left(-\frac{\sigma^2}{2m_0}\right). \quad (119)$$

Dirlik’s method is applicable when the spectrum represents a Gaussian process. This is often a reasonable assumption for a wind turbine that is operating normally, in conditions that are not

²⁸Dirlik (1985)

²⁹Ragan and Manuel (2007), Gao and Moan (2008), Mrsnik *et al.* (2013), Park *et al.* (2014), Durodola *et al.* (2018)

statistically extreme. There are a couple important exceptions though. First, immediately in the vicinity of the rated wind speed, the wind turbine transitions between the below-rated (variable-speed/maximum power-point tracking) and above-rated (constant-power/blade-pitch) control modes. A spectrum collected under such conditions will reflect the nonlinear switching between the modes, and many response variables will be non-Gaussian. Second, Dirlik’s method performs poorly when there are deterministic alternating components like gravity loads in the signal. The reason for this is simple: an individual, or dominant, sinusoid does not have a Gaussian probability distribution.

The nonlinearity associated with control-mode transitions is not a problem that can be solved, in the general sense, within the framework of Dirlik’s method. If the signal is not Gaussian, then both the theory and the calibration on which (116) is based do not apply. Constraints on time and resources prevent us from pursuing a formal solution to the problem as part of the present memo. Such a solution would be based on a large number of time-domain simulations near the rated wind speed, which could be used to calibrate an equation like (116) or, alternatively, a machine-learning model like a neural network. This is a considerable complication from a design perspective, since it implies that the calibration is conducted not for a set of generic all-encompassing spectra, like Dirlik did, but rather for a particular wind turbine and environmental condition. For the present purpose of designing the state observer, we simply assume the problem away. The observer model is scheduled on the basis of the reliable and slowly-changing cluster wind speed. According to this scheduling variable, we adopt either the below-rated or above-rated model, and its corresponding Dirlik fatigue rate. In other words, the dynamics of the transition, including the saturation of variables like the blade pitch or generator power, are ignored.

The problem of deterministic alternating stresses, like those due to gravity, wind shear, or tower shadow, cannot simply be ignored. This is because they are present whenever the turbine is operating, and they dominate important fatigue-related metrics like stresses at the blade root and in the driveshaft. An ad-hoc method was developed specifically to handle the fatigue analysis of these components.³⁰ Since the method plays a central role in the observer’s fatigue estimates, we will give it a full description and review here.


Figure 16 shows spectra of nominal stresses (not including stress concentrations or complex geometry) at the root of a wind turbine blade during normal operation in a turbulent wind.³¹ The plot at left is for the point on the downwind side of the blade root, which is primarily influenced by out-of-plane (flapwise) loads; the plot at right is for a point on the cross-wind side of the blade root – aligned with the direction of rotation – for which in-plane loads, particularly gravity, dominate. Some of the features seen in the spectra are annotated. It is of particular importance here to distinguish between the response to stochastic turbulence – the “humps” in the plots – versus periodic effects – the “spikes”.

The crux of the problem is to count fatigue cycles of a signal with both Gaussian and periodic components. The total signal in time is simply the sum of the stochastic and periodic processes; they are independent. Unfortunately the principle of superposition cannot be used directly to establish cycle counts, for reasons that will be made clear in the following discussion. Rather, we must estimate the statistics of peaks and troughs in the combined signal. The method we shall use is ad-hoc; yet it is rational, based on a consideration of the frequency content and relative strengths of the stochastic and periodic signals.

Provided that the system is linear and the noise Gaussian, the signal may be decomposed into any number of constituent parts without further assumptions or loss of accuracy. The stress spectra in rotating coordinates, with pronounced peaks and valleys, lend themselves to decomposition by frequency band. Let us then decompose the signals in this way:

³⁰Merz (2015b)

³¹To facilitate the multi-blade coordinate transform, the rotor speed was held constant when generating these spectra; under variable-speed operation the peaks, and in particular the spikes associated with periodic loads, would smear out somewhat.



- (a) Low-frequency turbulence
- (b) Stochastic signal: nP rotational sampling of turbulence
- (c) Periodic signal: wind shear, tower shadow, residual gravity
- (d) $1P$ gravity
- (e) Tower resonance $\pm P$
- (f) Residual nP periodic signal
- (g) Collective edgewise mode

Figure 16: One-sided stress spectra at the blade root, during normal operation below the rated wind speed. At left: downwind location, primarily out-of-plane loads. At right: cross-wind location, primarily in-plane loads. The heavy black curve shows the spectra, and the dashed curve is its integral, this being the variance of the signal.

1. A low-frequency stochastic signal, such as (a) in Fig. 16.
2. Individual narrow-band stochastic peaks at nP , such as (b) in Fig. 16. Only those peaks that contribute perceptibly to the variance (dashed curves in Fig. 16) need to be considered; in this example, the $1P$ peak is sufficient.
3. Individual components of the periodic signal, such as (c) in Fig. 16. Again, only those peaks that contribute perceptibly to the variance need to be considered.
4. A high-frequency residual, not bothering to distinguish between stochastic and periodic effects.

Denote these processes as respectively σ_L , σ_S , σ_P , and σ_H . Note that

$$\sigma(t) = \sigma_L(t) + \sigma_S(t) + \sigma_P(t) + \sigma_H(t). \quad (120)$$

We now count cycles for each of these processes *independently*: (1) and (4) according to Dirlik’s method, (2) according to a narrow-band approximation (Rayleigh), and (3) by the trivial consideration of a sinusoidal signal. Denote these cycle counts (low, $1P$ stochastic, $1P$ periodic, and high) as n_L , n_S , n_P , and n_H .

The cycles are combined in the following order:

1. Combine n_S and n_P , assuming that these are two sinusoidal signals with a random phase offset.
2. Combine n_H with (n_S, n_P) , based on their relative amplitudes and separation in frequency.
3. Combine (n_H, n_S, n_P) with n_L , based on their relative amplitudes and separation in frequency.

Consider σ_S and σ_P , with the corresponding n_S and n_P . These are a stochastic process and a sinusoid, in the vicinity of the same frequency band. The stochastic process is narrow-banded, which means that over a single oscillation it is approximately sinusoidal, near the $1P$ frequency. The probability distribution of the amplitude is given directly by the cycle count n_S . The relative phase of the stochastic and periodic sinusoids will drift: it is random and uniformly distributed between 0 and 2π . Say that the amplitude of the k^{th} stochastic cycle is $\hat{\sigma}_{S,k}$. The peak amplitude $\hat{\sigma}_m$ of the sum of the periodic and stochastic sinusoids, having the same frequency, is

$$\hat{\sigma}_m = |A \cos \theta + B \sin \theta| \quad (121)$$

where

$$A = \hat{\sigma}_{S,k} + \hat{\sigma}_P \cos \alpha, \quad B = -\hat{\sigma}_P \sin \alpha \quad (122)$$

and

$$\theta = \tan^{-1} \left(\frac{B}{A} \right). \quad (123)$$

For a given phase α , (121) maps the stochastic cycle count of the k^{th} stress bin into a *different* stress bin, the m^{th} . The computation is repeated for $n = 181$ phase angles, $\alpha = 0, \pi/180, 2\pi/180, \dots, \pi$, each with probability $1/n$.

Define a “cycle count density” ρ , such that the number of cycles within a given stress bin can be computed as

$$n(\sigma_k) = \rho(\sigma_k) \Delta\sigma, \quad (124)$$

where $\Delta\sigma$ is the bin width. Normalizing the cycle counts in this way ensures convergence to a particular profile, as the bin width is refined; in the limit, n goes to zero, while ρ remains finite. It is also convenient to consider the rates

$$\frac{dn}{dt} = \frac{d\rho}{dt} \Delta\sigma, \quad (125)$$

that is, the cycle counts and cycle-count densities accumulated per unit time.

Figures 17 and 18 show cycle-count density rates $\dot{\rho}$ derived from the spectra of Fig. 16. In each figure the column at left contains sample time-series of the isolated stochastic and periodic signals, together with the $\dot{\rho}$ for the stochastic signal.³² The column at right shows the equivalent quantities for the combined signal.

Starting with Fig. 17, from the plot at left it is evident that a Rayleigh distribution provides a proper estimate of $\dot{\rho}$ for the stochastic signal. In comparison with rainflow counting (and also Dirlik’s method), the narrow-band approximation neglects some small-amplitude cycles that are not expected to contribute to fatigue. Combining the stochastic and periodic signals “stretches” the Rayleigh distribution, and this effect is captured properly by the spectral method. There are some small oscillations in the spectral method’s curve associated with the discrete binning of stresses and phase angles. Otherwise, the procedure combining stochastic and periodic cycles works nicely, accurately capturing the all-important mid- to high-amplitude stress cycles, which are responsible for the majority of fatigue damage.

Figure 18 displays essentially the same features as Fig. 17, and the comments from the previous paragraph apply.

Let us now consider the combination of cycles from two random signals of different frequency bands, say “low” and “high” frequencies. These may each be wide-band processes, or they may contain a dominant periodic component. When high- and low-frequency signals are combined, some of the high-frequency cycles may disappear – depending on the relative amplitudes of the signals. Figure 19 provides an elementary example. Although the high-frequency wave makes the total signal “bumpy”, it does not alternate throughout the central portion, and therefore does not accumulate high-frequency cycles in this region.

³²The $\dot{\rho}$ for the periodic signal is a delta function – a spike – at the signal’s amplitude, and this does not need to be plotted.




Figure 17: Cycle count density rate curves from the stochastic and periodic blade root stress spectra near the 1P frequency, along with samples of the corresponding time series. Stresses at this location are dominated by flapwise bending.




Figure 18: Cycle count density rates and samples of time series, at a location dominated by edgewise bending.




Figure 19: How cycles may be lost when low and high frequency signals are combined.




Figure 20: A half rainflow cycle A-A', approximated as part of a sinusoid.




Figure 21: A cycle approximated as a three-level square wave.

A simple and conservative method is offered in order to combine the cycle counts of signals from different frequency bands. This is based on the three-level approximation described in Merz (2015c) for which a full derivation will be given here. The peculiarities of rainflow cycle counting make it difficult to arrive at more elaborate or formal method that is applicable for wide-band, potentially non-Gaussian signals. The present method can undoubtedly be improved through a more formal treatment with fewer assumptions,³³ and it is recommended to do so when time and resources permit.

Consider any peak of the low-frequency signal. The occurrence of peaks has a characteristic frequency

$$f_L = \sqrt{\frac{m_4}{m_2}}, \tag{126}$$

where the moments are given by (117). We conceptually think of the particular low-frequency peak as being a sine wave with this characteristic frequency, like the red curve in Fig. 20, representing the half-cycle A-A'. Now we represent this sine wave as a three-level square-wave signal, Fig. 21. On top of this low-frequency square wave lie some of the high-frequency peaks. How many? Well, the characteristic peak frequency f_H for the high-frequency signal can also be computed by (126). Then, over the period $T_L/3 = 1/(3f_L)$ at which the square wave is at its maximum, we can expect $f_H/3f_L$ high-frequency peaks.

³³Bishop (1988) provides a clue. He defines a list of criteria that are satisfied when a rainflow cycle is formed, and then formulates these in terms of transition probabilities, iterating the joint distributions (assumed stationary) to convergence.




Figure 22: A high-frequency process superposed with the three-level square wave, showing the amplitude σ_L of the low-frequency rainflow half-cycle.

Figure 22 shows an example, this being a case in which the high-frequency process is wide-band. A couple of the challenging aspects of rainflow counting are evident. First, the effective amplitude of the *low-frequency* half-cycle is taken from the single highest peak of the combined curve, which has its own statistical distribution, distinct from the distribution of the individual peaks. Second, if the high-frequency process is wide-band, then there will be many peaks that are irrelevant to the height of the combined signal: n_H is misleadingly high.

It is conservative, from the perspective of rainflow counting, to model a wide-band process as a narrow-band process with the same variance. In the case of the example in Fig. 22, such an approximation would be quite reasonable, since there is clearly a dominant narrow-band process containing most of the energy. Making a narrow-band assumption, we define the characteristic frequency according to the number of *zero crossings*,

$$f'_H = \sqrt{\frac{m_2}{m_0}}. \tag{127}$$

The peaks then follow a Rayleigh distribution, (119). A further simplification is possible if we allow for the fact that adjacent peaks of a narrow-band distribution tend to be correlated, and assume that the “beating” frequency is of the same order or lower than the characteristic frequency of the low-frequency signal. Then we do not need to consider the joint probabilities over a number of peaks: we can simply use the Rayleigh distribution. Note that these assumptions lead to an *ad hoc* method, one that we can show to be valid for the particular type of problem at hand, but not of general applicability.

Consider a high-frequency stress bin σ_j occurring at rate $\dot{n}_H(\sigma_j)$, and a low-frequency stress bin σ_k occurring at rate $\dot{n}_L(\sigma_k)$. These two events are assumed to be uncorrelated, which is implicit in the nature of a stochastic spectrum, though this does not hold true if the two events are both periodic signals. Here we can assume that at least one of the two events is stochastic, with a random phase. For each cycle of the low-frequency signal, one half-cycle of the high-frequency signal will augment the peak, and a second half-cycle will augment the trough. The remaining high-frequency cycles remain unchanged.

The rate at which high-frequency cycles augment the low-frequency cycles is f_L : one half-cycle at the peak, one half-cycle at the trough. The *fraction* of high-frequency cycles that augment a




Figure 23: Combined (n_S, n_P, n_H) cycles, expressed as density rate curves. Counts obtained by the rainflow method are compared with those from Dirlik’s equation and the proposed three-level method. The plot at left shows a point dominated by flapwise moments, while at the right is a point dominated by edgewise moments.

low-frequency cycle is then

$$\gamma = \frac{f_L}{f_H}. \quad (128)$$

The fraction $\gamma \dot{n}_H(\sigma_j)$ is subtracted from the original bin σ_j and mapped to a new bin, $\sigma_p = \sigma_j + \sigma_k$, where σ_k is distributed according to $\dot{n}_L(\sigma_k)$. This leaves $(1 - \gamma) \dot{n}_H(\sigma_j)$ in the original bin.

Figure 23 shows results derived from the spectra of Fig. 16, when the approach is used to combine the high-frequency cycles n_H with $1P$ cycles (n_S, n_P) . Considering first the flapwise stresses (left-hand plot), the nature of the rainflow and spectral curves differ in the low-amplitude range. Neither of the spectral methods correctly capture the way in which the periodic component, which here is relatively mild (Fig. 17), dissipates some of the small-amplitude cycles.³⁴ This is to be expected, since the Dirlik method has no consideration of periodic effects, while the three-level method conservatively neglects this dissipation. At high stress amplitudes, the rainflow count follows a Rayleigh distribution in the tail, which agrees precisely with that of Dirlik’s method. The three-level method overpredicts the occurrence of extreme cycles by approximately a factor of two. This is also to be expected, since the three-level square wave extends the time spent at the maximum and minimum values of the low-frequency signal. Figure 24 provides a different perspective of the high-amplitude tail, here on a linear scale.

The three-level method is not needed in cases without a strong periodic component; in this case, Dirlik’s method can be used directly. The purpose of the three-level method is to provide reasonable predictions in cases such as the one in the right-hand plot of Fig. 23, which has a strong $1P$ alternating stress due to gravity. Here Dirlik’s method fails: one must explicitly incorporate the periodic signal into the counting algorithm. The three-level approach correctly, albeit slightly conservatively, predicts the way in which high-frequency stress cycles augment the gravity cycles. The residual small-amplitude cycles are overpredicted, by design.

The strength of the three-level method is that it provides a reasonable and slightly conservative estimate when the periodic component is small, and also a reasonable and slightly conservative estimate when the periodic component is large. Although applying the three-level method is not necessary in the former case, the method will provide the correct trends as periodic components become more pronounced, for instance at points on the blade root in between the two extremes of Fig. 16.

The final step in the cycle-counting analysis is to combine (n_H, n_S, n_P) with n_L ; that is, to incorporate the low-frequency stochastic signal into that of Fig. 23. To do this we apply the three-phase

³⁴Refer to Fig. 19.




Figure 24: A close-up of the tail of the left-hand plot in Fig. 23, with a linear scale.




Figure 25: Combined (n_L, n_S, n_P, n_H) cycles, expressed as density rate curves, for the point dominated by flapwise moments.

assumption yet again. The one catch is that a frequency of $1P$ is used for f_H in (128), as this represents the rate associated with the major cycles of the combined (n_H, n_S, n_P) signal; that is, those cycles that would be likely to augment each peak or trough of the low-frequency signal.

Figure 25 shows the final results of the cycle counting, for the point dominated by flapwise moments. The errors associated with the Dirlik and 3-level methods are of comparable magnitude, but the 3-level method has the advantage of being on the conservative side. The two methods agree on the high-amplitude cycles, which have an important influence on the fatigue life. There is evidently an advantage to an approach – whether simplified like the 3-level method, or something more elaborate – that treats the periodic part of the signal explicitly.

The final cycle counts for the case dominated by edgewise moments are indistinguishable from those in the right-hand plot of Fig. 23.

For non-rotating components like the foundation, where stochastic signals dominate over periodic, a straightforward application of Dirlik’s method gives a good estimate of the fatigue cycle counts.

4.7.3 Actuator wear

Constraints on time and scope prevent us from developing a proper model of actuator wear, based on the physics associated with the potential modes of degradation and failure. As an expedient for preliminary studies, a surrogate for actuator wear is taken to be the total distance travelled,

$$\mathcal{D} = \int_0^t \left| \frac{d\alpha}{d\tau} \right| d\tau, \quad (129)$$

where α is the actuator angle: β in the case of the pitch actuators, or χ for yaw. The rate of accumulation of wear is then

$$\dot{\mathcal{D}} = \left| \frac{d\alpha}{dt} \right|. \quad (130)$$

The mean speed of the pitch and yaw actuators is zero, so $\dot{\alpha}$ is characterized by its variance, which is also equal to the variance of $|\dot{\alpha}|$. A convenient measure of the actuator wear rate is therefore the standard deviation of $\dot{\alpha}$, which can be derived from the spectrum as

$$\sigma_{\dot{\mathcal{D}}} = \int_{-\infty}^{\infty} S_{\dot{\alpha}}(f) df = (2\pi)^2 \int_{-\infty}^{\infty} f^2 S_{\alpha}(f) df. \quad (131)$$

It is recommended to revisit this in the future, and implement a method for estimating the wear rate based on a physical model of the damage mechanisms in pitch and yaw actuation systems.

4.8 Tuning and performance

The TotalControl Reference Wind Power Plant³⁵ is adopted as a case to demonstrate the operation of the state observer. The turbines in this plant are an offshore version of the DTU 10 MW Reference Wind Turbine; a description of the version used here, including the support structure and electrical components, can be found in Anaya-Lara *et al.*³⁶

The model used in this section is linearized about an operating point of a 10 m/s wind speed, and either maximum power-point tracking (nominal operation), or an operator power command of 6 MW.

4.8.1 Setup for testing the observer

The observer can be tested with a setup like that sketched in Fig. 26. A physical model of the wind turbine is linked with the observer model. The inputs to the wind turbine are simplified: a power command, uniform wind speed, wind direction, and waterline ocean wave force. These simplified inputs make it straightforward to check the observer’s estimates against the “true” values.

The wind turbine model is closed-loop, including the turbine’s controller. It outputs sensor measurements and control commands, as well as a set of additional outputs: the rotor-average induction factor a , rotor thrust F_T , available aerodynamic power P_a , and a set of structural displacements \mathbf{q} at the blade root and foundation. When it comes to the sensor measurements, one simplification has been made with respect to Fig. 6: the voltage and current measurements have been replaced by a single value of electric (active) power, and the reactive power is not considered at present.

The sensor measurements from the turbine are fed to the state observer, which outputs the same quantities. If the observer is functioning as it should, the estimated outputs will match the same values from the turbine model.

³⁵Andersen *et al.* (2018), Merz *et al.* (2019). The wind turbine specification is taken from Bak *et al.* (2013).

³⁶Anaya-Lara O, *et al.* (2018). *Offshore Wind Energy Technology*. Wiley.




Figure 26: The architecture used to test the observer.

4.8.2 Full and reduced-order observer models

The observer is based on an open-loop wind turbine model with 296 states. It is feasible to solve for the Kalman filter gains and implement an observer based on the full model. It is not necessary to retain all 296 states, though, since many of these are associated with a high-frequency response that has no impact on the wind plant controller. Better to eliminate some of the unnecessary states. This is done with the following procedure:

1. Perform a modal transformation of the open-loop matrices. Retain those modes that contribute to the input-to-sensor transfer functions, following the procedure described in Merz (2020).
2. Perform a Y-transformation³⁷ to make the resulting matrices real.
3. Reduce the order of the resulting model using sequential perturbation approximation,³⁸ preserving the accuracy of the transfer functions *up to a given cutoff frequency* to within a given tolerance. Here the cutoff frequency was set to 0.5 Hz.
4. Perform a second modal transformation followed by a Y-transformation, this time of the reduced matrices. Make the modes with eigenvalues above the cutoff frequency quasi-steady, by partitioning the (modal) matrices and solving for the steady-state values of the high-frequency modes. Specifically,

$$\frac{d\mathbf{x}_1}{dt} = (\mathbf{A}_{11} - \mathbf{A}_{12}\mathbf{A}_{22}^{-1}\mathbf{A}_{21})\mathbf{x}_1 + (\mathbf{B}_1 - \mathbf{A}_{12}\mathbf{A}_{22}^{-1}\mathbf{B}_2)\mathbf{u}, \quad (132)$$

$$\mathbf{y} = (\mathbf{C}_1 - \mathbf{C}_2\mathbf{A}_{22}^{-1}\mathbf{A}_{21})\mathbf{x}_1 + (\mathbf{D} - \mathbf{C}_2\mathbf{A}_{22}^{-1}\mathbf{B}_2)\mathbf{u}, \quad (133)$$

with “1” indicating the retained modes, and “2” the quasi-steady modes.

Note that the reduction procedure does not alter the inputs \mathbf{u} or outputs \mathbf{y} ; the interface variables are unchanged. The final version of the turbine observer, used to generate the results in this report, had 26 states.

³⁷Merz (2020), Stevens and Lewis (2003)

³⁸The Octave function `spamodred` was used; this is based on Sreeram *et al.* (1995) and other references cited therein.

Table III: Observer gains along the diagonals; other entries are zero. The variable corresponding to each gain is noted.

Q		R		K_I	
V	1.0	Ω	5×10^{-5}	Ω	0.01
θ_V	0.1	$v_{n,x}$	0.001	$v_{n,x}$	0.10
F_w	1.0	$v_{n,y}$	0.001	$v_{n,y}$	0.10
		V	1.0	V	0.10
		θ_V	1×10^{-4}	θ_V	0.10

4.8.3 Tuning and dynamic response

The weights \mathbf{Q} and \mathbf{R} in (35), as well as the gains \mathbf{K}_I in (37), were tuned by trial-and-error. Good performance was obtained with the gains listed in Table III: good tracking while not being overly sensitive, and limited overshoot. Of particular note is the large noise placed on the anemometer wind speed measurement V , in the \mathbf{R} matrix. This is in effect telling the observer that we do not trust the anemometer, and the majority of weight should be given to the estimate of wind speed obtained through the combination of the rotor speed, blade pitch, and electric power sensors. Also of note is the *small* noise placed on the anemometer wind direction measurement. This was necessary, because the misalignment of wind is only marginally observable through the standard sensor measurements. That is to say, even though the anemometer measurement is noisy, it is nonetheless more reliable than the estimate obtained from the other sensors. When the anemometer measurement was discounted (via a higher value in the \mathbf{R} matrix), it was found that errant wind-direction changes were predicted, which degraded some of the other estimates. In view of the comparatively high gain on the anemometer measurement, it is necessary that this is pre-filtered, say, with a low-pass filter, and calibrated, before feeding it to the observer. The wind direction measurement can also be improved by a clustering methodology, taking advantage of measurements collected at multiple wind turbines in the vicinity.³⁹

To illustrate the performance, Figs. 27 through 30 plot the response of the turbine and observer outputs to step changes in the power command \hat{P} and the wind speed V . The oscillations seen in many of the signals are associated with the first resonant modes of the tower, and the active damping of these using blade pitch and electrical power. The key result is that the observer provides an accurate estimate of the key non-sensor quantities used in control, in particular the wind speed and rotor thrust. The integral action on the residual sensor errors is evident as gradual corrections in signals like the wind speed.

5 Examples of controller tuning and performance

The controller from Section 3, with the observer from Section 4, were linked with a model of the TotalControl Reference Wind Power Plant, described in Merz *et al.* (2019). Due to the importance of nonlinearity to the controller’s function – particularly the saturation of the power-tracking integral – a special simulation setup was required. Each wind turbine model is linearized about an operating point, which for the examples of this section is taken to be a wind speed of 10 m/s and power command of 6 MW. Working with linearized models, both for the turbines and the observers, enables them to be reduced, and diagonalized, using the algorithm of Section 4.8.2. The matrix multiplications required to advance the state equations in time are fast for such sparse matrices, which made simulation of the entire plant feasible in Octave.⁴⁰

Tuning of the plant controller was conducted on simplified models; the process and final tuning are described in Section 5.1. Section 5.2 looks at the clustering of wind turbines, estimation of the

³⁹Annoni *et al.* (2019), *op.cit.*

⁴⁰Octave is an interpreted code, which is slow in comparison with compiled codes like C or Julia.




Figure 27: The response of sensor outputs to a unit step change in the electric power command. Dark lines: turbine model. Light lines: observed quantities. Ω : rotor speed (rad/s), β : blade pitch (rad), χ : nacelle yaw angle (rad), P_e : electric power (MW), $v_{n,x}$: nacelle fore-aft velocity (m/s), $v_{n,y}$: nacelle side-to-side velocity (m/s), V : wind speed (m/s), θ_V : wind direction (rad).




Figure 28: The response of auxiliary outputs to a unit step change in the electric power command. a : axial aerodynamic induction factor, F_T : rotor thrust (MN), P_a : available power (MW), $i_{s,q}^\theta$: q-axis current at the transformer terminals (kA), $\theta_{F,x}$: elastic side-to-side rotation of the foundation mudline node (rad), $\theta_{F,y}$: elastic fore-aft rotation of the foundation mudline node (rad), $\theta_{b,y}$: elastic flapwise rotation of the blade root node (rad), $\theta_{b,z}$: elastic edgewise rotation of the blade root node (rad).




Figure 29: The response of sensor outputs to a unit step change in the wind speed.




Figure 30: The response of auxiliary outputs to a unit step change in the wind speed.

Table IV: Default tuning of the controller, for the analyses in this section.

Variable	Value	Units	Description
K_P	1.0	(-)	Proportional gain on power tracking
K_I	0.4	s^{-1}	Integral gain on power tracking
K_T	6.0	m/s	Gain on thrust compensation
ω_T	0.03	Hz	Thrust low-pass filter corner frequency
w_η	0.10	(-)	Normalized width of saturation band on K_I
ω_c	0.05	Hz	Low-pass filter corner frequency for inputs
ω_n	0.24	Hz	Tower notch filter frequency for inputs
ζ_1	0.04	(-)	Lower damping ratio on notch filter
ζ_2	0.40	(-)	Higher damping ratio on notch filter

cluster-average wind speed, and turbulence cascade. Section 5.3 illustrates how theoretical spectra are obtained from the turbulence cascade and used to generate an estimate of component wear rate. Section 5.4 wraps up with a simulation of the Total Control Reference Wind Power Plant, including aspects of operation as a function of wear rate.

The simulations in Sections 5.2 through 5.4 make use of synthetic wind field data, averaged over the rotor area, with realistic coherence across the wind plant. This data was generated according to the empirical spectra of Viguera-Rodríguez, *et al.*⁴¹ These spectra are representative of offshore sites in the North Sea, where there is more energy at very low frequencies – that is, the weather is more variable – than in the standard wind spectra used to design wind turbines (Burton *et al.* 2001).

5.1 Tuning, frequency response, and step response

Table IV lists the tunable scalar parameters in the controller, along with their values after tuning for the DTU 10 MW wind turbine. The parameters w_η , ω_n , ζ_1 , and ζ_2 were straightforward to determine, and require little elaboration. The other parameters have a crucial influence on the controller behavior and system dynamics, and interact with each other – that is, they cannot be set independently. All the gains were tuned by trial-and-error using a closed-loop dynamic model consisting of the wind turbine, plant controller, and a “grid” that fed the power back with a gain of N_t , the number of turbines in the plant. This then represents the worst-case scenario where a disturbance causes all the turbines to respond in unison; under normal operation, the feedback mechanism through the grid is more diffuse. Tuning was first done using a linear transfer-function model, and then nonlinear simulations were run with elementary inputs (steps, ramps) to see the performance in the time domain, and test features like saturation.

The low-pass filter frequency ω_c is applied to all the inputs, and sets the upper bound on the bandwidth of the controller response. The value of 0.05 Hz is a tradeoff between the speed of the power-tracking response and the desire to reject spurious signals, in particular the degree to which high-frequency dynamics adjacent to the tower notch frequency are excited by the feed-through and proportional-gain pathways. The gains K_P and K_I were tuned according to the classic criteria, a rapid rise and slight overshoot to a step response, implying a damping ratio on the control-response mode of around 0.5. The gain K_T and low-pass frequency ω_T were set by scanning the (K_T, ω_T) space, and finding the highest K_T that reduced the thrust over the frequency band $0 \leq f < \omega_T$, without introducing undesirable resonance in the power-tracking response. Note that the gains K_P and K_I are such that, when $\lambda N_t \approx 1$, they dominate K_T .

Figure 31 illustrates the performance of the controller in terms of transfer functions from windspeed and power command, to electric power, rotor thrust, rotor speed, and blade pitch. The high-frequency part of the transfer functions, above 0.1 Hz, is essentially unaffected by the choice of the plant control

⁴¹Sørensen *et al.* (2002, 2008), Viguera-Rodríguez *et al.* (2010, 2011)




Figure 31: Transfer functions between input wind speed and power commands, and output electric power P_e (MW), rotor thrust F_T (MN), rotor speed Ω (rad/s), and blade pitch β (rad). The upper two plots show the response under the default gains: full power tracking and thrust rejection. The lower left plot shows the influence of reducing the power-tracking gains to zero. For reference, the lower-right plot shows the response to changes in the wind speed, under a constant power command.

tuning; this is due to the low-pass and notch filters on the plant control inputs. The interesting portion of the curves, for the present purposes, is that below 0.1 Hz.

The two plots at top show the response to fluctuations in wind speed and plant power command, where the latter has been normalized such that each turbine receives a power command of unit amplitude. Starting with the power-tracking response (top-right plot), there is a power-control mode with a frequency of 0.09 Hz and damping ratio of around 0.4. The timescale of the response is on the order of 10 seconds. The electrical equipment would support a faster, more aggressive response, but this would come at the cost of exciting tower motions, and would be detrimental to the drivetrain; then it would be necessary to continually evaluate the tradeoffs between power tracking and structural fatigue. It is simpler and more appropriate for a baseline controller to avoid such unwanted interactions through filtering.

As for the response to a perturbation in wind speed, this is seen in the top-left plot. Despite the fact that the power-tracking controller is acting against the thrust compensation, there is nonetheless a reduction in thrust over a narrow frequency band around 0.03 to 0.08 Hz; this can be seen by comparing against the lower-right plot, which shows the response when the power command is held constant. The thrust compensation obviously comes at the cost of increased power fluctuations; and the integral effect of the power-tracking controller eventually overwhelms the thrust compensation, such that there is a negligible reduction in thrust below 0.01 Hz. Given the characteristic timescales of atmospheric




Figure 32: The response of electrical power (at top) and rotor thrust (at bottom) to unit step inputs in wind speed (at left) and power command (at right), with different gains.

turbulence, where much of the energy is below 0.01 Hz, it is expected that the default tuning behaves essentially as a power-tracking controller, without much thrust mitigation. The situation changes if the power-tracking gains are reduced. The plot at lower left shows the extreme case where power tracking is disabled. In this case the thrust compensation is effective at low frequencies.

The spike in electrical power at the tower resonant frequency (0.24 Hz) is associated with the active damping of tower side-to-side vibration, which is part of the turbine’s internal controller, not influenced by the plant controller.

Figure 32 provides an alternate look at the turbine response, in the time domain. A step in either wind speed or power command was input at $t = 30$ s. The plots show the response of the electrical power (upper plots) and rotor thrust (lower plots). Different gains give different responses. The gray curve (barely visible behind the red in the left-hand plots) is the case where there is no change in the power command sent to the turbine. The red curve is the case where only power-command tracking is active, without any thrust compensation. The blue curve, on the other hand, is the case with thrust compensation and no power-command tracking. Finally, the black curve is the case where all the gains are set to their nominal values.

Looking at the right-hand plots, it is evident that the introduction of thrust compensation degrades the power-tracking performance, although not severely. Both tunings are equally responsive, in terms of rise-time, but the damping of the power-tracking mode is somewhat reduced in the presence of thrust compensation. This would appear to result in extra thrust cycles, which could be detrimental; but keep in mind that this is a step function, which excites high-frequency dynamics, whereas a more usual input would be a gradual ramp. The timescale of ramps in the power command should evidently be 10 seconds or more, to avoid causing any additional oscillations in thrust.




Figure 33: An example of the clustering algorithm for separating “turbulence” from “weather”, applied to simulated data with proper plant-wide correlations. The black line is the cluster wind speed, drawn from the individual turbine wind speed estimates, shown in gray. The turbine wind speed series, for which the cluster wind speed is determined, is highlighted in red.

The plots at left show the response to a step change in wind speed. Such a sudden change in the wind speed excites all sorts of aerodynamic and structural modes in the wind turbine, quite apart from any effect of the plant controller. In particular the first tower side-to-side frequency is visible in the plot of electrical power, as this is being used for active damping. In terms of the plant-control response, the thrust compensation (blue curve) wants to increase the power set-point, in order to keep the thrust near its original value.⁴² However, when power tracking is active (black curve), it pulls the power back to its original level, and the thrust drifts accordingly.

Overall, it can be said that the closed-loop dynamics of the nominal tuning are within performance envelopes usually considered acceptable for PI controllers. Operation is stable under the combination of power tracking and thrust compensation, or either extreme: pure power-tracking or pure thrust compensation. By design, with the nominal tuning, the power-tracking controller dominates the thrust compensation; this is adjusted by scheduling down the power-tracking gains.

5.2 Clustering, weather, and turbulence cascade

Figure 33 shows the clustering algorithm of Section 4.5 at work. The black line is the cluster wind speed for the turbine whose local wind speed is shown by the red line. Recall that the cluster wind speed is not simply the average, but rather the maximum likelihood computed by a simplified Bayesian method. This makes it robust to sensor errors; Fig. 34 is an example. The figure also shows the response to a step input in the wind speed; with the current tuning, the characteristic response time of μ_V is about 1 minute.

The black line in Fig. 33 is the “weather”; the difference between the black and red lines is the “turbulence”, Fig. 35. The goal is to identify the spectral parameters of Section 4.6 (used further to derive the instantaneous damage rate), and for this the recent history of turbulent fluctuations is stored, called here the “turbulence cascade”. At each timestep⁴³ the value of the turbulence is added to the top of the turbulence cascade, and the last value removed, advancing one step in time. The turbulence cascade is therefore a window of a specified length in time on the curve in Fig. 35, the

⁴²It is perhaps counterintuitive that an increase in the wind speed should result, without compensation, in a *decrease* in the thrust. This is due to the blade pitch control; indeed, the static thrust-versus-wind-speed curve of a wind turbine has a negative slope in the constant-power, pitch-controlled regime.

⁴³The turbulence cascade may be updated with a timestep that is longer than that of the simulation timestep. In the present case, the update timestep was 1 s, compared with a simulation timestep of 0.05 s.




Figure 34: The cluster wind speed estimate with a step input of -5 m/s at $t = 10$ minutes. Two sensors are faulty and are returning zero; the cluster wind speed (black curve, see Fig. 33 for the legend) is to be compared with plain averaging of the signals (blue curve).




Figure 35: The time series of turbulence derived from Fig. 33, showing the window (red box) containing the turbulence cascade at the present time (dashed line).

front edge of which is located at the present time. An example of the window is shown as a red box, with the present time being the dashed line.

The window should be long enough that a reasonable approximation to the m_0 and m_2 spectral moments can be derived. On the other hand, the window cannot be too long, because the response to transient conditions will be too slow. For the present analyses a timestep of 1 s and a window of $2^9 = 512$ measurements (ca. 8.5 minutes) was chosen, based on the character of the wind speed fluctuations in Fig. 35. In a real application, a shorter window could be justified, in order to detect cases where a wind turbine finds itself in partial wake overlap, with a meandering wake from an upstream turbine causing transient fluctuations in the effective wind speed. Wake effects are not modelled in the present time series of turbulence, and we leave this as an open question for later investigation.




Figure 36: Single-point turbulence intensity and length scale for each turbine, derived from the turbulence cascade after a 10-minute burn-in period. The error bars indicate the standard deviation.

5.3 Spectra for estimating component wear rate

The parameters describing the *single-point* turbulence spectrum at each turbine, according to (48) through (51), are shown in Fig. 36. These are reasonable single-point values, given the nature of the input wind data. The length scale is larger than the typical design values used with a Von Karman spectrum, however this is consistent with the spectra from which the wind time series were generated.

A numerical example of the chain of operations from the theoretical wind spectrum to fatigue or component wear rates was given already in Section 4,⁴⁴ and need not be repeated here. There is the question of arriving at a single scalar “damage rate” \dot{D}^* , for input to the controller, from the wear rates of various components. We shall side-step the issue here, and leave the details to future work. Broadly speaking, the problem is similar to design optimization, where a cost or objective function is defined based on a collection of individual metrics. This is often taken as a weighted sum-of-squares; the extremum among the individual metrics; or as a more elaborate attempt (often based on dubious assumptions) to formulate the actual financial cost.

5.4 Controlling the TotalControl Reference Wind Power Plant

It remains to demonstrate how the turbines work in concert to balance the plant power output, while rejecting low-frequency thrust loads. According to the controller of Section 3, the degree to which a given turbine contributes to power balancing is set as a function of the damage rate \dot{D}^* . Here we shall consider a generic, normalized damage rate that takes a value between 0 (low damage) and 1 (high damage). Let the relationship between \dot{D}^* and α_i be

$$\alpha_1 = 1 - 0.1\dot{D}^* \quad \text{and} \quad \alpha_2 = 1 - \dot{D}^*. \tag{134}$$

That is, as the damage rate increases from 0 to 1, the nominal share of the turbine’s power, for determining the target thrust, drops from 1 to 0.9; and the power-tracking gain factor drops from 1 to 0, such that turbines with the highest damage metric do not contribute to balancing plant power.

Four scenarios are considered. The first is without any plant control; simply a uniform power set-point of 6 MW is sent to each wind turbine. The next scenario is when the plant controller is operating, and all damage rates are zero. In the third scenario, Turbine 1 is given a damage rate of 1, with the rest being zero. Finally, in the fourth scenario the turbines are initialized with random damage rates between 0 and 1, which are then held constant throughout the simulation period.

⁴⁴More details on the turbulence model can be found in Merz (2020).




Figure 37: The power at the PCC, where the set-point is 192 MW, or 6 MW per turbine on average. Gray line (flat at 192 MW): constant set-point. Red: zero damage. Blue: Turbine 1 unit damage. Black: random damage.

The power tracking control is excellent. Figure 37 shows the power at the PCC during the four scenarios. The generator power controls act quickly, and in the absence of any commands the power is held steady at 6 MW.⁴⁵ When the power-tracking and thrust mitigation control strategy is introduced, the power fluctuations increase, though they do not exceed 0.5%. There is essentially no difference in the overall power-tracking performance, for different distributions of damage rate among the turbines in the plant. This is to be expected, since a decrease in the power-tracking gains at one turbine is counterbalanced by a similar increase at another turbine.

Figure 38 provides another view of the way in which the turbines help to balance each other. These plots are rainflow-count exceedance curves: guidance on how to read the curves is found in the figure caption; they provide a nice way to view the statistical aspects of the fluctuations in a time series, and have a connection with the severity of component wear. The difference between the gray curve and the other curves illustrates how implementation of a thrust-reducing control strategy reduces the severity of fluctuations in the thrust, while introducing fluctuations in the electric power. That much is obvious; what is interesting is how the black and red curves match closely. The black curve shows the case where the damage rate \dot{D}^* at Turbine 1 is set to 0, and for the red curve \dot{D}^* is 1; this is the only difference between the two scenarios. Why does the level of damage at Turbine 1 not affect its response; that is, why is it not forced to contribute more to power tracking, when its damage is 0? The answer is that the other turbines in the plant are taking up the slack. Turbine 1 with $\dot{D}^* = 1$ contributes nothing to power tracking; yet there are still 31 other turbines who are contributing, and the loss of one turbine makes little difference. To state things another way, the thrust-mitigation function, which for a single turbine appeared to be dominated by the power-tracking function (Figs. 31 and 32), is in fact fully effective when the turbine is operating as part of a plant.

Now the question is: How does the controller work when different turbines have different damage rates? In Scenario 4, the turbines have been assigned a random value of \dot{D}^* , between 0 and 1. The results are nicely summarized in Fig. 39, showing the mean and standard deviation of each turbine's

⁴⁵Not all sources of high-frequency noise are included in the electrical model, so the degree of uniformity may be excessive.




Figure 38: Rainflow-count exceedance curves of thrust and power for Turbine 1. Gray: constant power command (Scenario 1). Black: damage = 0 (Scenario 2). Red: Turbine 1 damage = 1 (Scenario 3). The gray curve is near zero in the right-hand plot. For an example of how to interpret these curves, the gray curve in the left-hand plot indicates that there are about 100 total cycles; 40 cycles that exceed an amplitude of 0.02; 10 cycles that exceed an amplitude of 0.11; and so on.

rotor thrust and electric power, plotted according to its damage rate. There is a significant synergy effect by coordinating the operation of wind turbines in a plant: the severity in rotor thrust fluctuations can be reduced for over half the turbines, with only a few turbines – those that can best tolerate it – experiencing an increase in the severity of loading.⁴⁶

These results, obtained for rotor thrust, could likely be repeated for other types of “load” components, like driveshaft torque, blade root moments, or blade pitch actuator usage. An attempt to simultaneously mitigate several load components, using only the electrical power command, would require tradeoffs. This would best be handled by a more advanced control algorithm capable of acting according to a cost function, specifying how the different load components should be traded. In designing such a controller, linear-quadratic (LQR) synthesis about a steady-state operating point would be the simplest place to start.

6 Power systems design and analysis with PSCAD

It is important to understand how the behaviour of the control strategies might affect the dynamics and operation of the electric grid. Power System Computer Aided Design (PSCAD) is a simulation software program built specifically for this purpose, and can provide a detailed analysis of system transients and other events affecting the response of offshore wind power plant electrical infrastructure.

This capability is used to provide a non-linear representation of the offshore electrical infrastructure of the TotalControl Reference Wind Plant (TC-RWP). The result is a high-fidelity electrical model of the TC-RWP, which captures the details of high frequency dynamics of non-linear electrical components under varying wind farm control conditions.

A PSCAD model has been developed to interface with the hierarchical wind plant controller, together with the STAS aeroelastic wind turbine model. This section presents the architecture of the PSCAD model, the interface to the wind turbines and wind plant controller, and some preliminary numerical results.

⁴⁶These are the low-frequency fluctuations in thrust, based on a rotor-average wind speed input, akin to the frequency band (a) in Fig. 16. There are additional sources of loading acting at higher frequencies, which are not significantly affected, positively or negatively, by the plant-level controller.




Figure 39: Trends in the mean and standard deviation of thrust and power fluctuations, as a function of damage rate (Scenario 4). Nominal results for the case with a constant power command (Scenario 1) are shown in the background.

6.1 Electrical System

The electrical system on the PSCAD simulation program is modelled following the detailed description and parameters for the Reference Wind Power Plant presented in previous TotalControl deliverables.⁴⁷ They provide information on the electrical representation of the layout for the collection grid with cable sections parameters and of the high voltage AC export system.

6.1.1 Collection Grid

The TC-RWP is a virtual test bed for wind power plant control designed with symmetry to facilitate high-resolution numerical flow simulations. The layout consists of 32 turbines in a staggered pattern with a distance between rows and columns of 5 times the rotor diameter (Andersen 2018). Hence considering the spatial locations, the length of the cable sections between individual adjacent turbines is either $5D$ or $5.6D$ as shown in Figure 40.

The electrical layout of the TC-RWP, has been conducted according to the EERA DTOC inter-array design procedure (Endegnanew *et al.* 2013). The inter-array grid voltage is 66 kV, which is foreseen as the standard for the current generation of offshore wind power plants, with turbines approaching, and soon exceeding, 10 MW rating. It consists of two strings of 7 turbines, and three strings of 6 turbines with branching reducing the overall length of the thickest (500 mm^2) cable.

The cable model is a representation of XPLE 3-Core aluminium submarine cable. The thick and thin lines represent different cable types whose properties are summarised in Table II in Merz *et al.* (2019). The chosen cross-sections 95 mm^2 and 630 mm^2 are used based on the current rating.

Within the PSCAD standard simulation package, there are no 3-core XPLE cable models provided. However, the software offers enough flexibility to build such models, if there is enough detailed and accurate information available of the required cable parameters. The PSCAD simulation program provides two approaches to model the cable sections either through a) the Cable library or b) the Transmission line libraries.

- (a) Creating models with the Cable library: The choice of this approach depends on the complexity of the studied system. For example, the Cable library provides a model with detailed description of cable parameters such as cable formation (i.e. flat or trefoil), distance between conductors, laying depth conductor diameter, soil temperature, cable propagations etc. It is suitable to analyse

⁴⁷Andersen *et al.* (2018), Merz *et al.* (2019)




Figure 40: Collector grid cable section layout of the TC-RWP.

Parameters	Values
Resistance	0.246 mΩ/m
Inductive reactance	0.133 mΩ/m
Capacitive reactance	18.72 MΩ/m

Table V: Collection grid Pi-section cable parameters. From Table II in Merz et al. (2019).

complex cable designs. The response of the model to extremely high frequency events can be very good, provided that enough ‘real world’ information is available to fine tune the model.

(b) Creating models with the Transmission line library: This approach captures mainly the frequency response characteristics of the cable and suits analysis which focus around the resistance and reactance values of cables (i.e. R, Xc, XI). Three types of transmission line models are provided depending on the complexity of the analyses at hand (Jeevantha 2016):

- Pi Section Model: This is appropriate for system studies that are based on a single frequency but does not capture travelling waves within cables. For example, grid system studies at 50 Hz.
- Bergon Model: This is also appropriate for single frequencies but also captures travelling waves.
- Frequency dependent model: This captures a wider range of frequencies DC up to GHz and also reflections.

In this system study, most of the analysis undertaken builds on wind farm control and how it interacts with the active and reactive power requirements (PQ) of the network. Hence, the Pi section equivalent model was considered to have an appropriate response, and it was used to design the 0.99 km and 1.1 km cable sections. Table V summarises the details of Pi-sections based on the RLC (Resistance, Inductance, Capacitance) cable parameters.

Parameters	Values
Transformer rating	180 MVA
Impedance	3% (0.01 pu)
No-load losses	0.5% (0.005 pu)
Copper losses	3% (0.003 pu)
Magnetizing current (no load)	1%
Air core reactance	20% (0.2pu)

Table VI: Transformer data used for the model.

6.1.2 Export system

A schematic of the export system is given by Merz *et al.* (2019), and can also be seen in Fig. 48 below. Table II in Merz *et al.* (2019) provides properties of the submarine collection grid cables, as well as offshore and onshore sections of export cables. The PSCAD simulation program presents a high-fidelity representation of the export system with detailed model of a 180 MVA transformer and high voltage alternating current (HVAC) export cables. In the transformer model, the core non-linearity is represented with non-linear characteristics approximated based on the ‘knee point’, the ‘air core reactance’ and the magnetizing current at the rated voltage. The used parameters are listed in Table VI.

Similar to the approach taken in section 6.1.1, the transmission line representation was used to model the underground and submarine cables each 30 km in length. Here, the Bergon model was chosen as this is a sufficient representation for most power system studies.

For the shunt reactor, this is placed by the 400 kV/220 kV onshore transformer substation, and it incorporates enough inductance to compensate for both cables’ capacitance and associated reactive power. The shunt reactor is modelled based on the description of the X/R ratio of 500 provided. The system is modelled to be grid connected, with the point of common coupling (PCC) considered to be at the onshore substation.

6.2 Wind turbine model

The model for the wind turbine in the PSCAD simulation program is developed with an open framework to make it flexible to changes further down the line. This will facilitate analyses and investigations planned in other related activities of the project, tasks 4.1.3 and 4.1.4. The model design is split into three main systems of the wind turbine, i.e. the ‘aero-mechanical’ power train, electrical power train and switch gear arrangements. Figure 41 shows a graphical representation of this set up.

A description for an approach for modelling the wind turbine is presented in the following subsections. A simplified representation of the DTU 10 MW reference wind turbine is used, modelled initially to investigate the behaviour of the power network design with regards to the control interactions between the electrical power train and the baseline wind farm control dispatch actions.

6.2.1 Switch Gear

A simplified arrangement for switch gear design for connecting adjacent wind turbines is modelled as shown in Fig. 42. This is a typical design of those found in offshore wind farm electrical systems and represents the connections between the wind turbine transformer and switch gear. Three modes of operation are considered; when the wind turbine is connected, disconnected and isolated under fault conditions.

The PSCAD simulation program also provides high fidelity representations of switch gear which cover the electrical transient characteristics of the breakers. This will allow to explore how breakers react when switching transients occur that may affect the dynamic behaviour of closed loop controllers.




Figure 41: High level description of the Wind Turbine Model in PSCAD.




Figure 42: Wind turbine string switch gear arrangement representation in PSCAD.

This could be useful for further work to be carried out in the future activities of the project, such as those planned in task 4.2 aimed to better the understanding of electrical, mechanical and control interactions of wind power plants.

6.2.2 Electrical power train

The description of the electrical power train model is provided below, which covers the design of the grid-side converter and active and reactive power control loops. The features of the power train are designed to achieve control of the wind turbine at its AC terminals. In order to achieve this, the control systems of the converter must provide, at least, the following functions:

- Active power control to maintain DC voltage level on the converter capacitor.
- Reactive power control either through voltage regulation or power factor at the wind turbine AC terminals.
- Converter AC current control with the aid of PI controller for active and reactive currents.




Figure 43: Simplified representation of electrical power train.

- A grid voltage synchronisation technique for synchronising the converter voltages to the collection grid voltage (i.e. voltages at the terminals).

Regarding the first function, i.e. active power control to maintain DC voltage level, it is worth noting that modelling its behaviour accurately is key to represent properly how the wind turbine’s converter will respond to high frequency events like grid faults. The scope of work within this deliverable though, is to understand how the wind farm control and wind turbine control would behave under closed loop conditions when scaled up to a full scale wind farm with 32 wind turbines, such as those in TC-RWP. Therefore, in order to simplify the model and speed up the simulation, the converter characteristics are assumed to be ideal and the converter is modelled as controlled voltage source with power set-point capability as shown Figure 43.

Regarding the other three control functions, i.e. reactive power, current control and grid synchronisation, Figure 44 presents the block diagram of the whole control system incorporating those three control functions. On the left-hand side, it gives all measured voltage and current signals, including voltages and current at wind turbine terminals V_{WT} . Then the AC signals are transformed to their corresponding d-q representations taking the wind turbine terminals as the reference vector. Grid voltage synchronisation is achieved using a vector-based Phase Locked Loop (PLL) scheme. The d-axis and q-axis values for current control are measured values from the wind turbine AC systems. They should be evaluated according to whether they are for reactive power control, for voltage regulation or power factor correction.

In this study, PI based control is used as shown in Figure 45. The current controller adjusts the AC terminal voltage by regulating the reactive current flow between the converter and wind turbine AC connection point. Using two feedback PI controllers with decoupling loops, the measured converter current d and q components i_{dWT} , i_{qWT} follow their respective references i_{dref}^* , i_{qref}^* , set by an outer power-based controller. The output control signals generate the d-q voltage elements V_{cd} , V_{cq} , to give the required voltage magnitude and phase angle.

All the individual 32 turbines in the TC-RWP are modelled with this capability, comprising a total of 128 PI controllers tuned to control the individual actions of the active and reactive current paths for each turbine.

The outer control loop is a power based controller built to handle active and reactive power set-points as shown in Figure 45. These are modelled also with PI controller to provide reactive and active power set-point control for each turbine. The reactive power control is modelled to receive power set-point commands but can be substituted if required with voltage regulation loop to provide voltage support to maintain the magnitude of the wind turbine terminal voltage.

The proposed structure for the wind turbine ‘aero-mechanical’ model is shown in Figure 46. This




Figure 44: Block diagram of control scheme.




Figure 45: Control loops of Grid Side Converter.



Figure 46: Block diagram of model for aero-mechanical power train.

structure has been chosen because it allows each block to be modified individually to accommodate the requirements of different models when collaborating with other modelling activities happening in parallel within WP4.

Three possible ways to approach the detailed modelling of this block diagram in the PSCAD simulation program have been thought of. These are described below as options 1, 2 and 3.

Option 1: Time series data with look-up tables at each stage.

This option takes into consideration effects of the different modelling approaches in other work packages. Activities of other work packages can interact at different levels following the structure of the wind turbine model described.

a) Wind model. WP4 partners could interact at the wind model by providing a time series of effective wind speeds for each of the 32 wind turbines. The varying advanced control modelling approaches i.e. wakes, surrogate, predictive, etc., would need a simple model that applies these effects transforming this to a single point wind speed. An effective wind speed can then be passed to the PSCAD model as a time series data (resolution of data can be decided depending on the case studies).

b) Rotor aerodynamics model. The structure of this part of the model will depend on how it is required to respond to requests from the wind farm controller. A possible approach is splitting this into two lookup tables. The first lookup table uses the effective wind speed to determine the current operating state (rotor speed and blade collective pitch angle) of the wind turbine. The second lookup table combines that with the tip-speed ratio (TSR) to determine the C_p power coefficient of the wind turbine. If any other areas such as dynamics are to be captured in this model, the effect of this would need to be reflected within these two lookup tables.

c) Drivetrain model. A simple solution can be used for modelling the wind turbine drivetrain by multiplying the speed and divide the torque at the low-speed shaft by the gearbox ratio, giving speed and torque at the high-speed shaft. Mechanical losses may also be considered in a simplified way. The disadvantage of this approach is that the dynamics of the control system are not simulated. However, this reduces the need for additional PI control in the drivetrain. As the primary purpose of the baseline controller is to examine the electrical system wind farm closed loop controller, this might be a suitable approach to test this. If there were reasons to try to provide interaction at this stage, considerations would be to make on how to reflect the required changes in the drivetrain.

Option 2: Time series data of wind speeds and generator lookup table

This option will work with wind speed time series data that represent the operation of the wind farm under the different control techniques to be investigated. The model could be simple solution with lookup tables of the DTU 10 MW turbine from results of our SCADA and Bladed Models information that can work on PSCAD simulation environment. In such a scenario, the performance of the model would be validated with information provided to act like a digital twin. This could be with time series data of wind speeds and generator characteristics of each turbine for which representative lookup tables can be developed, as illustrated in Figure 47.

Option 3: Transfer function representation of wind turbine aerodynamics

This model allows the wind turbine aerodynamics of the power train to be modelled if they can be represented as transfer functions. The PSCAD model can represent time domain characteristics of




Figure 47: An example of option 2 with lookup tables of wind turbine speed and power curves.

transfer functions in s-domain. This would need to be expressed as either first or second order transfer function or a series of transfer functions. This option would be appropriate and suitable for models where transfer functions already exist such as the STAS wind turbine model to use for the future case studies in Section 7.

6.3 Analysis of WPP model

The current state-of-the-art wind power plant (WPP) has generally two control levels. The wind turbine controller which addresses each wind turbine separately ensuring a safe and optimised operation of the individual wind turbine under given wind conditions and the WPP controller which addresses the whole system level tasks including grid code compliance at the connection point (PCC) of the WPP to the electrical network. The main functionalities of wind power plant controllers are:

- To receive set-points for the whole WPP from the grid operator regarding:
 - Curtailment (reducing active power P to a specified value in specified time)
 - Reactive power Q, power factor or voltage control set-points
- To calculate how to distribute the required demands of the grid operator within the plant:
 - Distributing the overall plant active power reduction requirement to individual wind turbines
 - The same for reactive power set-points to the individual wind turbines

Although WPP control exists, control functions are progressing towards more coordinated approaches where the primary objectives include increasing wind farm power production, reducing turbine loads and providing ancillary grid services in an optimised way.

The activities within this work package look to develop a PSCAD model of the TC-RWP with a hierarchical wind power plant supervisory control which allows for these more advanced control strategies to be applied. This is envisaged to act as baseline to investigate other activities within this work package which investigate novel control strategies such as model-predictive plant controllers, surrogate models, and dynamic wake induction control.

The TC-RWP PSCAD model is developed with these considerations in mind. The baseline is set to be the state-of-art described above with the capability to receive set-points for the whole WPP from the grid operator and calculate how to distribute the required demands of the grid operator. In this case, there is no specific wind farm control theory applied at the PCC and it is left open and flexible to investigate different control strategies such as the STAS based power plant controller to be




Figure 48: An illustration of representation of the hierarchical baseline PSCAD TC-RWP model.

used in the case studies outlined in Section 7. With this baseline framework, control strategies such as PI closed loop control and model predictive supervisory control can be applied within the PSCAD model. An illustration of the representation on the PSCAD model is as shown in Figure 48.

This consists of the electrical system with the collection grid and export system (Section 6.1) and the closed-loop based wind turbine controller (Section 6.2). An open-loop based WPP with equal power distribution dispatch across all turbines is used here to validate the operation of the system. This model at later stages will implement a version of the Observer and Plant Control law described in Sections 3 and 4 as part of the activities of the case studies in Section 7.

Figure 49 shows results for power production across all turbines and the response of the wind turbine controller to a sudden step response change in power.

For power production, a representation of the power curve of a wind turbine is used to represent its electrical power output at different wind speeds. Figure 50 shows a generic example of a turbine power curve. This power production results imitates the turbine at the cut-in wind speed. The turbine has zero power output when the wind speed is less than the cut-in speed. Power output increases as the wind speed increases above the cut-in speed and it achieves the rated power output at the rated wind speed.

The power response of the wind turbine controllers is also tuned to represent the fast performance of the wind turbine reacting to power changes within a second after set-point commands are received. In this case, at worst case scenario of a step change in power demand from no power to rated power, the wind turbines achieve a steady state response within 500 ms.




Figure 49: Simulation Results: (a) Power production with the wind turbines, (b) wind turbine controller electrical power response.




Figure 50: Power curve of a wind turbine.

7 Conclusions and Outlook

A supervisory control algorithm for large wind power plants has been designed. It is intended as a baseline for comparison against more advanced algorithms. The controller tracks a total power command specified by the plant operator, while compensating for fluctuations in rotor thrust at individual wind turbines. By taking advantage of the fact that each wind turbine is its own independently-controlled generating unit, the plant controller is able to reduce thrust fluctuations on highly-stressed wind turbines, while tightly following the commanded power, provided that this is below the total power available in the given wind conditions.

The control architecture is simple, based on PI-type algorithms with a filter cascade. However, this simplicity is attained because the controller acts on physically-meaningful inputs, which are not available directly from the raw sensor data. A state observer is therefore necessary. A candidate design for a state observer has been developed, following but one out of a variety of feasible approaches to the problem. The observer architecture is based on a physical model of the wind turbine. A noteworthy feature is its ability to provide online predictions of the fatigue damage or component wear rate, using cycle-counts derived from a spectral model of the dynamics. This includes a probabilistic clustering algorithm to separate local turbulence from larger-scale transients in the atmospheric flow through the plant, and thereby derive a set of analytical turbulence spectra that represent the conditions seen by each wind turbine. The observer is well-suited to the design and tuning of plant control algorithms, and could be useful in an operational setting; however once the system architecture is fixed and a database of operational data is available,⁴⁸ data-driven algorithms would likely outperform the proposed one.

The controller, including the observer, was tested on a model of the TotalControl Reference Wind Power Plant. This demonstrated that the controller satisfies the performance specifications, and indeed has the potential to provide a significant reduction in the fluctuating thrust loads experienced by operating wind turbines. On balance, the load reductions on highly-stressed turbines far outweigh the load increases on low-stressed turbines: there is a definite synergy effect in coordinating the operation of wind turbines across a large wind power plant.

7.1 Outlook

The controller of Sections 2 through 5 will be used in the next phase of TotalControl to study the system dynamics of grid-connected wind power plants. This section outlines plans for the work to be conducted in this area, which will culminate in Deliverable 4.6, on electrical/mechanical/control interactions in large wind power plants.

In summary, the work in the next phase is expected to be structured in three steps:

1. Integrate STAS wind turbine representation into TC-RWP model in PSCAD.
2. Explore electro-mechanical-control interactions in grid faults.
3. Explore the impact of ‘bad designs’ on electro-mechanical-control interactions.

The following three subsections elaborate on each of these separately

7.2 Integrated model for case studies

The case studies to be investigated require two modelling tasks:

Model A: Updated electrical module of the STAS turbine model The current implementation of transformer and converter in the electrical module of the STAS Turbine is simplified (Merz 2019). These simple models are sufficient to analyse the comparatively low-frequency dynamics that are relevant for turbine control and active power control at the plant level. However,

⁴⁸ ...either actual operational data, or from high-fidelity simulations...




Figure 51: Block diagram of integrated WT/WPP control/PSCAD model.

the analysis of electrical interactions between turbines and the grid requires a more detailed representation of the electrical turbine components. A first step will therefore be to refine the transformer and converter models in the STAS Turbine.

Model B: Integrated PSCAD and STAS model This would build further on the baseline hierarchical model described in Section 6 with an addition of a simplified representation of the STAS Plant Controller and STAS wind turbine model. This would be integrating STAS Turbines (state-of-the-art representation of aeroelastic, electrical and control dynamics of a wind turbine) into the PSCAD grid model.

Figure 51 shows the foreseen integrated model. The currents and voltages at the turbine terminals will be the interface variables between the STAS Turbines and the PSCAD grid model. The plant controller uses SCADA information from the turbine models together with an external active power to generate power set-points for the turbine controllers.

The integrated model will provide relevant aeroelastic behaviour of wind turbines, a WPP control function that dispatches the active power demand, and a detailed model of the electrical grid system in the WPP. This constitutes a representation of the dynamics in a WPP which is beyond common practice. The system dynamics will be analysed in case studies. Both models STAS and PSCAD could be somewhat validated and are close to sufficient representations and undertake analysis for different scenarios. The performance of both models is to be compared, i.e. 1) the linear STAS model, and 2) the non-linear PSCAD model with linear representations of wind plant control and mechanical power train in STAS.




Figure 52: Annotated frequency trace of the partial black-out event in August 2019 (National Grid ESO 2019).

7.3 Case studies on grid transients and faults

Different case studies will be simulated in PSCAD to analyse the advanced wind farm control response to grid disturbances and faults. Both STAS wind farm controller and drivetrain STAS models will be integrated within a representation of the TC-RWP in an Electromagnetic Transients Program (EMTP). This will enable the team to explore the interactions of the electrical and mechanical components of the WPP with control functions that are of particular interest.

PSCAD is built for this sort of analysis as it captures non-linear high frequency oscillation events and provides a close to real-life representation. This makes PSCAD better suited to build complex controllers than similar packages for power systems in Matlab. The integrated model (Model B) from Section 7.2 will be used for this activity. In most of the power system studies carried out with EMTP tools in the time domain, the wind turbines are aggregated into a single unit represented as a controllable voltage source. Some studies incorporate an aggregated full back to back power converter connected wind turbine. In this case, it is our intention to represent every single wind turbine independently to an appropriate level of accuracy, i.e. complex enough to replicate the wind turbine dynamics that influence electrical grid dynamics, and light enough so it can run reasonably quick in a workstation.

As an example of the type of case study that would be interesting to replicate and to better understand, Figure 52 shows the annotated frequency trace of the partial black-out event that occurred in Eastern England in August 2019. The disconnection of a big offshore wind farm was a key element in the chain of events, which likely originated from the wind turbine/control system command to disconnect the turbines from the grid. The overcurrent capacity of the converter switches is also a key element in this situation, which is in turn directly related to the turbine level control loop.

The exact characteristics of the case studies will need to be better enunciated, but it is clear that developing combined electro-mechanical models where both the electrical system and the turbine mechanical dynamics are represented accurately enough for its behaviour being representative in high frequency events is of interest for the industry.

7.4 Case studies on WPP control and grid interactions

Another set of case studies will focus on the interaction of the WPP(s) with the regional grid.

Two main approaches will be followed:

1. Generate a deeper understanding of the system by analysing unfavourable designs and interactions.
2. Investigate the interactions of generating units connected to the same grid.

The first approach builds on deliberately constructing “bad WPP(s)” and searching for scenarios with unfavourable interactions between the WT/WPP controls and low-frequent grid dynamics. The analysis of critical setup and behaviour can support design decisions that are more favourable.

The second approach targets the integration of WPPs into the wider power system. For that, the integrated WPP model with baseline plant controller will be connected to a model of a regional grid with additional power generators. The influence of neighbouring generation units (wind, solar, conventional, power system stabilizer (PSS) functions) and their control can be studied. This includes interarea oscillations. Focus will be on the impact between WPP control and the other units.

While wind power plants have been traditionally operated with the aim to maximize power production, this operational strategy will be challenged in the future with a higher share of renewable energy in the grid. Wind power plants are expected to participate to a larger extent in the electricity markets and to be exposed to the increasingly volatile electricity prices. This provides incentives to curtail power generation and instead alleviate structural loads or provide ancillary services. Changing the operational strategy affects the requested control functionality. New control functions are required that trade off the conflicting objectives (loads, power, grid support). Suite of WT/WPP control functions can be studied in the integrated model.

Furthermore, the impact of transmission system design and grid topologies can be studied in different scenarios. This is not only relevant for WPPs but also for the integration of other variable renewable energy into the power system.

Acknowledgements

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

References

- Anaya-Lara *et al.* (2018). *Offshore Wind Energy Technology*. Wiley.
- Andersen S *et al.* (2018). *Reference Wind Power Plant*. TotalControl Deliverable D1.3.
- Annoni *et al.* (2019). Wind direction estimation using SCADA data with consensus-based optimization. *Wind Energy Science* 4: 355-368.
- Bak C, *et al.* 2013. *Description of the DTU 10 MW Reference Wind Turbine*. DTU Wind Energy Report-I-0092, Technical University of Denmark.
- Bishop NWM (1988). *The Use of Frequency Domain Parameters to Predict Structural Fatigue*. PhD Thesis, Department of Engineering, University of Warwick, Coventry, England.
- Burton T, *et al.* (2001). *Wind Energy Handbook*. Wiley.
- Connell JR (1982). The spectrum of wind speed fluctuations encountered by a rotating blade of a wind energy conversion system. *Solar Energy* 29: 363-375.

- Dirlik T (1985). *Application of Computers in Fatigue Analysis*. PhD Thesis, Department of Engineering, University of Warwick, Coventry, England.
- Downing SD, Socie DF (1982). Simple rainflow counting algorithms. *International Journal of Fatigue* 4: 31-40.
- Durodola JF, *et al.* (2018). Artificial neural network for random fatigue loading analysis including the effect of mean stress. *International Journal of Fatigue* 111: 321-332.
- Endegnanew AG, *et al.* (2013). *Design procedure for inter-array electric design. Report D2.2, EERA DTOC*.
- Fossen TI (1994). *Guidance and Control of Ocean Vehicles*. Wiley.
- Gao Z, Moan T (2008). Frequency-domain fatigue analysis of wide-band stationary Gaussian processes using a trimodal spectral formulation. *International Journal of Fatigue* 30: 1944-1955.
- Hewer GA (1971). An Iterative Technique for the Computation of the Steady State Gains for the Discrete Optimal Regulator. *IEEE Transactions on Automatic Control* 16: 382-384.
- Hsu CS (2010). *Cell-to-Cell Mapping: A Method of Global Analysis for Nonlinear Systems*. Springer-Verlag.
- Jeewantha S (2016), *et al.* *Modelling Cables and Transmission Lines with PSCAD/EMTDC (Webinar, November 2016)*.
- Kleinman DL (1968). On an Iterative Technique for Riccati Equation Computations. *IEEE Transactions on Automatic Control* 13: 114-115.
- Kristensen L, Frandsen S (1982). Model for power spectra of the blade of a wind turbine measured from the moving frame of reference. *Journal of Wind Engineering and Industrial Aerodynamics* 10: 249-262.
- Kundur P (1994). *Power System Stability and Control*. McGraw-Hill.
- Madjidian D (2016). Scalable minimum fatigue control of dispatchable wind farms. *Wind Energy* 19: 1933-1944.
- Merz KO (2015a). Environmental loads for frequency-domain aeroelastic analysis of offshore wind turbines. Memo AN 15.12.19, SINTEF Energy Research.
- Merz KO (2015b). Design verification of the drivetrain, support structure, and controller for a direct-drive, offshore version of the DTU 10 MW Reference Wind Turbine. Memo AN 15.12.68, SINTEF Energy Research.
- Merz KO (2015c). Rapid optimization of stall-regulated wind turbine blades using a frequency-domain method: Part 1, loads analysis. *Wind Energy* 18: 1703-1723.
- Merz KO (2018). *STAS Aeroelastic 1.0 – Theory Manual*. Report 2018:00834, SINTEF Energy Research.
- Merz KO (2019). *STAS Electric 1.0 – Theory Manual*. Memo AN 19.12.07, SINTEF Energy Research.
- Merz KO (2020). *Development of an LQR framework for rapid prototyping of offshore wind turbine controllers, with application to active load control*. Report 2020:00257, SINTEF Energy Research / TotalControl Deliverable D3.3.

- Merz KO, Pedersen MD (2018). Offshore wind turbine controls. Chapter 5 of Anaya-Lara *et al.* (2018).
- Merz KO, *et al.* (2019). *An Electromechanical Model of the TotalControl Reference Wind Power Plant*. Report 2019:00342, SINTEF Energy Research / TotalControl Deliverable D1.5.
- Mrsnik M, *et al.* (2013). Frequency-domain methods for a vibration-fatigue-life estimation – Application to real data. *International Journal of Fatigue* 47: 8-17.
- National Grid ESO (2019). *Technical Report on the events of 9 August 2019*. ESO final technical report.
<https://www.nationalgrideso.com/document/152346/download> (Accessed October 2020).
- Park J-B, *et al.* (2014). A new fatigue prediction model for marine structures subject to wide band stress processes. *Ocean Engineering* 76: 144-151.
- Pitoiset X, Preumont A (2000). Spectral methods for multiaxial random fatigue analysis of metallic structures. *International Journal of Fatigue* 22: 541-550.
- Ragan P, Manuel L (2007). Comparing estimates of wind turbine fatigue loads using time-domain and spectral methods. *Wind Engineering* 31: 83-99.
- Smilden E (2019). *Structural control of offshore wind turbines – Increasing the role of control design in offshore wind farm development*. PhD Thesis, Department of Marine Technology, Norwegian University of Science and Technology.
- Sreeram *et al.* (1995). New results on frequency weighted balanced reduction technique. Proceedings of the American Control Conference, Seattle, WA, USA, June 1995.
- Stengel RF (1994). *Optimal Control and Estimation*. Dover.
- Stevens BL, Lewis FL (2003). *Aircraft Control and Simulation*. Second Edition, Wiley.
- Sørensen P, *et al.* (2002). Wind models for simulation of power fluctuations from wind farms. *Journal of Wind Engineering and Industrial Aerodynamics* 90: 1381-1402.
- Sørensen P, *et al.* (2008). Modelling of power fluctuations from large offshore wind farms. *Wind Energy* 11: 29-43.
- Vigueras-Rodriguez A, *et al.* (2010). Wind model for low frequency power fluctuations in offshore wind farms. *Wind Energy* 13: 471-482.
- Vigueras-Rodriguez A, *et al.* (2011). Spectral coherence model for power fluctuations in a wind farm. *Journal of Wind Engineering and Industrial Aerodynamics* 102: 14-21.
- Zhou K, Doyle JC, Glover K (1996). *Robust and Optimal Control*. Prentice Hall.