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

Title: Validation of optimized control schemes

Deliverable no.: 2.7

Delivery date: (17.08.2022)
Lead beneficiary: DTU
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
Ervin Bossanyi	DNV	ervin.bossanyi@dnv.com
Gunner Chr. Larsen	DTU	gula@dtu.dk
Renzo Ruisi	DNV	
Mads M. Pedersen	DTU	mmpe@dtu

Acknowledgements/Contributions:		
Name	Name	Name
Anders Sommer, VF	Peder Enevoldsen. SGRE	

Document information

Version	Date	Description			
			Prepared by	Reviewed by	Approved by
1	12.7.2022		Gunner Chr. Larsen	Ervin Bossanyi	

TABLE OF CONTENTS

1. EXECUTIVE SUMMARY

In this deliverable we present a description of the Lillgrund offshore WPP, including site as well as WT description. This WPP serves as the demonstration case for an intended full-scale validation of optimized open-loop WPP control schedules conditioned on ambient wind speed and wind direction. The developed WPP controller includes an entire row of Lillgrund WTs (i.e. Row A), and Row A power output is the objective function used in the optimization.

The WPP schedules are developed using the DNV and DTU numerical models LongSim and Fuga, respectively. The developed optimal controllers are tested by comparing LongSim results, based on Fuga optimized input, with Fuga optimized results and vice versa. These models are very different in their nature — Fuga being a linearized RANS CFD solver, and LongSim being based on a range of engineering wake models and allowing for unsteady WPP flow modelling — but nevertheless the overall computed Row A power gains are reasonably similar, which gives confidence to the developed WPP schedules, even though the two models predict significantly different distributions of the power gains over wind speeds and directions.

Based on the numerical analysis, a simplified WPP controller for implementation in the Row A WTs is developed and tested against the more detailed optimized controller. The test predicts only a small reduction of performance of the simplified controller compared to the full controller, and on this basis it is decided to use the simplified WPP controller for the full-scale WT implementation. The test, toggling between base control and optimal control in 50 minute successive time slots, has been running since 09-05-2022 and is still ongoing, because a huge amount of data is needed to obtain robust statistics for the comparison of measured and numerical data. Initial results of the validation are finally shown.

2. INTRODUCTION

Lillgrund offshore wind power plant (WPP) has been extensively used in the TotalControl [1] project, both as a reference wind farm for example calculations and also for campaigns of field measurements. In 2020, a field measurement campaign used three long-range scanning LiDARs, one to measure the wind inflow and two to characterize the wakes behind several wind turbines (WTs) as their control settings were changed – both yaw misalignment and axial induction settings [11].

In the final stages of the project, an extensive axial induction control experiment has been implemented on the wind farm for field validation of optimal open loop WPP control schedules derived from numerical simulations using the platforms LongSim [9], [10] and Fuga [5] ,[6].

This report describes the design of this WPP controller, the testing of the controller using realistic time-domain simulations, and the design and implementation of the field test experiment on the actual WPP. This is also described in [12], and further details are available in [2]. Finally, some preliminary results of the full-scale field testing of the WPP control scheme are presented, covering a period of three months, from the start of the testing on 09-05-2022 until 09-08-2022. The tests are still ongoing as more data is needed to obtain statistically significant results.

Beyond TotalControl, the plan is to extend these full-scale tests to optimal control of the entire WPP of 48 WTs.

A wake steering controller was also designed in [2], which could give greater power gains, but the effect on fatigue loads is not so benign, and practical issues with site hardware would anyway have precluded any plans for implementation within the time and budget of the TotalControl project.

3. THE LILLGRUND OFFSHORE WIND POWER PLANT

Lillgrund WPP is located near Copenhagen in the sea between Denmark and Sweden, and consists of 48 Siemens Gamesa 2.3MW WTs with rotor diameter of about 93m (SWT-2.3-93).

3.1 SITE AND WPP LAYOUT

The WTs are arranged in eight regularly spaced rows with some gaps, as shown in Figure 1. The gaps were not planned but resulted from challenging foundation conditions. One characteristic of the Lillgrund WPP, which makes it especially suited for WPP flow and control studies, is the unusually small inter-turbine spacings — as low as 3.3 diameters in one direction.

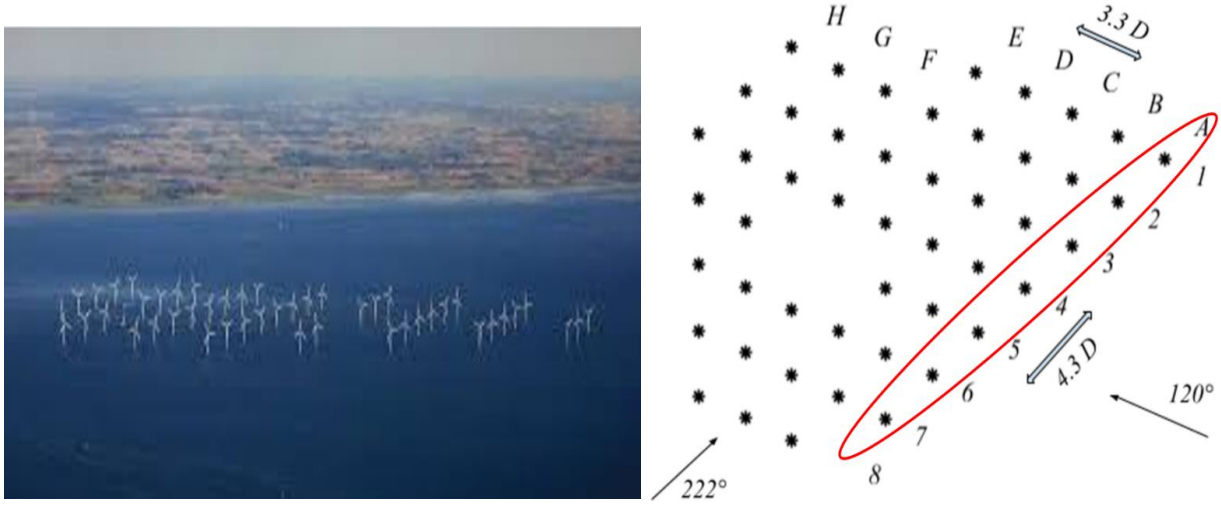

Figure 1: Lillgrund wind farm layout showing WT numbering as well the row of test WTs.

The prevailing wind is south-westerly, with fairly low turbulence intensity (although higher than on some offshore sites due to the proximity of land), and atmospheric stability split fairly evenly between stable and unstable stratification, though the ratio varies with the season and time of day.

3.2 WT MODEL

A *Bladed* aeroelastic model of the SWT-2.3-93 wind WT was constructed using information provided directly by the turbine manufacturer SGRE, and indirectly by DTU. Although some assumptions had to be made where the information supplied was uncertain or insufficient, the model gave a reasonable match to information which was provided on WT performance, such as natural frequencies etc. Further details are not publicly available.

The WT controller follows a fairly conventional operating curve: below rated wind speed, the blade pitch is constant, and the WT power is set as a function of rotor speed using a look-up table, which allows power to be maximised by tracking the optimum tip speed ratio as far as possible, up to the maximum rotor speed of 16 rpm. Above rated wind speed, the pitch angle increases to maintain the rated power of 2.3 MW at 16 rpm rotor speed.

4. WPP CONTROLLER DESIGN

The axial induction controller design has two parts: firstly using power and thrust coefficient tables calculated by the aeroelastic code Bladed to define different WT controller set-points designed to reduce rotor thrust, and hence wake losses, and secondly, using an optimization loop within the wind farm code LongSim to find the distribution of these set-points across the turbines in the farm which maximizes the total farm power.

4.1 WPP SETTINGS FOR AXIAL INDUCTION CONTROL

Since wake losses are essentially caused by rotor thrust, WT set-points are sought which reduce the rotor thrust by modifying the tip speed ratio and pitch angle, so that the power coefficient C_P remains as high as possible for a given level of thrust coefficient (C_T) reduction. In this way the controller design space is reduced from two to one design variable per WT [8] and hence considerably reducing the computational burden associated with the controller optimization.

The conventional set-points used by the turbine controller are the fine pitch angle β_0 and the 'optimal mode gain' K as defined in [3] - this defines the generator torque Q as a function of the measured generator speed ω which achieves the desired values of C_P and tip speed ratio λ

$$Q = K\omega^2 - L, \text{ where } K = \frac{\frac{\pi}{2}\rho R^5 C_P}{\lambda^3 G^3}$$

where ρ is the air density, R the rotor radius, G the gearbox ratio and L represents any mechanical drive train losses referred to the high speed shaft.

For any desired C_T reduction, an iteration loop calculates the optimal λ and β_0 which maximize C_P, and the corresponding value of K is then calculated as above. This is easily converted to a power-rpm lookup table, as required for this type of WT, with power demand calculated as $Q\omega\varepsilon$, where ε is the electrical efficiency. Set-points were calculated for thrust reductions from zero to 16% compared to the 'greedy' strategy (which maximizes power without regard to thrust, as for an isolated turbine). Figure 2 shows the resulting power and thrust coefficient curves, and the power-rpm lookup tables and pitch angles needed to achieve this are shown in Figure 3. 'Mode A' represents the nominal setting assumed to be used on site, although this information has subsequently been found to be outdated: the actual nominal setting is currently much closer to the 'Optimised' setting in Figure 2.

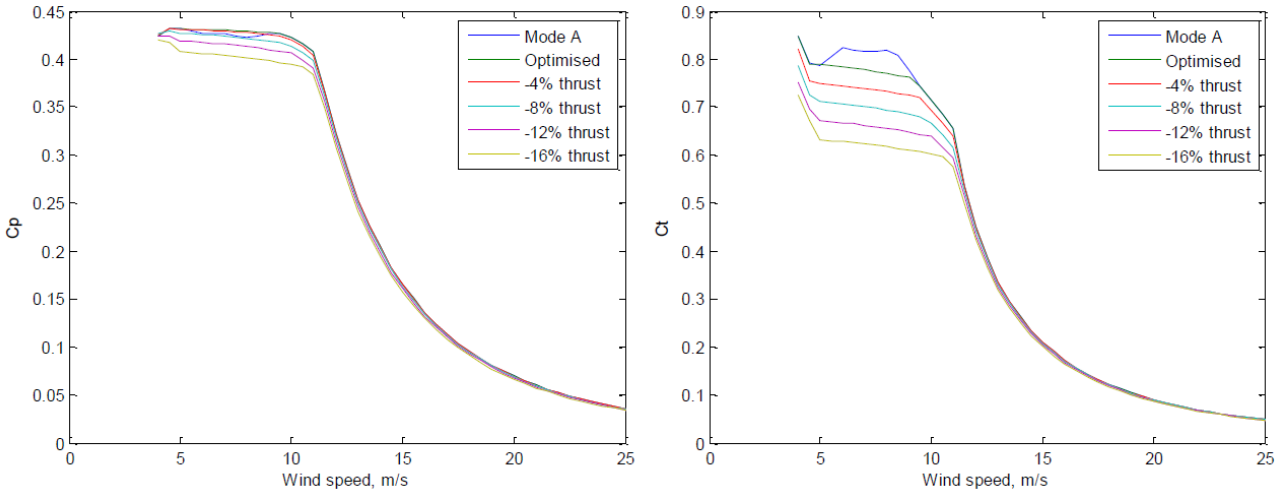

Figure 2: Effect of the optimal thrust reduction set-points on power coefficient (left) and thrust coefficient (right).

ΔC_T	K	β_0
Mode A	-	-1°
0%	0.4221	-1°
-4%	0.4638	-0.54°
-8%	0.4835	0.09°
-12%	0.5093	0.66°
-16%	0.5596	1.12°

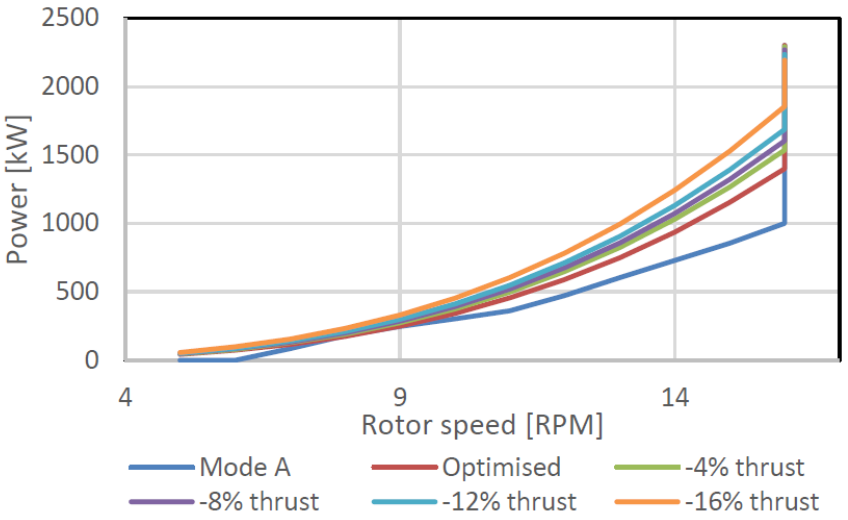

Figure 3: Controller settings for the optimal thrust reduction set-points.

4.2 WPP SET-POINTS FOR AXIAL INDUCTION CONTROL

To design the axial induction controller for the WPP, the LongSim platform is used to perform a steady-state optimization in which the thrust reduction set-points of each turbine are adjusted until a defined merit function is maximized. A number of control designs for Lillgrund WPP are documented in [2], including both wake steering and axial induction control, optimized either for power maximization only, or for a weighted combination of power and some key fatigue loads.

For practical reasons, the field tests had to be limited to axial induction control only, using a merit function for power maximization only. Axial induction control, even if optimized only for power maximization, already reduces loads significantly, as shown in [2], so power maximization is an appropriate goal. Moreover, although several turbines are equipped with load sensors, full evaluation of a strategy designed for power and loads would need load sensors on all the turbines.

Axial induction control is most effective when the wind is blowing directly along the rows of turbines, in below-rated wind speeds. Since south-westerly winds prevail at Lillgrund, set-points were optimised for wind directions from 208° to 235° in steps of 1°, and for wind speeds from 5 to 12 m/s in steps of 1 m/s. Ideally the set-points would be changed as a function of not only wind speed and direction, but also turbulence intensity and atmospheric stability. However, it is less straightforward to measure turbulence intensity in real time, and atmospheric stability measurement would need additional sensors. Therefore, the set-points were calculated only for a typical turbulence intensity of 7%, and for slightly unstable atmospheric conditions with an Obukhov length of -200m (informed by an analysis of historical data from the nearby Drogden lighthouse). The resulting set-points are still predicted to work reasonably well, even if the turbulence intensity and atmospheric stability are different. Even in quite stable conditions, with Obukhov length 35m, the predicted gains were actually slightly greater than in the slightly unstable case; they could, however, be increased significantly further by re-optimising the set-points for these conditions, but to be conservative the set-points calculated for -200m were adopted for all conditions.

To simplify the practical implementation as far as possible, two further variations were investigated. Firstly, a controller for a single row of WTs — i.e. Row A (see Figure 1) — was designed, with the set-points of just those seven WTs optimized to maximize the total power from that row. The idea was to test this simpler setup first, and if successful, then move on to a test where all WTs would be controlled to maximize the power production of the entire farm. Secondly, a major simplification was to choose *fixed* set-points for each of the Row A WTs, i.e. not even varying these with wind speed and direction, which would provide some power gain over a certain range of wind conditions. This would be very straightforward to implement, not even requiring any estimated wind speed and direction other than to disable the control when the conditions were outside the range when some benefit could be expected.

These fixed set-points for each turbine were chosen by calculating a Gaussian-weighted average of the set-points around the point of greatest benefit — namely 8m/s and 222° — with Gaussian standard deviations of 10° and 3m/s, respectively. The resulting set-points, which are shown in Table 1, were then used irrespective of the actual wind speed and direction. Figure 4 shows contours of the predicted power gains for the full Row A controller (left) and the simplified controller with constant set-points (right, with the same colour scale). The power gains with the simplified controller are only slightly reduced, so the simplification seems worthwhile, especially as it removes any uncertainties that might arise from wind speed and direction estimation errors.

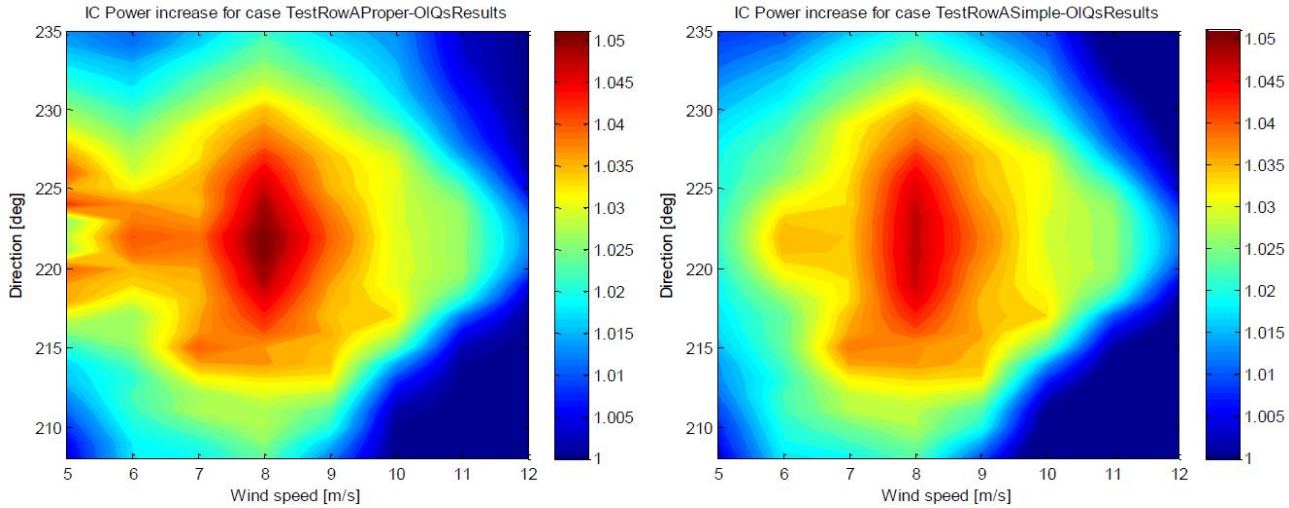

Figure 4: Energy gain contours for full controller (left) and simplified controller (right). The colours show the percentage gain in the power production of turbines A-01 to A-07.

Table 1: Constant Row A set-points for the simplified controller.

Turbine	Thrust reduction setpoint
Turbine A-01	0%
Turbine A-02	-6.6%
Turbine A-03	-5.0%
Turbine A-04	-3.8%
Turbine A-05	-3.8%
Turbine A-06	-4.6%
Turbine A-07	-9.0%

5. PRE-IMPLEMENTATION TESTING OF THE PROPOSED CONTROLLER

The controller set-points were generated by means of steady-state analyses using DNV's LongSim code. This relies on critical assumptions about the behavior of the turbine wakes. The controller is entirely open-loop, so there is no feedback from measurements to help compensate for any modelling inaccuracies. Since real revenue loss could result from an unsuccessful test, it was considered important to corroborate the predicted power gains by means of additional modelling.

Sensitivity of performance to uncertain atmospheric conditions has already been mentioned. The final controller has been simplified to make it independent of atmospheric conditions, and has been tested

against a range of conditions different from those for which it was designed (including radically different atmospheric stability). The LongSim platform incorporates a range of different engineering wake models, but a modified stability-dependent Ainslie model combined with particular models for wake superposition (using the sum of absolute deficits) and added turbulence (Quarton-Ainslie) has been used here. This combination has generally given good results when validated against a wide range of wind farm datasets [4]. To further increase confidence, two-way comparisons between LongSim and the very different Fuga model from DTU have been carried out, as described in Section 5.1. Finally, in Section 5.2 the controller, which has been designed for steady-state conditions, is further tested using LongSim's dynamic time-domain simulation capability, to give a better idea of how the controller might behave in practice under the influence of unsteady WPP flow conditions.

5.1 COMPARISON AGAINST FUGA

The Fuga model [5], [6] uses a linearized Navier-Stokes formulation of fluid flow with the individual turbines represented as actuator discs as derived from a detailed aerodynamic model HAWC2 [13]. This is entirely different from the engineering approach of LongSim. A two-way comparison was carried out for the Row A control: firstly, selected optimal set-points generated by Fuga for particular wind conditions were evaluated in steady-state LongSim calculations; and secondly, the effect of the DNV set-points was evaluated using Fuga.

The first comparison was not straightforward, because Fuga generates (as well as pitch angle) a rotor speed and power for any given wind speed, rather than an implementable controller set-point (the optimal mode gain or the equivalent power-rpm curve as described above). Therefore, power and thrust curves for each WT were generated, which would match the DTU settings as closely as possible at the wind speed seen by that particular WT (as predicted by Fuga).

The LongSim runs were carried out for three cases: the base case, the DTU settings, and DNV settings re-optimized for neutral conditions (since Fuga was run under the assumption of neutral conditions). The power gains are shown in Table 2, for five different combinations of wind speed and direction and associated with 6% ambient turbulence intensity. The gains are quite similar, although the LongSim gains are always slightly higher, but this might be expected, since they are calculated using the same code with which they were optimized.

Table 2: Comparison of power gains from Fuga-optimized and LongSim-optimized set-points when run in LongSim.

Direction, deg	Wind speed, m/s	Gain (%) with DTU settings	Gain (%) with DNV settings
217	6	3.44	3.71
220	12	1.65	1.88
221	8	2.96	4.44
224	11	1.94	2.26
227	6	3.72	5.09

For the second comparison, the full controller set-points generated by LongSim for Row A were tested by DTU using Fuga. The results are shown in Figure 5. The left plot is the same as in Figure 4, and corresponding Fuga result is shown on the right, with the same color scale. The LongSim steady calculations used wake parameters representative of 10-minute average results, which implicitly include meandering. Therefore the Fuga results have been post-processed with a Gaussian direction smearing of 2.75° standard deviation to account for wake meandering (estimated using a simplistic model linking lateral turbulence intensity with direction variability). The results show some clear differences, although there are enough similarities to confirm that, according to either code, reasonable power gains should be expected with the planned field test.

With Fuga, the power gains fall off rather faster as the wind direction moves away from the directly-aligned 222° direction, and also at the higher wind speeds. However, Fuga *also predicts much higher power gains* at the lowest wind speeds, which could be due to different assumptions about turbine behavior in the cut-in region, which can make a big difference: Fuga assumed zero power and a low thrust coefficient below a cut-in speed of 4 m/s, while LongSim assumed zero thrust coefficient once the power dropped to zero (thus nearer to 3 m/s). In practice, the supervisory control logic around cut-in wind speed can make a big difference here, because the thrust coefficient remains high, even at zero power, if the rotor speed does not immediately decrease to idling speed by feathering the pitch. Supervisory control details were not available here; in any case it is the interaction of this logic with the low-frequency variability of the wind speed, which is important. Note also that in this region, large differences in fractional power gain may be insignificant because the actual power production (the denominator) is very small.

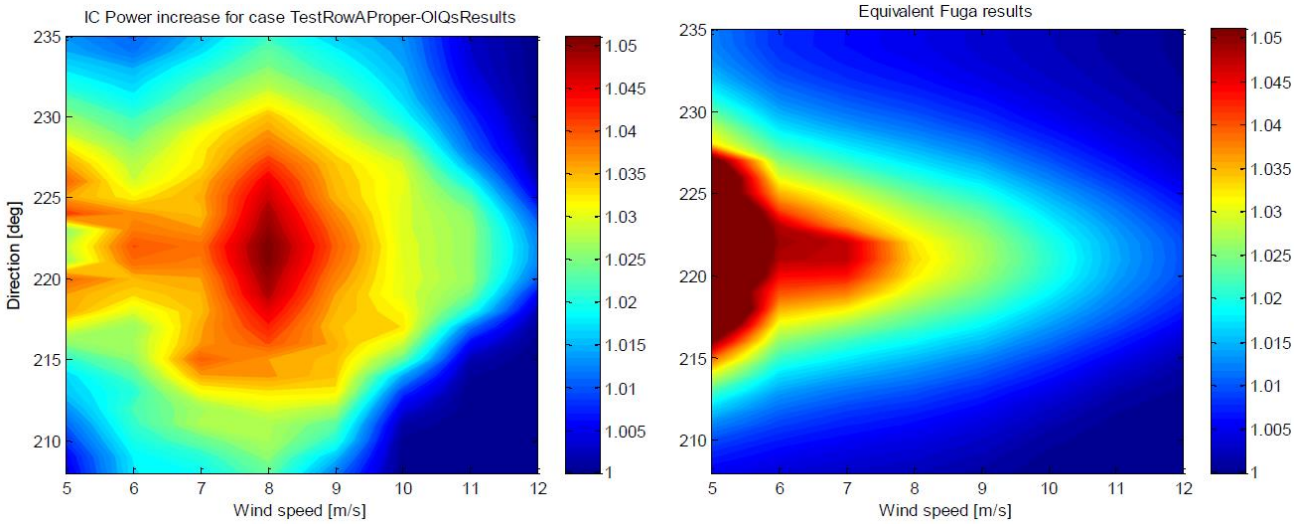

Figure 5: Energy gain contours for DNV set-points predicted by LongSim (left) and Fuga (right).

In summary, although there are some clear differences between LongSim and Fuga predictions, the level of agreement between these very different codes is good enough to give increased confidence that the

set-points proposed for the field test should give positive results, both codes predicting power gains in the range 0% – 5% over this range of wind conditions.

5.2 DYNAMIC TIME-DOMAIN SIMULATION

The proposed set-points were generated with LongSim using steady-state calculations, but real life can be very different from this. For a realistic test of how the controller might behave, LongSim's dynamic time-domain simulation capability was used. This simulates the effects of realistic *time-varying* meteorological conditions on short and long timescales, including slow and fast changes of wind speed and direction, spatial variations, turbulence, wake meandering and advection, as well as any relevant dynamic behaviour of the wind WT and WPP controllers. Since wake meandering is modelled explicitly, the wake model used in these simulations represents an *instantaneous* wake profile; this is derived from the profile used for the steady calculations by using the inverse of the meandering correction proposed by Ainslie [7].

The simulation is driven by a stochastic wind field covering the whole WPP. Here this was generated from a 22-hour time history of single-point wind measurements derived from historical SCADA data from Lillgrund. The single-point wind conditions are shown in Figure 6, while Figure 7 shows a snapshot of the wind field at one point in time. The variations in the wind field are suitably correlated in space and time, evolving as they advect downstream. WT wakes are embedded in the wind field, and their advection and meandering are driven by the three components of low-frequency turbulence in the wind field. High-frequency turbulence is added at each turbine position and is assumed to be uncorrelated between WTs.

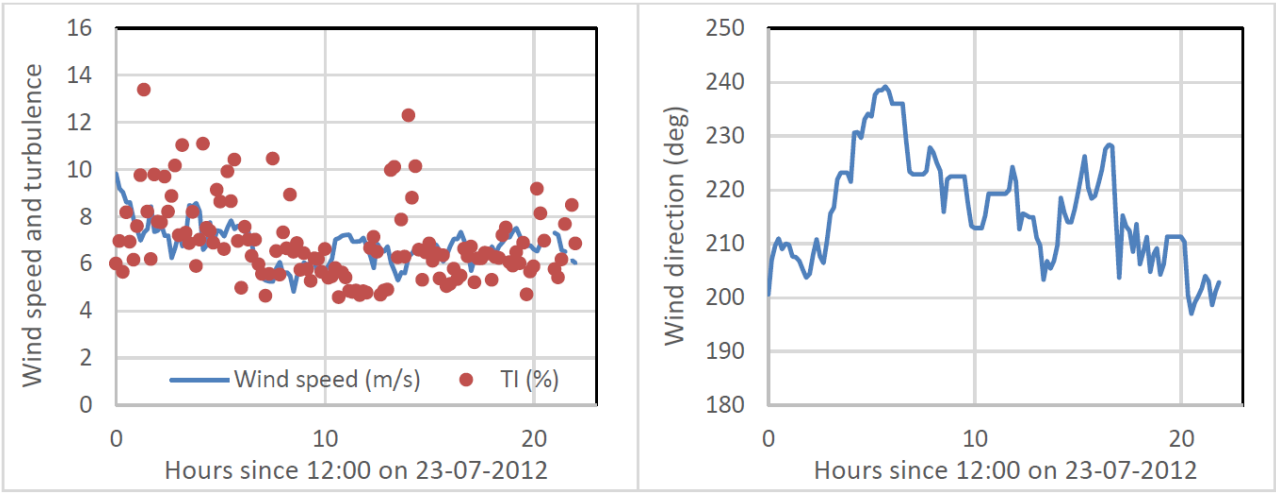

Figure 6: Wind conditions for time-domain simulation.

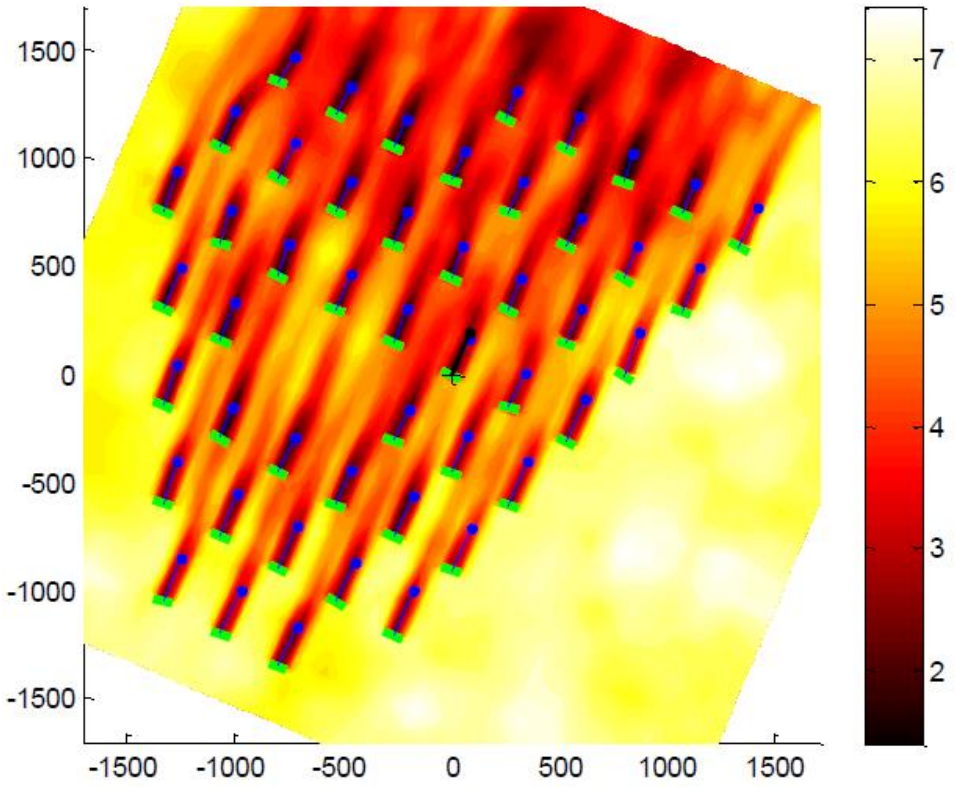

Figure 7: Contours of wind speed magnitude: snapshot at one point in time.

Simulation outputs include time histories of key variables at each turbine such as wind speed, power, pitch, rpm, yaw, etc. as well as various loads, as well as total wind farm power. Simulations were run for the base case with no control (all turbines operating with the nominal 'Mode A' settings), the optimized full Row A controller, and the simplified Row A controller to be used for the initial field tests. Figure 8 shows the percentage power gain of the developed Row A controller compared to the base case, for the full controller (left plot) and the simplified controller (right plot). The red lines show the quasi-static power gain interpolated from the steady-state results of Figure 4 at the wind conditions for each point in the simulation, while the blue line is the dynamic simulation result. The difference is due to all the dynamic effects captured by the simulation, such as spatial and temporal wind variations, wake meandering, etc., as well as the effect of WT control actions. In this case, a realistic yaw control algorithm resulted in about 20 discrete yaw adjustments per hour for each turbine on average. The interaction of supervisory control logic with low-frequency wind variations is explicitly modelled. As explained above, the effect of cut-in logic can be crucial to the power gains at low wind speed, but unfortunately, no details were available on the supervisory control logic used by these turbines, so this effect could not be included.

For the full controller, there is also the effect of a wind estimation algorithm used by the WPP controller (in this case using an average of measured wind speed and direction at all un-waked WTs), which is used to find the appropriate set-points that are updated every minute – this is not needed for the simplified controller, because the set-points for each turbine are fixed.

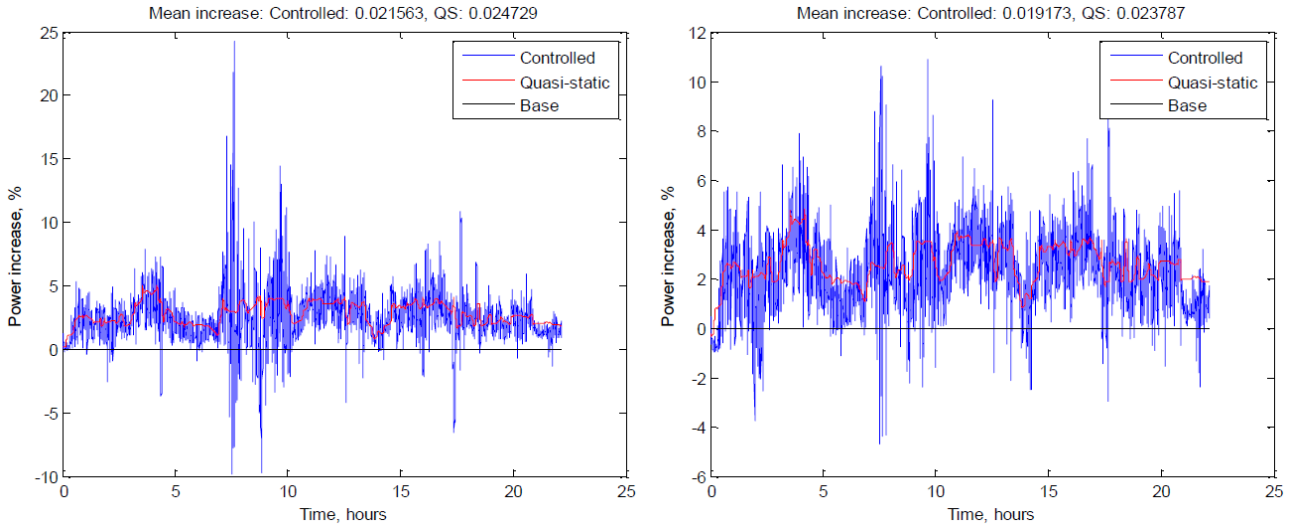

Figure 8: Power increase over base case: full controller (left) and simplified controller (right).

Clearly the dynamic effects at any point in time lead to considerable variations in the power gain, giving an indication of the scatter that can be expected in the field test results. Of course, the scatter will be further increased, because field data with controller toggled on and off will never be for identical wind conditions, as in these simulations. The simplified controller seems to avoid some of the more extreme variations, which might prove helpful. These results show that for this particular 22-hour simulated period, an average power gain of 1.92% is predicted with the simplified controller, compared to 2.38% if all dynamic effects are ignored (i.e. the quasi-static result). This gain is only slightly lower than with the full controller, where 2.16% increase is predicted, compared to the corresponding quasi-static result of 2.47%.

These increases are with respect to the 'Mode A' nominal controller setting. Subsequently it was found that the current nominal setting has been improved, to be more like the 'Optimal' setting described above. This increases the nominal power output by roughly 1%, meaning that the power gains from the field tests are expected to be smaller by an equivalent amount, i.e. around 1% rather than 2%.

In Figure 9, simulated damage equivalent loads (DELs) are plotted for three of the turbines in Row A for the base case and the simplified controller. Loads are higher for the middle and downstream turbines, A-04 and A-01 respectively, than for the upstream turbine A-07, because of wake turbulence, and for all three turbines, the simplified controller significantly reduces the loading compared to the base case. The figure shows the blade root out of plane bending moment DEL with Wöhler exponent 10 (appropriate for GRP). The tower base fore-aft overturning moment (Wöhler exponent 4, for steel) shows a very similar result (not shown). The loads are calculated using LongSim's surrogate loads model, as described in [2]. This gives some confidence that the proposed controller should be safe to implement in terms of fatigue loading, and is even likely to reduce loads significantly. Although only a small number of WTs are equipped with load sensors, it is hoped that the field test data, in addition to power production, may provide an opportunity for some validation of the loads model.

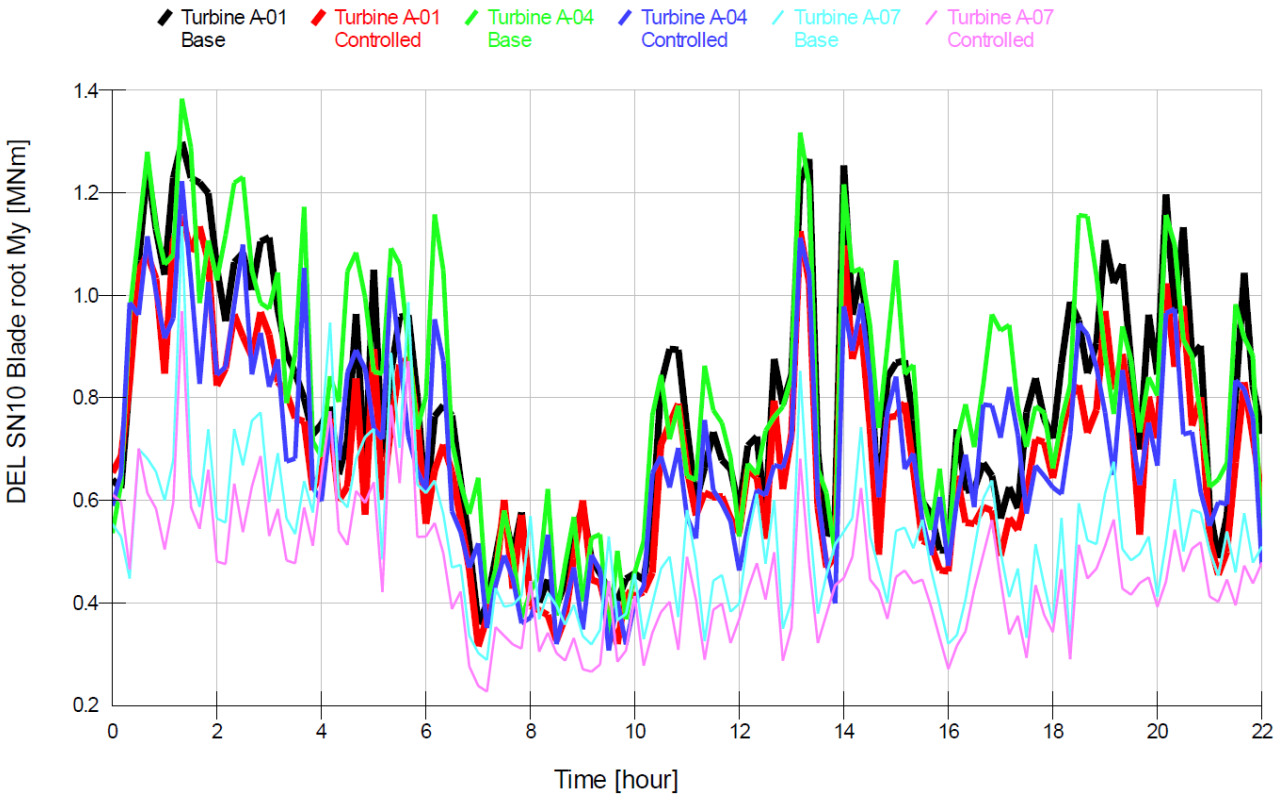

Figure 9: Damage equivalent loads: blade root bending moment at three Row A turbines.

6. FIELD TEST IMPLEMENTATION

The Row A simple controller was implemented at Lillgrund wind farm during December 2021. The implementation required only software updates, allowing the fine pitch angle and power-rpm curves to be changed for the individual (Row A) WTs. A small modification to the power-rpm curves in Figure 3 was implemented, making all curves the same as the existing Mode A curve below 8 rpm, in order to avoid issues around cut-in. This modification only affects performance at very low wind speed, and a LongSim analysis was performed to show that the effect on predicted power gains was very small, giving a result almost indistinguishable from Figure 4 (right hand plot).

The test is being run as a toggle test, in which the Row A turbine set-points are switched between the nominal setting and the simplified Row A controller settings every 50 minutes. This time interval was selected to be long enough for a reasonable period of settled operation after the wake effect has had time to travel the length of the wind farm, but short enough to expect periods of operation with the controller toggled off and on during *similar meteorological conditions*. A toggle interval which is not an exact fraction of a day also helps avoids any bias due to synchronisation of toggle switches with diurnal cycles in wind conditions such as atmospheric stability. The test is only active when the wind direction is in the south-west quadrant. The toggle testing was switched on in late December 2021, and data collection was scheduled to continue for some months at least, in order to be able to detect the expected small power increases in a statistically significant way. However, unfortunately there was initially a fault in the implementation, meaning that only changes in WT rotational speed was implemented, but the corresponding changes to the fine pitch settings were not happening. The error was identified by analysing the early results, and corrected as from 09-05-2022. The toggle testing has continued uninterrupted since then.

7. FIELD TEST RESULTS

Some preliminary field test results are presented here, based on analysis of the data recorded in the period from 09-05-2022 until 09-08-2022. The quantity of data is not sufficient to reach robust conclusions about the effectiveness of the control scheme, but it is enough to demonstrate that the test is proceeding correctly, and perhaps already gives some tentative indications about the performance of the controller. Further analysis will be undertaken once more data has been collected.

Figure 10 shows a sample of measured pitch angles from the Row A turbines (10-minute averages). Three distinct periods can be identified. The first few hours show baseline operation, which occurs when the wind direction is outside of the south-west quadrant, and all the turbines have the same fine pitch setting of around -0.2°. In the next period up to hour 1285, the pitch angles all increase above fine pitch, which indicates operation above rated wind speed. Then, starting from hour 1285, we clearly see the effect of toggling between nominal ('Control Off') and optimised 'Control On' settings. During 'Control On' periods, the fine pitch angles of the seven turbines in Row A are all different, matching the values corresponding to the optimised settings of Table 1. The wind speed increases above rated after hour 1305, causing the pitch angles to increase again. This gives a clear indication that the toggling of the fine pitch settings is implemented correctly.

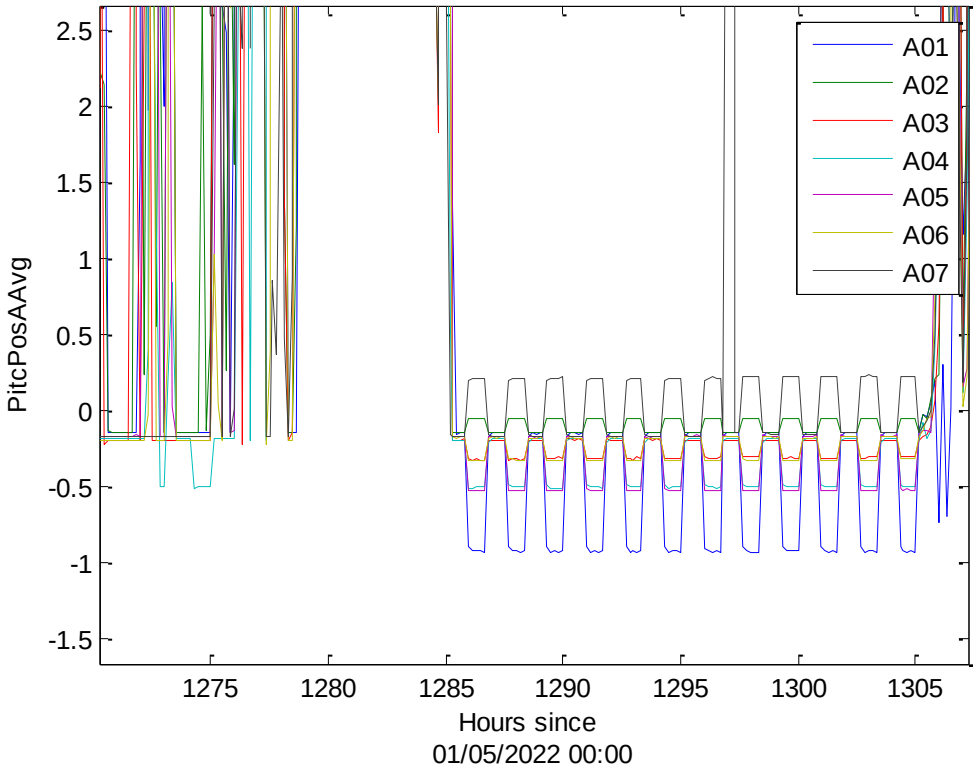

Figure 10: Sample of measured pitch angles at the Row A turbines.

Figure 11 plots power versus RPM for turbine A07 (10-minute average points for the period from 09-05-2022 until 04-07-2022). The points are clustered into two distinct groups, following two separate curves. The lower curve is for the nominal setting ('Control Off'), and the upper curve matches the 9% thrust reduction setting for turbine A07 with 'Control On'. The solid lines show the expected curves for the two settings. Similar plots were produced for the other Row A turbines, but the difference between the two curves is largest for turbine A07. These plots confirm that the toggling of the power-RPM settings is also correctly implemented.

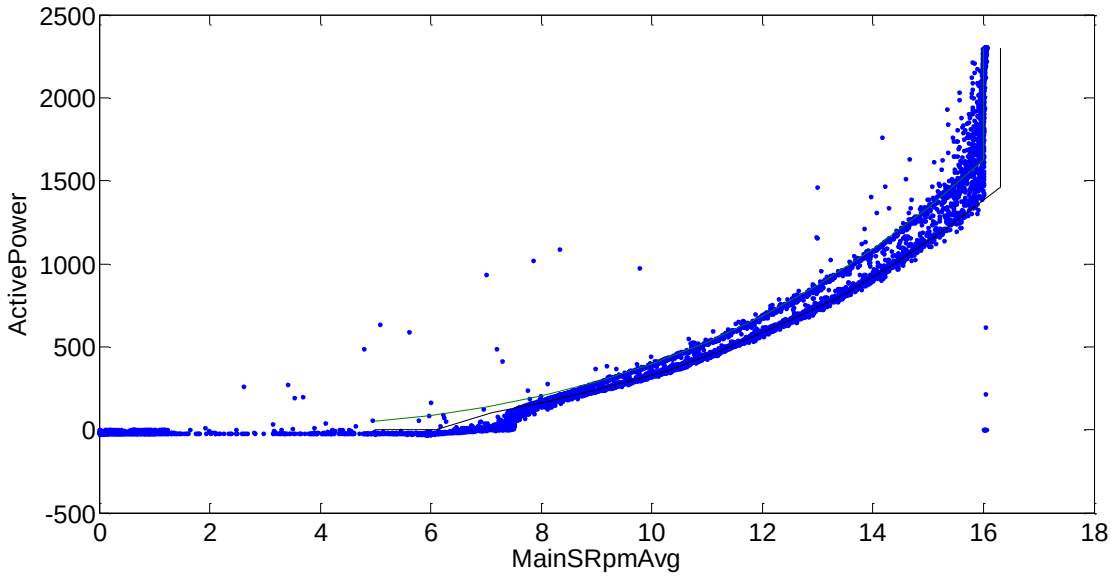

Figure 11: Measured power vs RPM for turbine A07 (10-minute average points, and expected curves).

Having confirmed the correct toggling action, the measured 10-minute average data was then analysed using the following steps:

1. Determine a measure of incident wind speed for each 10-minute datapoint. This reference wind speed has been calculated as the mean of the wind speeds measured at reference turbines B08, C08 and D08, since these three turbines are un-waked for wind directions of interest, and remain in nominal operation throughout, so their wind speed estimates are not affecting by toggling. Turbine A07 is not included, since its wind speed calibration would only be valid for nominal operation; also it is affected by the wake of turbine B08 for some wind directions of interest. A wind speed tolerance of 1 m/s is defined, and if the wind speed of one of the reference turbines differs from the mean of the three by more than this, the mean is recalculated just using the other two.
2. A similar procedure is used to determine a measure of wind direction for each 10-minute datapoint using the same three reference turbines. In this case, because of the 0/360° discontinuity in direction, averages are always calculated for the sine and cosine of each direction, and the arctangent is used to convert back to an average direction. However, before doing this, a calibration offset is calculated to account for an apparent systematic difference between the direction measurements at the three reference turbines: for each reference turbine, the mean direction over the whole period is calculated, and the calibration offset is taken as the difference between that mean and the mean for all three turbines. After applying these offsets, the mean direction is recalculated. As for wind speed, a wind direction tolerance of 10 degrees is defined, and if the wind direction of one of the reference turbines differs from the mean of the three by more than this, the mean is recalculated just using the other two. This is especially important because there are occasional periods when one of direction signals seems to 'stick'. Finally, a 'Wind direction difference' filter is defined so that points, where the wind direction at turbine A07 differs from the mean by more than the tolerance, can be rejected.
3. Four further filters are applied as follows. A wind direction filter only allows points where the wind direction is between 208° and 235°, as no power gains are expected outside this region. A wind speed filter only allows points with a reference wind speed between 3 and 13 m/s. A pitch filter rejects points where the pitch angle of any of the Row A turbines is less than -1° (as the fine pitch is always greater than this) or greater than 30° (meaning that the turbine is either stopped or generating well above rated, where no gains are expected). Finally, a power filter rejects points where any of the Row A or reference turbines is generating no power. These filters are needed, because there was no

turbine status signal available. About 95% of points are rejected by the filtering, about half because the wind direction is out of range. About 40% of points are rejected by the power filter, and 20% by the pitch filter, but of course many of these would have occurred when the wind direction was out of range anyway.

4. No signal was available to give the toggle status, so the following method was adopted to decide whether the control was on or off for each 10-minute datapoint. A preliminary indication of toggle on status is if the pitch angle of a Row A turbine is less than -0.8°, since the fine pitch for turbine A01 should be -1°, whereas the fine pitch of all turbines should be -0.2° during toggle off. This cannot be used in higher winds, because the pitch increase above fine pitch. Since the toggle period is 50 minutes, the first occurrence was located of five consecutive 10-minute points where the turbine A01 pitch angle was less than -0.8°. This was found on 9th May, which was the first day of toggling operation. The toggling was then assumed to occur exactly every 50 minutes for the rest of the dataset. A check towards the end of the period confirmed that the assumed toggle on periods were still exactly synchronised with periods of five consecutive points where the pitch of turbine A01 was less than -0.8°. Finally, since wakes are expected to take several minutes to advect along the length of Row A (about 7.5 minutes at 10 m/s and more at lower wind speeds), the first 10-minute datapoint of each toggle on or toggle off period was rejected, as this would represent a period of transition. See Figure 12 for an illustration confirming that the last four of each five toggle points is indeed selected correctly.

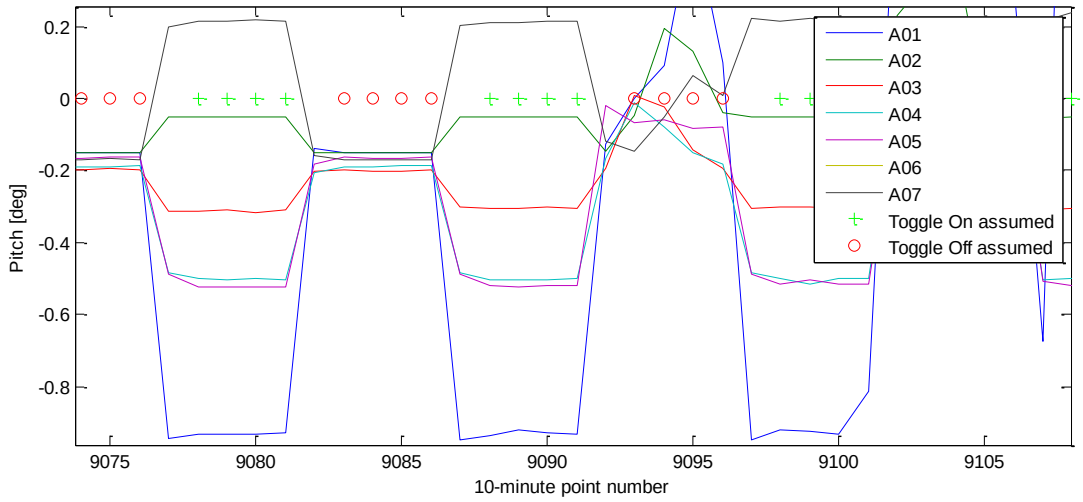

Figure 12: Plot confirming synchronisation of toggle periods towards the end of the period.

5. The total power of Row A for each datapoint was normalised by a reference power calculated as the mean power of the three reference turbines (B8, C8 and D8), to give a normalised power measure.

6. Having filtered the data as above and split it into two groups, toggle on and toggle off, the normalised power was binned by wind speed (1 m/s bins) and wind direction (5° bins). In each bin, the mean Toggle On normalised power was divided by the mean Toggle Off normalised power, to give a measure of the power gain.

The filtering yielded 549 valid datapoints, of which 296 were Toggle Off and 253 were Toggle On. When binning by both wind speed and direction, many bins are very sparsely populated, so only bins with at least three Toggle On and three Toggle Off points are shown (still a very small number for statistically meaningful results). Figure 13 shows the numbers of points in these remaining bins. Many bins remain empty, and no bin has more than 17 On or Off points. The normalised power increase is shown in Figure 14, while Figure 15 shows two possible measures of the reliability of the results in each bin: the left hand plot shows the minimum of the number of Off points or On points, whichever is less, in each bin, while the right hand plot shows the absolute difference between On and Off normalised power divided by the

standard deviation of points in the bin (actually the average of the standard deviations of On points and Off points).

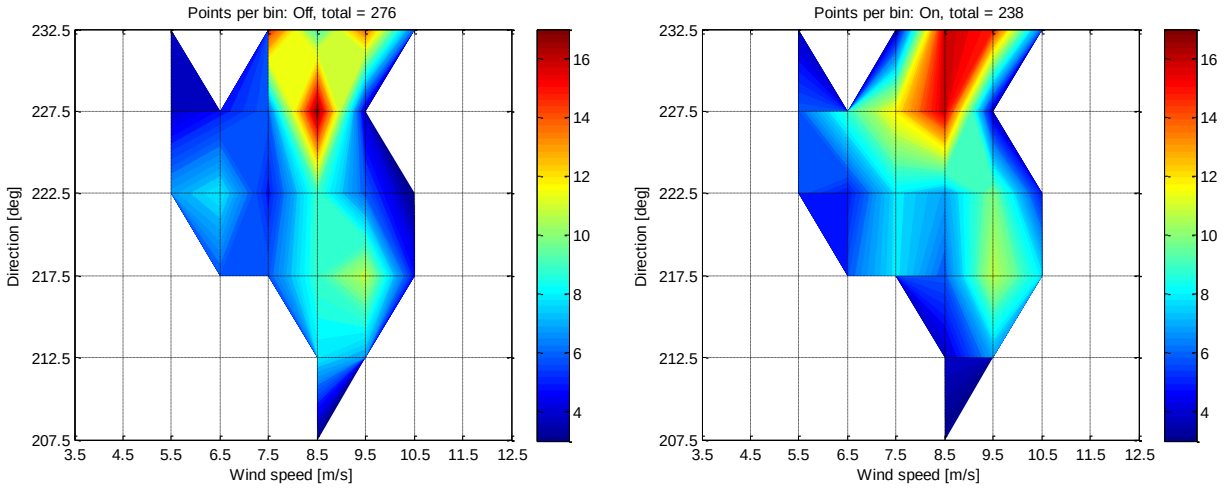

Figure 13: Numbers of Toggle Off and Toggle On points in the wind speed and direction bins.

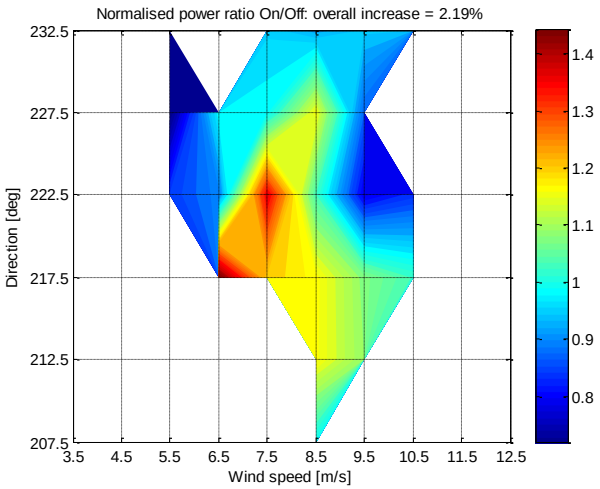

Figure 14: Normalised power gain.

From this, it is clear that there is not yet sufficient data to draw firm conclusions, but the indications are nevertheless encouraging. Over all of the populated bins in Figure 14, the mean power increase is about 2.2%, although this includes some thinly populated bins and some where the absolute difference in normalised power is small compared to standard deviation of the bin. Gains between 10% and 50% appear in some of the central bins, which is where positive gains are expected, although there are also some significantly negative gains away from the centre. The actual gains are shown in Table 3, with the minimum On or Off points in brackets; colour coding is used for positive and negative gains, and bold text for bins where the standard deviation based reliability measure exceeds 0.8. This begins to give a consistent picture of significant gains in the central part of the domain, and losses around the edges. This pattern compares well with the predictions of Figure 5, although the achieved gains are perhaps higher in the centre, and the losses are greater at the edges. This points to the desirability of moving on to a full dynamic control scheme with directionally-varying setpoints.

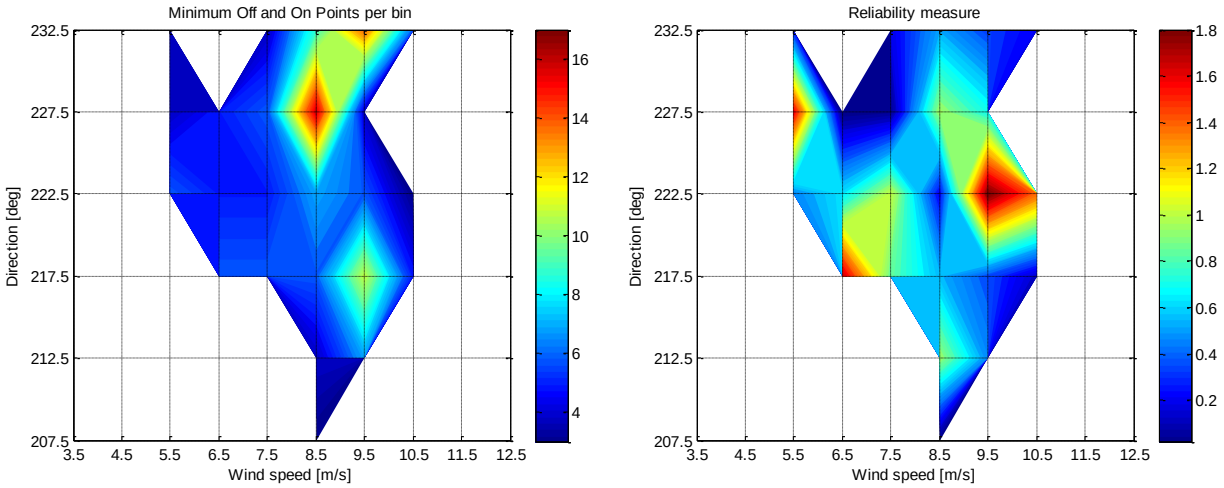

Figure 15: Reliability of bins: Minimum points per bin (left), and normalised power difference compared to standard deviation (right).

Table 3: Normalised power gain (ratio) in each bin, with minimum points in brackets.

	207.5°	212.5°	217.5°	222.5°	227.5°	232.5°
4.5 m/s		1.05 (6)				
5.5 m/s				0.85 (6)	**0.72 (4)**	0.94 (4)
6.5 m/s			**1.46 (6)**	0.90 (5)	0.99 (5)	
7.5 m/s			**1.21 (6)**	**1.37 (5)**	0.99 (6)	0.96 (4)
8.5 m/s	1.00 (3)	**1.17 (4)**	1.17 (6)	1.04 (7)	**1.16 (16)**	0.93 (9)
9.5 m/s		1.09 (7)	1.09 (11)	**0.80 (6)**	0.88 (4)	0.95 (14)
10.5 m/s			1.03 (5)	**0.84 (3)**		0.96 (6)

The data supplied from the wind farm for the present analysis did not so far include any measure of turbulence intensity. As this is a crucial factor for the wake dissipation, it would be useful if some turbulence intensity data could be made available to allow a more comprehensive final analysis once more data has been collected.

8. Conclusion

This report has described the design of an axial induction controller, which has recently been deployed for toggle testing at Lillgrund offshore wind farm. The controller, which is designed to increase the power output of one row of turbines on the farm (and should also reduce their fatigue loading significantly) has been made as simple as possible for ease of practical implementation and testing.

The controller was first tested using two independent and very different models, before being implemented on site. Data collection on site has been ongoing since the last few days of 2021, but due to an implementation error the validation study could only be based on data collected from 09-05-2022 and onwards.

The results from the two different models, although not the same, broadly concur in the level of gains that can be expected from the controller. Dynamic time-domain simulation has also confirmed that reasonable gains can be expected. The most important validation, of course, is the field test results, and these are showing initially promising results, although considerably more data still needs to be collected before robust conclusions can be drawn.

The field test is therefore continuing beyond the TotalControl project, and furthermore, once more data has been collected and analysed, it is considered to extend the field testing to include another row. Later on it is also considered to test a similar but dynamic controller for the whole wind farm, where the setpoints will change with wind direction.

REFERENCES

[1] TotalControl project, www.totalcontrolproject.eu.

[2] Optimization of WPP set-points, TotalControl deliverable D2.3, 12/04/2021, https://www.totalcontrolproject.eu/-/media/Sites/TotalControl/Publications/Public-deliverables/D2-3-Optimization-of-WPP-set-points.ashx?la=da&hash=186E4B9321CA8BC99FB4416A01847D1240BA22A6

[3] Burton T., Jenkins N., Bossanyi E., Sharpe D. and Graham M., Wind energy handbook, 3rd edition, 2021 (John Wiley & Sons).

[4] Neubert, A et al, Validating time-domain energy assessments: wake modelling considerations and recommendations, Proc. Resource Assessment & Analysis of Operating Wind Farms, WindEurope 2021.

[5] Ott S, Berg J and Nielsen M (2011). Linearized CFD models for wakes. Risø-R-1772(EN). Risø National Lab. for Sustainable Energy, 39p.

[6] Ott, S., Laan, P. van der and Larsen, G.C., Upgrade of Fuga (2019). TotalControl deliverable D1.7, https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5c75e36da&appId=PPGMS.

[7] Ainslie J. F., Calculating the flow field in the wake of wind turbines, J. Wind Engineering and Industrial Aerodynamics **27**, 1988.

[8] Vitulli, J. A., Larsen, G. C., Pedersen, M. M., Ott, S. and Friis-Møller, M. (2019): Optimal open loop wind farm control. Journal of Physics - Conference Series. 1256, 12 p., 012027.

[9] Bossanyi E., Dynamic wind power plant simulator for wind farm controller testing, Proc. Wind Energy Science conference, Copenhagen, June 2017.

[10] Bossanyi E., Combining induction control and wake steering for wind farm energy and fatigue loads optimization, Proc. Torque from Wind conference, Milan 2018, (IOP proceedings).

[11] Simon E. and Zamanbin A., Lillgrund measurement campaign dataset, TotalControl deliverable D1.1.

[12] Bossanyi, E.A., Ruisi, R., Larsen, G. Chr. and Pedersen, M.M., Axial induction control design for a field test at Lillgrund wind farm, Journal of Physics Conference Series: The Science of Making Torque from Wind (TORQUE 2022), IOP Publishing, DOI 10.1088/1742-6596/2265/4/042032, https://iopscience.iop.org/article/10.1088/1742-6596/2265/4/042032

[13] Larsen, T.J., HAWC2aero - the user's manual, Risø-R-631(ver. 1-1) (EN), Risø National Laboratory, DTU, 2008.