PAPER • **OPEN ACCESS**

Control co-design of a large offshore wind farm considering the effect of wind extractability

Journal of Physics: Conference Series **2767** (2024) 092026 doi:10.1088/1742-6596/2767/9/092026

Control co-design of a large offshore wind farm considering the effect of wind extractability

M L Pahus, T Nishino, A Kirby and C R Vogel

Department of Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, UK

E-mail: takafumi.nishino@eng.ox.ac.uk

Abstract. We present a prototype of a novel control co-design (CCD) method for large offshore wind farms. A traditional wind farm internal flow model using a Gaussian wake model and a Boolean yaw optimisation method is coupled, via the two-scale momentum theory, with an analytical model of 'wind extractability', to account for the farm-scale blockage effect that depends on the farm size and atmospheric boundary layer height. We also introduce a 'gridded' turbine layout optimisation into the CCD, which allows us to quickly find optimal angles of the primary axes of a regular turbine array to maximise the annual energy production (AEP) of a given number of turbines at a given offshore farm site. The developed CCD method is used to demonstrate an important trade-off between the minimum number of turbines and minimum farm area required to achieve a certain AEP at the Seagreen offshore farm site as an example. Our results also suggest that the yaw optimisation could help reduce the number of turbines required, but only when we aim to achieve a high AEP with a small farm area. The developed model could be further coupled with a financial model of wind farms in future studies.

1. Introduction

Wind energy is one of the leading sources of renewable energy. As a crucial part of the green energy transition, the demand for more and larger wind farms is projected to grow for many years to come. For large wind farms, the so-called farm blockage [1, 2] or induction [3] is known to affect their power production, but traditional 'wakes-only' models cannot capture this effect. Not only does this mean that current industry-standard wind farm models tend to overestimate power production levels, it also implies that the current way we design and control wind farms is not truly optimal. For example, a recent study by NREL on the 'control co-design' (CCD) of the next generation of wind farms [4] presents a computationally efficient method that enables a coupled optimisation of turbine layout and yaw control for wake steering, but it does not account for the farm blockage effect. As the impact of turbine layout on large wind farm power could be altered significantly by the 'two-scale interaction' of (turbine-scale) wake and (farm-scale) induction effects [5], it is important to develop a CCD method that considers the effect of mesoscale atmospheric response to large wind farms.

In this study, we propose a novel CCD method for large wind farm optimisation based on the two-scale momentum theory [3]. We start with an analysis of an 'unoptimised' wind farm (resembling the Seagreen offshore wind farm) using a methodology similar to [6] but with a more efficient and robust computational method. To account for the farm blockage effect, an existing wind farm model (PyWake [7]) is coupled with a novel 'wind extractability' model [8] predicting the strength of atmospheric response to the farm. Following this 'baseline' analysis,

two pillars of wind farm optimisation are investigated: (i) control, in the form of deliberate turbine yaw misalignment, and (ii) layout design. For the latter, a novel 'gridded' approach is developed, allowing quick prediction of the trade-off between the minimum number of turbines and minimum farm area required for a given annual energy production (AEP). Finally, results of a new CCD method combining the yaw control and 'gridded' layout design are presented to demonstrate the importance of wind extractability for the optimisation of large wind farms.

2. Theoretical background

The two-scale momentum theory for large wind farms, first proposed in [9] and then generalised in [3] and validated in [5], describes the power of a wind farm as a function of farm-average wind speed (analogous to the actuator disc theory describing the power of a turbine as a function of its axial induction factor). The non-dimensional farm momentum (NDFM) equation derived from the two-scale theory is

$$C_T^* \frac{\lambda}{C_{f0}} \beta^2 + \beta^\gamma = M \tag{1}$$

where

$$C_T^* = \frac{\sum_{i=1}^{n} T_i}{\frac{1}{2}\rho U_F^2 (nA)} \tag{2}$$

is a farm-average 'internal' thrust coefficient of turbines, n is the number of turbines, $\lambda = nA/S$ is the array density (where A is the rotor swept area and S the farm area), C_{f0} is an undisturbed sea-surface friction coefficient, γ is a friction exponent (assumed to be 2.0), $\beta = U_F/U_{F0}$ is the farm wind-speed reduction factor (where U_F and U_{F0} are the farm-average wind speeds for the cases with and without turbines present, respectively) and M is the momentum availability factor. Note that Eq. (1) describes, in a non-dimensionalised form, that the total turbine drag (first term of the left hand side) and the sea-surface friction drag (second term of the left hand side) are balanced by the momentum made available by the atmosphere (right hand side). This equation allows us to calculate a 'correct' value of β for a given set of C_T^*, λ, C_{f0}, γ and M.

2.1. Wind extractability factor

A recent study by Patel et al. [10] demonstrates that the momentum availability factor M in (1) can be estimated numerically from 'twin' mesoscale weather simulations (i.e., with and without a wind farm present). Their results also suggest that, for a realistic range of β values ($1 \geq \beta \geq 0.8$, depending largely on λ/C_{f0}), there is an approximately linear relationship between M and the farm induction factor $(1-\beta)$ as

$$M = 1 + \zeta(1-\beta) \tag{3}$$

where ζ is called the wind extractability factor. Although such 'twin' weather simulations allow us to estimate M (or ζ) as a time-dependent parameter (e.g., hourly) for a given offshore site, this approach seems too computationally expensive to be employed for wind farm optimisation (to maximise the AEP, for example). More recently, Kirby et al. [8] derived an analytical model of M using simple quasi-one-dimensional flow assumptions (considering only the effects of net advection, pressure gradient forcing and turbulent entrainment) and validated it against existing large-eddy simulations (LES) of a large wind farm in a conventionally neutral boundary layer (CNBL). They also showed that the analytical model of M could be simplified further to derive an analytical model of ζ, which could be expressed approximately as

$$\zeta = 1.18 + 2.18\frac{h_0}{C_{f0}L} \tag{4}$$

where h_0 is the 'undisturbed' atmospheric boundary layer (ABL) thickness without the farm present, and L is the streamwise length of the farm. In section 4.4 we adopt this analytical wind extractability model to consider the impact of farm-size ratio (L/h_0) on farm performance.

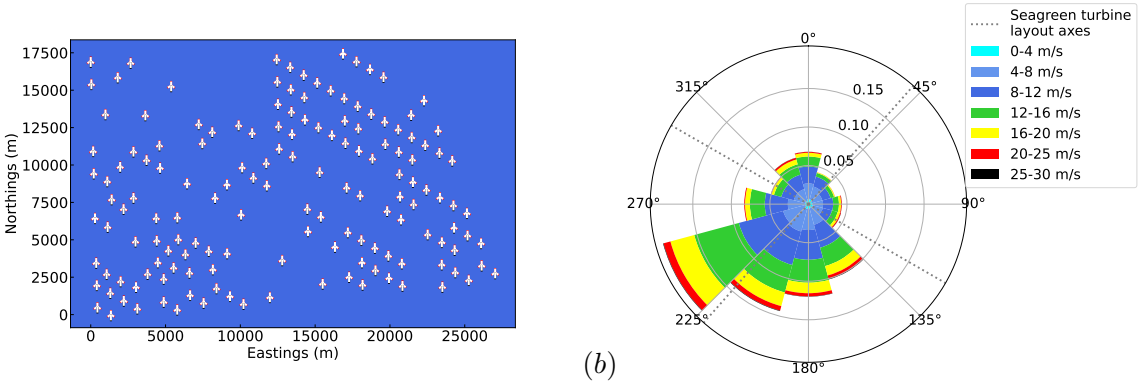

(*a*) (*b*)

Figure 1. Input data for the Seagreen offshore wind farm: (*a*) turbine layout and (*b*) wind rose.

3. Methodology

3.1. Baseline analysis

The methodology used for the 'baseline' analysis is similar to our previous study in [6]. We use the turbine layout and wind data from the Seagreen offshore wind farm (full configuration with 150 turbines), as shown in Fig. 1, but the turbine thrust and power curves are taken from the DTU 10 MW reference turbine [11] rather than the Vestas V164-10 MW turbine. The main difference from [6] is that we use PyWake in this study (instead of FLORIS used in [6]) as our 'internal' flow model to calculate C_T^*. We have successfully reproduced the study reported in [6] to verify our new computational code using PyWake. Table 1 summarises the setup.

Table 1. Summary of setup for the Seagreen baseline analysis.

Model/Parameter	Value/Reference
Wake model	Bastankhah & Porté-Agel [12]
Superposition method	Root-sum-square [13]
Number of horizontal grid points for β	10,000
Horizontal grid resolution for β	220 m (uniform in both directions)
Convergence criterion for β	Less than 0.01% from 'correct' β in Eq. (1)

3.2. Yaw optimisation

PyWake currently offers three different wake deflection models [7] to predict the effect of yaw misalignment. In this study we use the 'JimenezWakeDeflection' model based on [14] together with Bastankhah and Porté-Agel's Gaussian wake model [12]. These models are simpler than some other recent models and fast enough to allow us to run a large number of flow cases. For the optimisation of yaw angles, we follow a recent study by Stanley et al. [15] using a Boolean (instead of continuous) control, i.e., each turbine in the farm is either yawed at a predetermined angle (such as +20°) or not yawed (0°). The basic procedure is as follows:

(i) Sort all turbines in a given farm (for a given wind direction) from 'most upstream' to 'most downstream' and then use Jensen's simple cone-shaped wake expansion model [16] to assess, for each turbine, whether any of its downstream turbines are in its wake.

(ii) Calculate the power of 'unoptimised' farm (for a given 'undisturbed' wind speed U_{F0} and wind extractability ζ) as in the baseline analysis (i.e., no turbine is yawed). A few iterations are required here to find the 'correct' farm-upstream wind speed (to satisfy Eq. (1)).

(iii) For each turbine that was found (in Step (i)) to have any other turbines in its wake, check whether the farm power can be increased by yawing that turbine. Here the farm-upstream wind speed may need to be updated again (to satisfy Eq. (1)) but usually this is not required since the impact of yawing a single turbine on the farm-average C_T^* is very small.

(iv) Repeat (iii) from upstream turbines to downstream turbines, to determine which turbines should be yawed to maximise the farm power (for a given wind direction, U_{F0} and ζ).

(v) Calculate the AEP using given statistical data of wind direction, U_{F0} and ζ.

3.3. Gridded layout optimisation

In addition to the yaw optimisation, in this study we use a 'gridded' turbine layout optimisation method. This is a computationally fast method (much faster than many other methods, such as those based on a genetic algorithm) and also allows us to evaluate the impacts of farm size and farm shape separately, as shown later in Sections 4.3 and 4.4. In this method we aim to find an optimal gridded layout with turbines laid out regularly along two primary axes. The steps of creating a gridded layout and its farm area are summarised below and illustrated in Fig. 2:

(i) Consider a parallelogram array of turbines, with a regular turbine spacing of s along both axes. There are n_{row} and n_{col} turbines along the two axes, giving a total of $n = n_{row} \times n_{col}$ turbines, and the angle between the two axes is θ_{diff} (where $60° \leq \theta_{diff} \leq 120°$).

(ii) Rectangularise the array, by shifting all turbines above the top turbine in the first column (marked orange in Fig. 2) downwards. If $\theta_{diff} > 90°$, turbines below the bottom turbine in the first column are shifted upwards instead.

(iii) Define a quasi-rectangular farm area (marked blue in Fig. 2). Any locations at which the distance from the nearest turbine is larger than $s/\sqrt{2}$ are excluded from the farm area.

(iv) Rotate the whole array, so that the angle of the first axis is θ_1 (where $0° \leq \theta_1 \leq 120°$) and the angle of the second axis is $\theta_2 = \theta_1 + \theta_{diff}$. See Fig. 1($b$) for the angle definition.

Note that θ_{diff} is kept between 60° and 120° to keep the two axes dominant. Outside this range, the two intended primary axes would no longer be those along which the shortest inter-turbine distance is found. Also note that, to reduce computational costs, the upper limit of θ_{diff} can be reduced from 120° to min(120°, 180°$-\ \theta_1$) due to symmetry. The quasi-rectangular farm area S is approximately $Ns^2\sin\theta_{diff}$, which can be used for the calculation of λ in Eq. (1).

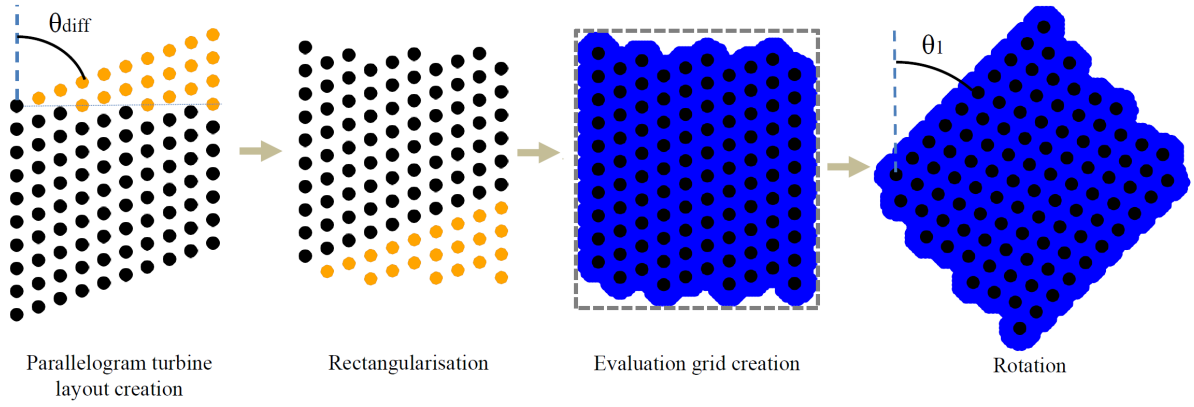

Figure 2. Steps of creating a gridded turbine layout and its quasi-rectangular farm area.

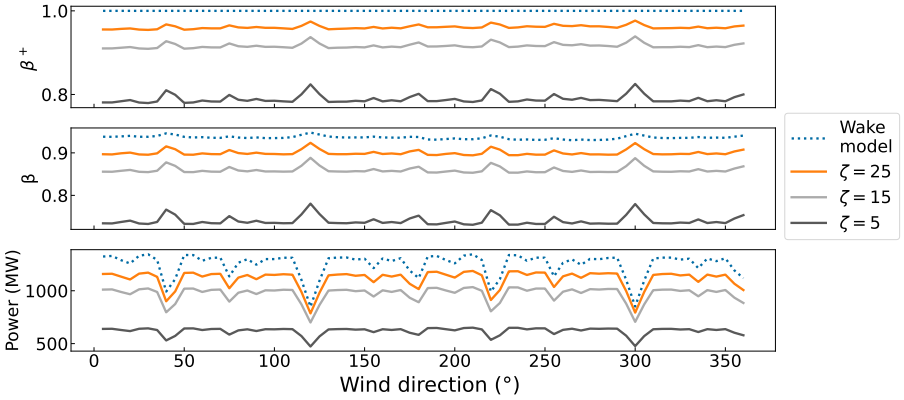

Figure 3. Effects of wind direction and wind extractability ζ on the farm-upstream wind speed reduction factor β^+ (top), farm-average wind speed reduction factor β (middle), and the power (bottom) of the Seagreen wind farm, for a fixed 'undisturbed' wind speed of $U_{F0} = 10$ m/s.

4. Results and discussion

4.1. Baseline cases (Seagreen)

Some example results of the baseline analysis are shown in Fig. 3, demonstrating the effects of wind direction and wind extractability. Here we define the farm-upstream wind speed reduction factor as $\beta^+ = U_F^+/U_{F0}$, where U_F^+ is the farm-upstream wind speed (i.e., the 'inflow' speed given to the wake calculation using PyWake). Note that $\beta^+ = 1$ when the farm blockage effect is not considered, as shown by the 'Wake model' results in Fig. 3. However, the farm-average wind speed reduction factor β is less than 1 even when the farm blockage is not considered (as expected from the fact that the wake model predicts only the 'deficit' of flow behind each turbine and not the acceleration of flow bypassing the turbine).

As already reported in the previous study [6], the farm blockage effect can be predicted using the two-scale coupling, i.e., by correcting the value of U_F^+ iteratively so that Eq. (1) is satisfied. Results for three different wind extractability scenarios, with a fixed ζ value of 5, 15 and 25, respectively, are plotted in Fig. 3, to confirm the trend that both β^+ and β (and thus the farm power) decrease as ζ decreases. (Note that this range of ζ was found by Patel et al. [10] to be typical for a large offshore farm; however, our recent study [8] suggests that $\zeta = 5$ is excessively low and unlikely to happen in reality.) It can also be seen that the farm power is lower when the wind direction is aligned with the two primary axes of turbine layout (40°/220° and 120°/300° for the Seagreen wind farm, see Fig. 1). However, the farm blockage effect is also smaller (i.e., the value of β^+ is higher) at these wind directions, due to the two-scale interaction [5].

The results presented in Fig. 3 are all for a fixed undisturbed wind speed of $U_{F0} = 10$ m/s, which is slightly lower than the rated wind speed for the DTU 10 MW turbine (11.4 m/s). Figure 4 shows the 'probability-weighted' average turbine power (averaged over all wind directions using the probability data shown in Fig. 1(b)) at different wind speeds. The wake model (without considering the farm blockage) predicts that the average turbine power will reach the rated power (10 MW) at a higher U_{F0} than the rated wind speed, due to the wake effects. When the farm blockage is considered, the average turbine power decreases further and an even higher U_{F0} is required for it to reach the rated power (e.g., 15 m/s at $\zeta = 5$).

Figure 4 also shows the AEP predicted using the same wind data shown in Fig. 1(b), again for hypothetical 'fixed' wind extractability scenarios (where ζ is assumed to stay at a certain value for the whole year). As expected, the AEP increases with the assumed value of ζ. Interestingly,

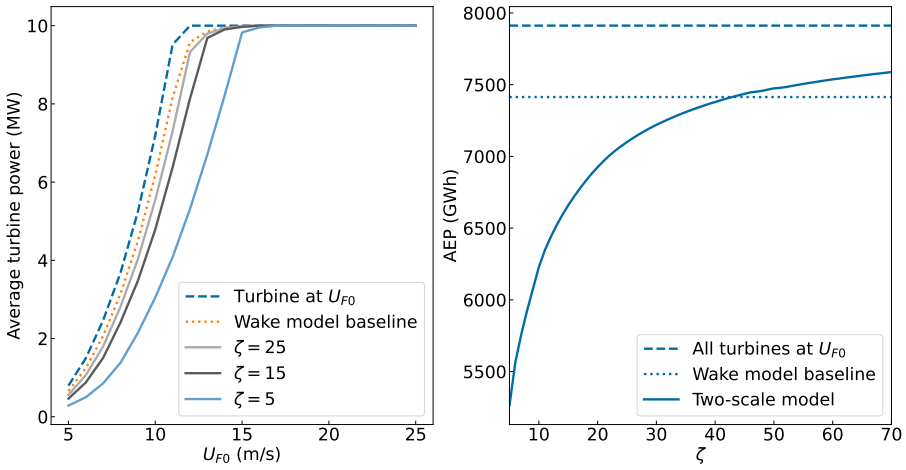

Figure 4. Effects of wind extractability ζ on the 'probability-weighted' average turbine power for all wind directions (left) and the AEP (right) predicted for the Seagreen wind farm. The AEP prediction for a given value of ζ takes about a minute on a standard desktop computer.

the two-scale coupled model gives a higher AEP than the original wake model prediction when ζ is very high (at $\zeta > 43$ in this case). This result may look somewhat counter-intuitive as the original wake model does not account for the farm blockage effect, but this happens essentially because the current wake model ignores any local acceleration of flow bypassing each turbine. This means that, when ζ is very high, the wake model under-predicts β and thus the two-scale coupling based on Eq. (1) results in a β^+ value of larger than 1 to satisfy the correct farm-scale momentum balance (with a wrong wake model prediction of the farm-internal flow). This suggests the importance of developing a better farm-internal flow model that accounts for not only the local deceleration of flow upstream and downstream of each turbine, but also the local acceleration of flow bypassing each turbine, in future studies.

4.2. Yaw optimised cases (Seagreen)

Figure 5 shows some results for the yaw optimised cases without considering the farm blockage effect. Here we have tested four different Boolean yaw control schemes, labelled as '$+20$', '± 20', '$+20, \pm 5$' and '$\pm 20, \pm 5$'. The scheme '$+20$', which is the scheme proposed in [15], allows each turbine to be either yawed at $+20°$ or not yawed, whereas the scheme '± 20' allows each turbine to be yawed at $+20°$ or $-20°$ or not yawed. In the schemes '$+20, \pm 5$' and '$\pm 20, \pm 5$', we follow the same procedure as '$+20$' and '± 20' during the first pass-through, and then allow a further addition or subtraction of $5°$ for each turbine during the second pass-through (i.e., we repeat Steps (iii) and (iv) twice). Hence, in the scheme '$+20, \pm 5$', for example, each turbine will be either yawed at $-5°$, $5°$, $15°$, $20°$ or $25°$, or not yawed, at the end of yaw optimisation.

As expected, the yaw optimisation is most effective when the wind direction is aligned with the two primary axes of turbine layout ($40°/220°$ and $120°/300°$) and the wind speed is below the rated speed (11.4 m/s). The farm power increases by about 10% in such cases, as shown on the left side of Fig. 5. A closer look at these plots also reveals that allowing for several different yaw angles (such as the scheme '$\pm 20, \pm 5$') results in power improvements at a wider range of wind directions, and this leads to a better improvement of the AEP (of up to about 1% for the case '$\pm 20, \pm 5$') as shown on the right side of Fig. 5. It should be noted, however, that these predictions of power improvements depend on the number of bins used for the wind direction

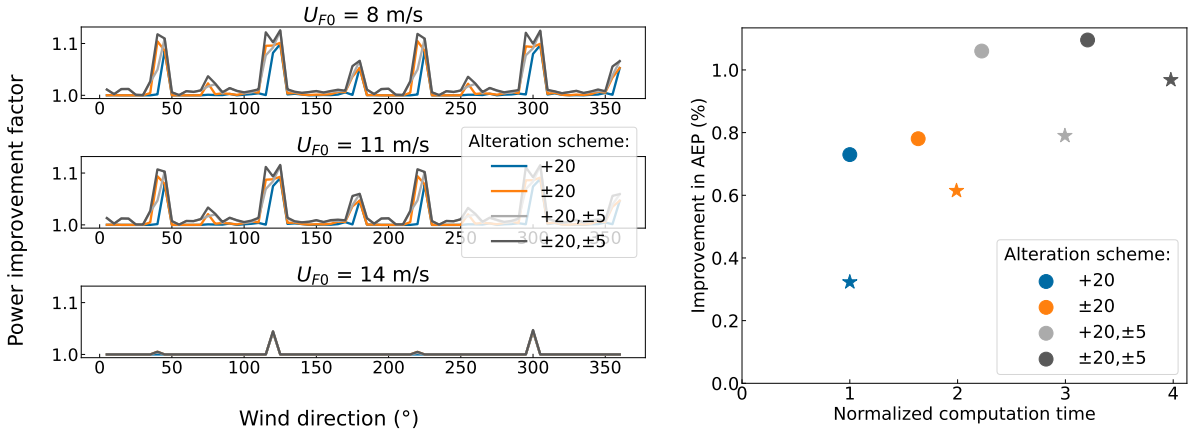

Figure 5. Power improvements achieved by four different yaw control schemes, for the Seagreen wind farm without farm blockage. Left: effects of wind speed and direction on the improvement factor. Right: AEP improvements plotted against required computational time. The circles and stars show the results using a 30°-binned and 5°-binned wind rose, respectively.

data. As the original Seagreen wind rose shown in Fig. 1(b) has only 12 bins (i.e., every 30°) for the wind direction, here we have created and used a refined wind rose with 72 bins (i.e., every 5°) simply by splitting the existing bins to smaller ones. If we use the original 30°-binned wind rose, the improvements in AEP are over-predicted, as shown in Fig. 5 (right). Nevertheless, in the rest of the paper we adopt the scheme '+20' with the original 30°-binned data (unless otherwise noted) to save computational costs. Note that the computational time plotted in Fig. 5 (right) has been normalised by that for the scheme '+20' for the 30°-binned cases and for the 5°-binned cases separately. The cost for the scheme '+20' using the 5°-binned data is about 6 times more than the same scheme using the 30°-binned data.

Now we combine the Boolean yaw optimisation method (the scheme '+20') with the two-scale approach. Figure 6 compares how the yaw optimisation affects the AEP predicted by the wake model (without the farm blockage) and the two-scale model (with the farm blockage) at five different 'fixed' wind extractability scenarios (ζ = 5, 10, 15, 20 and 25). It can be seen that the

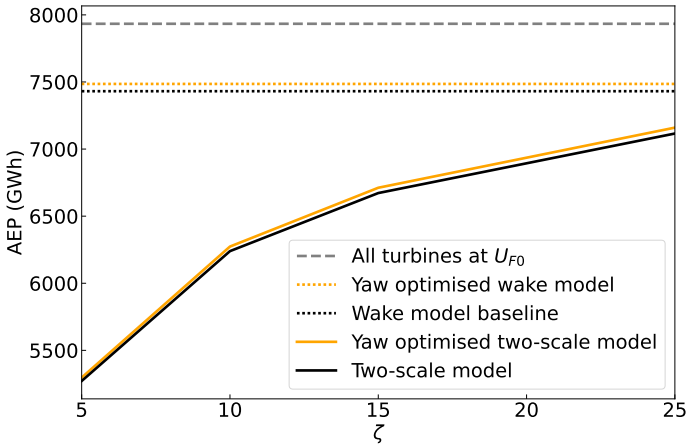

Figure 6. Effects of yaw optimisation on the AEP predicted for the Seagreen wind farm.

improvement in AEP due to yaw optimisation is much smaller than the possible impact of wind extractability (even though, as noted earlier, $\zeta = 5$ is excessively small and unlikely to happen in reality). These results further support the argument that future optimisation strategies for large wind farms should focus on finding methods to enhance the entrainment of momentum and energy from higher atmospheric layers into the farm [17, 18].

4.3. Optimal gridded layout

Now we show some results of 'gridded' layout optimisation. Here we still use the Seagreen wind data shown in Fig. 1(b) and the turbine thrust and power curves from the DTU 10 MW turbine but no longer use the Seagreen turbine layout shown in Fig. 1(a). By using the two-scale coupled approach (without the yaw optimisation in this subsection), we find an optimal gridded layout at the Seagreen site for a given wind extractability scenario.

Figure 7 shows the effects of different layouts (angles of the two primary axes) on the AEP for three 'fixed' wind extractability scenarios (with $\zeta = 5$, 15 and 25, respectively). Note that the number of turbines is fixed at $N = 100$ and the farm area is fixed at $S = 315$ km^2 in this analysis, and the AEP has been normalised by the 'max AEP', which is the AEP achieved by the same number of individual turbines having an undisturbed inflow speed (U_{F0}). It can be seen that there is little impact of wind extractability on the optimal layout axes, even though the AEP (for a given layout) is affected significantly by ζ, and not surprisingly, the optimal axis angles found here (approximately 35° and 125° for all three scenarios) are very close to those for the original Seagreen turbine layout (40° and 120° as shown in Fig. 1).

Although the results shown in Fig. 7 are for a fixed number of turbines and a fixed farm area, we have repeated the same analysis for different farm sizes and found that the optimal axis angles are still very similar. This suggests that the optimal axis angles (to maximise the AEP) depend predominantly on the wind rose (and not on the wind extractability) for a given farm site. This is an important finding, since this allows us to quickly find optimal axis angles first without considering the impact of wind extractability, and then look into the trade-off between the minimum number of turbines and the minimum farm area required to achieve a certain AEP (using the two-scale coupled approach only for the turbine layouts having the 'known' optimal axis angles) to explore a truly optimal farm design, as discussed in the next subsection.

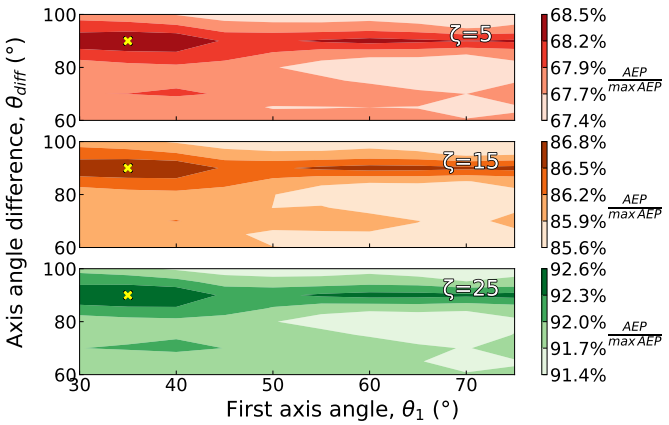

Figure 7. Effects of the angles of two primary axes of turbine layout on the AEP of an offshore wind farm at the Seagreen site (with a fixed number of turbines, $N = 100$, and a fixed farm area, $S = 315$ km^2) for $\zeta = 5$ (top), 15 (middle) and 25 (bottom), predicted using a 5°-binned wind rose. The yellow cross indicates the optimal layout with $\theta_1 = 35°$ and $\theta_2 = 125°$.

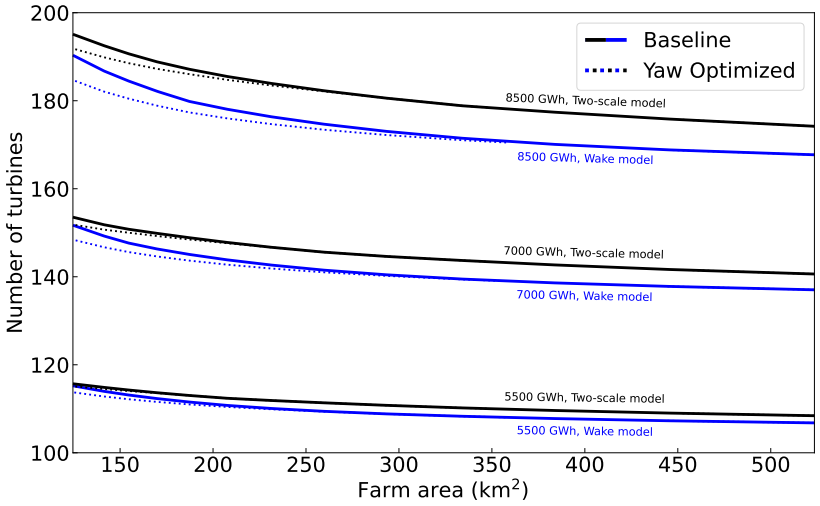

Figure 8. Contours of AEP for varying number of turbines and farm area, predicted by the wake model (blue lines) and the two-scale model (black lines) using a 10°-binned wind rose, for an optimal gridded layout for the Seagreen site (with the primary axes of 35°/215° and 125°/305°). Dotted and solid lines show the results with and without yaw optimisation, respectively.

4.4. Control co-design

Figure 8 shows combined effects of the number of turbines (N) and farm area (S) on the AEP of an offshore wind farm with an optimal gridded layout for the Seagreen site, predicted using the wake model (without farm blockage) and the two-scale model (with farm blockage), with and without the Boolean yaw optimisation (the scheme '+20'). For the two-scale model predictions, we have adopted Eq. (4) to account for the effect of farm size on the wind extractability factor ζ, assuming that $L = \sqrt{S}$ (regardless of the wind direction) and that the values of h_0 and C_{f0} are fixed at 500 m and 0.002, respectively. Note that Eq. (4) is based on a quasi-1D assumption, i.e., the farm width is much larger than the farm-layer thickness [8]. Also note that, in reality, the values of h_0 and C_{f0} are time-dependent and may have some correlations with the wind speed and/or direction, but here we have assumed typical offshore values for simplicity.

As can be seen from the figure, the general trends predicted by the wake model and the two-scale model are the same: the number of turbines required (to achieve a certain AEP) decreases as the farm area increases. This is because the power loss due to wake interactions is smaller when the turbines are more spread out. However, the two-scale model tends to predict a larger number of turbines required compared to the wake model, due to the farm blockage effect (or reduction of β^+) discussed earlier in Section 4.1. The differences between the wake model and two-scale model predictions tend to be small when the farm area is small, since ζ is high when S is small (i.e., the farm blockage is small for a small farm), but also when the area is very large ($S > 500$ km^2 in this case) since λ is small when S is large (i.e., the farm blockage should be small for a low-density farm). It can also be seen that the yaw optimisation reduces the number of turbines required, but only when we aim to achieve a high AEP with a small farm area.

It should be noted that, in a real wind farm design process, we usually need to consider not only the AEP but also the cost required, to minimise the levelised cost of electricity (LCoE), for example. The novel CCD method proposed in this study, which allows us to quickly predict the trade-off between the number of turbines and the farm area required to achieve a certain AEP at a given offshore farm site, could be combined with a financial model of wind turbines/farms in future studies, to further improve the design of future offshore wind farms.

5. Conclusions

In this study we have developed a novel CCD method for large offshore wind farms, taking into account the power loss due to farm-scale blockage as well as the power gain due to turbine yaw optimisation. The core idea demonstrated in this study is that a traditional 'engineering' wind farm flow model (such as PyWake used here) can be easily coupled with a recently proposed analytical 'wind extractability' model [8] to predict the farm blockage effect that depends on the farm size and the (undisturbed) atmospheric boundary layer height. We have also introduced a novel 'gridded' turbine layout optimisation, to quickly find optimal angles of the primary layout axes to maximise the AEP of a given number of turbines.

Our results suggest that, although the AEP achieved by a given layout depends significantly on the wind extractability, the optimal angles of the layout axes depend predominantly on the probability of wind speed and direction (and not the wind extractability or the farm size). This makes it much easier to predict the trade-off between the minimum number of turbines and the minimum farm area required for an optimal array to achieve a certain AEP at a given offshore farm site. We have demonstrated, by applying the developed two-scale method to the Seagreen offshore farm site, that this important trade-off can be easily predicted, for both cases with and without yaw optimisation. Our results also suggest that yaw optimisation could help reduce the number of turbines required, as expected, but only when we aim to achieve a high AEP with a small farm area. The developed model could be further coupled with a financial model of wind turbines/farms in future studies, to reduce the LCoE of future offshore wind farms.

Acknowledgements

We acknowledge the support of CRV's UKRI Future Leaders Fellowship (MR/V02504X/1) and AK's studentship from the NERC-Oxford Doctoral Training Partnership in Environmental Research (NE/S007474/1).

References

[1] Bleeg J, Purcell M, Ruisi R and Traiger E 2018 Wind farm blockage and the consequences of neglecting its impact on energy production *Energies* **11** 1609
[2] Sanchez Gomez M, Lundquist J K, Mirocha J D, Arthur R S, Muñoz-Esparza D and Robey R 2022 Can lidars assess wind plant blockage in simple terrain? A WRF-LES study *J. Renew. Sustain. Energ.* **14** 063303
[3] Nishino T and Dunstan T D 2020 Two-scale momentum theory for time-dependent modelling of large wind farms *J. Fluid Mech.* **894** A2
[4] Stanley A P J, Bay C J and Fleming P 2023 Enabling control co-design of the next generation of wind power plants *Wind Energ. Sci.* **8** 1341–1350
[5] Kirby A, Nishino T and Dunstan T D 2022 Two-scale interaction of wake and blockage effects in large wind farms *J. Fluid Mech.* **953** A39
[6] Legris L, Pahus M L, Nishino T and Perez-Campos E 2023 Prediction and mitigation of wind farm blockage losses considering mesoscale atmospheric response *Energies* **16** 386
[7] DTU Wind Energy https://gitlab.windenergy.dtu.dk/TOPFARM/PyWake
[8] Kirby A, Dunstan T D and Nishino T 2023 An analytical model of momentum availability for predicting large wind farm power *J. Fluid Mech.* **976** A24
[9] Nishino T 2016 Two-scale momentum theory for very large wind farms *J. Phys.: Conf. Ser.* **753** 032054
[10] Patel K, Dunstan T D and Nishino T 2021 Time-dependent upper limits to the performance of large wind farms due to mesoscale atmospheric response *Energies* **14** 6437
[11] Bac C et al. 2013 The DTU 10-MW reference wind turbine *Danish Wind Power Research 2013*
[12] Bastankhah M, Porté-Agel F 2014 A new analytical model for wind-turbine wakes *Renew. Energy* **70** 116–123
[13] Voutsinas S et al. 1990 On the analysis of wake effects in wind parks *Wind Engineering* **14** 204–219
[14] Jiménez Á, Crespo A and Migoya E 2010 Application of a LES technique to characterize the wake deflection of a wind turbine in yaw *Wind Energy* **13** 559–572
[15] Stanley A P J, Bay C, Mudafort R and Fleming P 2022 Fast yaw optimization for wind plant wake steering using Boolean yaw angles *Wind Energ. Sci.* **7** 741–757
[16] Jensen N O 1983 A note on wind turbine interaction *Technical Report Risø-M-2411* Risø National Laboratory
[17] Stevens R J A M 2023 Understanding wind farm power densities *J. Fluid Mech.* **958** F1
[18] Meyers J et al. 2022 Wind farm flow control: prospects and challenges *Wind Energ. Sci.* **7** 2271–2306