



Delft University of Technology

Dynamic wind farm flow control using free-vortex wake models

van den Broek, M.J.

DOI

[10.4233/uuid:8d2b92dc-51e9-4d8e-b1d8-c9e7bac211e7](https://doi.org/10.4233/uuid:8d2b92dc-51e9-4d8e-b1d8-c9e7bac211e7)

Publication date

2024

Document Version

Final published version

Citation (APA)

van den Broek, M. J. (2024). *Dynamic wind farm flow control using free-vortex wake models*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:8d2b92dc-51e9-4d8e-b1d8-c9e7bac211e7>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Dynamic wind farm flow control using free-vortex wake models

Maarten J. van den Broek



Dynamic wind farm flow control using free-vortex wake models

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology,
by the authority of the Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen,
chair of the Board for Doctorates,
to be defended publicly on
Friday 26 January 2024 at 12:30 o'clock

by

Maarten Jan VAN DEN BROEK

Master of Science in Mechanical Engineering,
Delft University of Technology, the Netherlands,
born in Zoetermeer, the Netherlands.

This dissertation has been approved by the promotor.

Composition of the doctoral committee:

Rector Magnificus	chairperson
Prof. dr. ir. J.W. van Wingerden	Delft University of Technology, promotor
Dr. ir. B. Sanderse	CWI, copromotor

Independent members:

Prof. dr. S.J. Watson	Delft University of Technology
Prof. dr. ir. A.C. Viré	Delft University of Technology
Prof. dr. L.Y. Pao	University of Colorado Boulder, US
Prof. dr. ir. J. Meyers	KU Leuven, Belgium
Dr. ir. J. Kreeft	Shell, the Netherlands



Keywords: wake steering, yaw misalignment, wind farm flow control, adjoint optimisation, economic model-predictive control, free-vortex wake

Printed by: Gildeprint

Style: Adapted from:
TU Delft House Style, with modifications by Moritz Beller
<https://github.com/Inventitech/phd-thesis-template>

Funding: This work is part of the research programme “Robust closed-loop wake steering for large densely spaced wind farms” with project number 17512, which is (partly) financed by the Dutch Research Council (NWO).

The author set this thesis in $\text{\LaTeX}$ using the Libertinus and Inconsolata fonts.

Copyright © 2024 by M.J. van den Broek

ISBN 978-94-6366-798-2

An electronic version of this dissertation is available at
<http://repository.tudelft.nl/>.

I'm going on an adventure!

Bilbo Baggins

Contents

Summary	ix
Samenvatting	xi
1 Introduction	1
1.1 In the face of climate change	3
1.2 On the growth of wind energy	4
1.3 Wind turbine wakes in wind farms	6
1.4 Wind farm flow control	8
1.5 Control for wake steering	10
1.6 This thesis	12
2 Free-vortex wake model	15
2.1 Introduction	17
2.2 Control-oriented free-vortex wake model	18
2.3 Optimisation for power maximisation	24
2.4 Results and discussion	27
2.5 Conclusions	39
2.A Partial derivatives	40
2.B Convergence study for numerical methods	48
3 Validation with wind tunnel data	51
3.1 Introduction	53
3.2 Free-vortex wake models	55
3.3 Parameter study and convergence	60
3.4 Validation with wind tunnel data	66
3.5 Results and discussion	71
3.6 Conclusions	80

4 Distributed wind farm control	83
4.1 Introduction	85
4.2 Model development	87
4.3 Controller synthesis	94
4.4 Methods for controller validation	98
4.5 Results and discussion	103
4.6 Conclusions	119
5 Conclusions and recommendations	121
5.1 Conclusions	123
5.2 Recommendations	125
References	127
Acknowledgments	139
Curriculum vitæ	141
List of publications	143

Summary

Wind energy plays an essential role in the transition to renewable energy sources as it provides a cost-effective alternative to fossil fuels. In that energy transition, the required growth of wind energy development means utilisation of the available wind resource should be maximised. This leads to the construction of large wind farms with densely spaced wind turbines to efficiently use the allocated space. Wind turbines operating in such close proximity inevitably experience losses in power production as aerodynamic wakes from upstream turbines impinge on downstream rotors. Part of this detrimental aerodynamic interaction is minimised through the optimisation of wind farm topology, which is, however, inflexible after construction and therefore has limited use in mitigating wake losses under time-varying conditions.

Wind farm flow control aims to control wind turbines in wind farms to reduce power losses and fatigue loading from wake effects to improve farm performance, whereas wind turbines are conventionally operated to maximise individual power production. One such strategy for cooperative control is wake steering, which makes use of the existing yaw degree of freedom that determines rotor alignment. Operation of the wind turbine rotor under intentional misalignment with the inflow reduces power production on that turbine but also leads to a lateral deflection of the wake. This allows redirecting the wind turbine wake away from downstream turbines, resulting in gains in wind farm power production.

The current state-of-the-art on wake steering control makes use of look-up tables based on prior optimisation with steady-state models. The validity of the models underlying these control strategies is limited under realistic, time-varying atmospheric conditions. Therefore, to assess the potential for improved power production from wake steering in wind farms under time-varying conditions, this thesis develops an economic model-predictive wind farm flow control strategy.

A control-oriented model of the wind turbine wake is formulated based on an actuator-disc rotor model and simulation with free-vortex wake methods. The associated novel derivation of the adjoint system enables gradient-based control optimisation for efficient implementation in an economic model-predictive control setting. The model captures the curled-wake dynamics that are essential to modelling wake deflection. Additionally, the control test cases achieve dynamic control results and wake steering under a time-varying inflow direction.

The free-vortex methods used in the model development have generally been used for rotor aerodynamics and near-wake predictions, with limited validation in the mid to far wake. We assess the accuracy of power predictions with the control-oriented wake model for wind turbines operating under yaw misalignment using three sets of data from wind tunnel experiments. The model achieves considerable accuracy in power predictions with steady misalignment and under dynamic control variations. The level of accuracy combined with the achieved computational efficiency demonstrates that the dynamic model has potential for integration in a wake steering controller.

Subsequently, model-based control optimisation with the free-vortex wake model is extended to the wind-farm scale by constructing a distributed optimisation approach. A network of free-vortex wake models is formed to parallelise control optimisation and account for the cumulative effects of wake steering in rows of wind turbines. The novel controller is tested in a large-eddy simulation environment and compared against a state-of-the-art look-up table approach based on steady-state model optimisation and an extension with wind direction preview. Under realistic variations in wind direction and wind speed, the preview-enabled look-up table controller yielded the largest gains in power production. The novel controller based on the free-vortex wake produced smaller gains in these conditions, while yielding more power under large changes in wind direction. Additionally, the novel controller demonstrated potential for a substantial reduction in the increased demands on yaw actuator usage from wake steering.

The control strategy and model development presented in this thesis advance the state-of-the-art on wake steering control, demonstrating that accounting for the temporal dynamics of wake propagation can improve the power production of wind farms under realistic inflow conditions. This contributes to the efficient operation of current and future wind farms by making effective use of existing control degrees of freedom.

Samenvatting

Windenergie speelt een centrale rol in de transitie naar hernieuwbare energiebronnen omdat het een kostenefficiënt alternatief is voor fossiele brandstoffen. In die energietransitie betekent de benodigde groei in de ontwikkeling van windenergie dat benutting van de beschikbare windbron gemaximaliseerd moet worden. Dit leidt tot constructie van grote windparken waar windturbines dicht op elkaar geplaatst worden om efficiënt de beschikbare ruimte te gebruiken. Windturbines die zo dicht bij elkaar staan ervaren onvermijdelijke verliezen in vermogensproductie omdat het aerodynamische zog van bovenwindse turbines benedenwindse rotors raakt. Deel van deze nadelige aerodynamische interactie wordt geminimaliseerd door de optimalisatie van de windparktopologie, welke echter inflexibel is na de constructie en daardoor beperkt bruikbaar is in het verminderen van de zogverliezen onder tijdsvariërende omstandigheden.

Windparkstromingsregelingen streven ernaar windturbines in windparken aan te sturen om de vermogensverliezen en vermoeiingsbelastingen door zogeffecten te verminderen en de prestatie van het park te verbeteren, terwijl windturbines conventioneel aangestuurd worden om individuele vermogensproductie te maximaliseren. Een strategie voor coöperatieve aansturing is zogsturing, waarbij de gierbewegingsvrijheid benut wordt die de oriëntatie van de rotor bepaalt. De exploitatie van een windturbinerotor met scheefstand ten opzichte van de instroom vermindert de vermogensproductie van de turbine, maar leidt ook tot een laterale afbuiging van het zog. Dit maakt het mogelijk om het windturbinezog om een benedenwindse turbine heen te sturen, wat resulteert in een toename van de vermogensproductie van het windpark.

De huidige, meest geavanceerde, zogsturingsregelingen maken gebruik van opzoektabellen gebaseerd op voorgaande optimalisatie met tijdsafhankelijke modellen. De validiteit van de modellen onderliggend aan deze regelingen is beperkt onder realistische, tijdsvariërende omstandigheden. Om het potentieel vast te stellen voor een toename in vermogensproductie door zogsturing in windparken onder tijdsvariërende atmosferische omstandigheden, ontwikkelt dit proefschrift een economische modelvoorspellende windparkstromingsregeling.

Een regelingsgericht model van het windturbinezog wordt geformuleerd op basis van model van de rotor als remmende schijf en simulatie met vrije-wervelmethoden. De bijbehorende nieuwe afleiding van het adjuncte systeem maakt gradiëntgebaseerde regelingsoptimalisatie mogelijk voor efficiënte implementatie

in een economische modelvoorspellende regelstructuur. Het model beschrijft de gekrulde zogdynamica die essentieel zijn voor het modelleren van de afbuiging van het zog. Daarbij worden in testscenario's dynamische regelsignalen gevonden en wordt zogsturing onder tijdsvariërende instroom bereikt.

De vrije-wervelmethoden die gebruikt zijn voor de modelontwikkeling zijn algemeen gebruikt voor rotoraerodynamica en voorspellingen in het zog nabij de rotor, met beperkte validatie voor het midden- tot verre deel van het zog. We stellen de nauwkeurigheid vast van vermogensvoorspellingen met het regelingsgerichte zogmodel voor windturbines geëxploiteerd met gierscheefstand met drie verzamelende data van windtunnelexperimenten. Het model bereikt noemenswaardige precisie in vermogensvoorspellingen met constante scheefstand en met dynamische variatie van de aansturing. De mate van nauwkeurigheid gecombineerd met de bereikte computationele efficiëntie demonstreert dat het dynamische model potentieel heeft voor integratie in een zogsturingsregeling.

Daaropvolgend wordt de modelgebaseerde regelingsoptimalisatie met het vrije-wervelmodel uitgebreid naar de windparkschaal door het construeren van een gedistribueerde optimalisatiemethode. Een netwerk van vrije-wervelmodellen wordt gevormd om de regelingsoptimalisatie te paralleliseren en daarbij rekening te houden met de cumulatieve effecten van zogsturing in rijen van wind turbines. De nieuwe regelaar wordt getest in een grote-wervelsimulatieomgeving en vergeleken met een geavanceerde benadering gebaseerd op opzoektabelen met tijdsafhankelijke modeloptimalisatie en een uitbreiding met vooruitblik op de windrichting. Bij realistische variaties in windrichting en windsnelheid levert de opzoektabelregelaar met windrichtingsvooruitblik de grootste winst in vermogensopbrengst. De nieuwe regelaar gebaseerd op het vrije-wervelmodel levert kleinere winsten in deze omstandigheden, maar brengt meer vermogen op bij grote windrichtingsveranderingen. Daarbij demonstreert de nieuwe regelaar het potentieel voor een substantiële vermindering van de toegenomen eisen aan de gieraandrijving door zogsturing.

De regelingsstrategie en modelontwikkeling die zijn gepresenteerd in dit proefschrift bevorderen de huidige staat van zogsturingsregelingen, waarbij gedemonstreerd wordt dat het rekening houden met de tijdsdynamica van het zog kan zorgen voor een toename van de vermogensproductie in windparken onder realistische instroomcondities. Dit draagt bij aan de efficiënte exploitatie van huidige en toekomstige windparken door effectief gebruik te maken van bestaande graden van vrijheid.

1

Introduction

The impact of climate change necessitates drastic action. Wind energy development is one aspect of the shift to renewable energy that is required. Wind farm flow control aims to improve the efficiency of wind farms, effectively using available degrees of freedom for improving power production and decreasing the cost of energy. Wake steering is a control strategy that makes use of intentional yaw misalignment to redirect wakes away from downstream turbines. This cooperative control strategy for wind farms has been demonstrated in simulation, wind tunnel studies, and field experiments. Existing literature and the industry state-of-the-art focus on steady-state results, whereas, in reality, the wind is naturally time-varying in both speed and direction. This thesis works towards wind farm flow control, using wake steering, under time-varying atmospheric conditions.

Chapter contents

1.1	In the face of climate change	3
1.2	On the growth of wind energy	4
1.3	Wind turbine wakes in wind farms	6
1.4	Wind farm flow control	8
1.5	Control for wake steering	10
1.6	This thesis	12
1.6.1	Outline	13

1.1 In the face of climate change

Climate change should need no introduction. The average global temperature has already risen by 1.1°C in 2011–2020 compared to 1850–1900, which is fast approaching the 1.5°C limit to warming that is targeted in global policy. Furthermore, climate change is not just about average warming – extremes in terms of heat and drought, as well as heavy rains and flooding are increasingly common. This impacts global populations from agriculture and changing ecosystems to infrastructure and displacement of people [1].

Now, climate change can be concluded with high confidence to be caused by human activity – primarily from greenhouse gas emissions – and it is likely that the 1.5°C limit will be exceeded in the coming years [1]. Figure 1.1 depicts the sharp decrease in net emissions that is required to achieve this goal. Systemic change is required – on all levels from individuals to sovereign nation-states and global corporations – as a species we must lower our footprint on the planet.

This means efficiently using available resources and reducing consumption, as well as coordinating a global effort to transition from fossil fuels to renewable energy to limit anthropogenic sources of climate change [3]. As an example of united efforts for change, the European Union committed to transitioning to 40 % of energy from renewable sources by 2030, coming from less than 20 % in 2019 [4]. Photovoltaics, wave energy, hydropower, and wind energy are just a few of the possible alternatives to fossil fuels. Each has their own strengths and limitations and they should be distributed based on local resource availability to form a robust, reliable, efficient renewable energy system.




Figure 1.1: CO₂-equivalent global emissions of greenhouse gases as presented by the IPCC [1], comparing modelled pathways that limit warming to 1.5°C with limited or no overshoot to the modelled trajectories assuming policies implemented by the end of 2020. Major reductions in net emissions are required to reach the targets for limiting the rise of global average temperature to 1.5°C . Data from [2].

1.2 On the growth of wind energy

One part of the transition to renewable energy sources is the expansion of wind utilisation. As an alternative to fossil fuels, the EU has set out to achieve more than 300 GW of installed offshore renewable energy by 2050 to move towards net zero emissions [4]. The Netherlands alone has committed to installing 21 GW of offshore wind power generation capacity by 2030 [5], for which the majority still needs to be constructed as illustrated in Figure 1.2. These may be ambitious targets, but without ambitious goals sufficient impact will not be achieved – they reflect the urgency of taking deliberate action.

“Offshore renewable energy is one of the most promising routes to increase future power generation in the coming years in a way that meets Europe’s decarbonisation objectives and expected rise in electricity demand in an affordable manner.” – European Commission [4]

Economically, both onshore and offshore wind are viable alternatives to fossil fuels. Onshore wind is one of the most cost-effective sources of energy and, although offshore wind energy is more expensive in terms of installation costs and maintenance, the drive to realise cost-effective wind power has driven the levelised cost-of-energy for offshore wind down by 59 % in the time from 2010 to 2022 [6].

As the wind industry has matured, tenders for wind farm construction have become highly competitive, putting pressure on developers and wind turbine manufacturers to operate on tight margins. In this competitive market, investments are required to build the necessary infrastructure for constructing wind turbines and developing wind farms. Raw materials such as steel and rare-earth metals




Figure 1.2: Development of wind energy capacity for power generation in the Netherlands [7]. With an installed offshore wind capacity of 4.7 GW by the end of 2023, another 17.9 GW of additional offshore wind power is planned to reach the Dutch 2030 targets [5].

are sensitive to inflation and geopolitical stresses, production facilities need to be constantly adjusted to ever-larger wind turbine rotors, and installation vessels and port facilities need to grow to provide sufficient capacity for offshore installation [8]. Protectionist laws such as the US Jones' act hamper international cooperation and limit the global developments of offshore wind [9].

Another side of infrastructure for a renewable energy system is the electrical grid. New connections are required for wind farms and stability needs to be guaranteed under varying loads. Consumer demand will need to be flexible to deal with periods of over- and undercapacity in renewable energy production due to weather variations. Price incentives can drive demand to balance supply and power-to-x solutions can provide storage of excess power for later use, for example by driving hydrogen production on surplus power.

Next to the infrastructure requirements, the wind industry is in dire need of personnel. By 2026, the wind industry is expected to need over half a million qualified technicians, for which the majority still need to be trained [10]. Labour shortages lead to a lack of employees, from technicians for construction and maintenance to highly-educated scientists and engineers for developing robust wind power systems and advancing the state-of-the-art.

The development of wind energy also places increasing demands on available




Figure 1.3: Terrain elevation [11] and mean wind speed at 100 m above sea level [12]. The combination of higher mean wind speeds offshore and relatively shallow water depths in the North Sea and the Baltic Sea make these areas particularly suitable for offshore wind energy.

space. The North Sea and Baltic Sea are particularly suited for offshore wind energy given low water depth and higher average wind speeds offshore as shown in Figure 1.3, which means that conventional bottom-fixed wind turbines are a suitable option here. However, even on the sea, space is not unlimited as illustrated in Figure 1.4, showing the current operational and planned wind farms in the Dutch North Sea. These plans need to consider other stakeholders such as shipping, mining, oil and gas, and fishing, as well as the protection of natural areas.

1.3 Wind turbine wakes in wind farms

The motivation for the work in this dissertation is in efficiently making use of the space that is available. For that, we need to go a little bit deeper into the aerodynamics of wind turbines, where the effects that matter range from global weather patterns all the way down to the aerofoils of wind turbine blades [13].

Wind turbine blades generate lift to drive rotation of the wind turbine rotor. In turn, this drives the generator for conversion of aerodynamic torque to electrical power. The aerodynamic power p extracted by a wind turbine is calculated as,

$$p = c_p \frac{1}{2} \rho A_r u_\infty^3. \quad (1.1)$$

which scales linearly with rotor area A_r and with the cube of inflow velocity u_∞ . The power coefficient c_p is a measure of the efficiency of the rotor in extracting aerodynamic power. For modern wind turbines, individual turbine performance is approaching the theoretical limits in turbine efficiency – the Betz limit [14]. The natural way forward is then an increase in size to generate more power per turbine, as a larger swept area directly yields larger aerodynamic power. Additionally, the cubic scaling with inflow velocity means that slightly higher wind speeds lead to significantly larger wind resource in terms of power.

As energy is extracted from the air, wind turbines leave behind a wake – a region of lower wind speed with higher turbulence levels. Turbines operating in waked conditions experience a reduction in power production and higher fatigue loads. These aerodynamic wake interactions lead to large losses in wind farms.

Figure 1.4 illustrates how wind turbine wakes affect wind farms. The available space is densely filled with wind turbines for maximal power extraction. For the dominant wind direction, the layout designed is such that aerodynamic interaction is minimal. However, major losses in power production appear as the wind direction changes and aligns with rows of wind turbines, such that wakes impinge on downstream rotors.




Figure 1.4: Overview of current and planned offshore wind farms in the Dutch North Sea, adapted from [5]. Wind farm layouts are designed to minimise wake interaction for the dominant wind directions and given space constraints. For dominant wind direction, wake interaction is minimal. However, fixed-bottom wind turbines are inflexible to adjust for time-varying atmospheric conditions leading to wake losses when the wind aligns with rows of wind turbines. Space allocated to offshore wind is limited so wind farms need to be designed and operated efficiently considering intra-farm wakes. Due to close proximity, inter-farm wakes are likely to have considerable impact on prospective energy yield.

1.4 Wind farm flow control

The control of wind turbines in wind farms is a broad field, for which a useful distinction can be made between wind farm control and wind farm flow control, where the former is concerned mostly with the electrical grid and the latter deals with mitigation of wake effects in wind farms. While layouts may be optimised to minimise the detrimental effects of aerodynamic interaction under the most common wind directions, they are inflexible to deal with time-varying atmospheric conditions [15]. Repositioning of floating wind turbines to avoid wake impingement [16] is, as of now, still a thing of the future.

Wind farm flow control techniques can be roughly divided in three categories: wake steering, induction control, and wake mixing techniques [17]. Induction control utilises lowering of the wind turbine thrust to reduce power extraction. In theory, this leaves more aerodynamic power for downstream turbines to improve farm production. However, the impact on annual energy production was minimal in wind tunnel testing and field experiments [18–20].

Wake mixing strategies are based on dynamic thrust variations to stimulate breakdown of the wind turbine wake and improve mixing with the free-stream flow. These are realised either through collective pitch variations [21, 22] or by individual pitch control [23]. The potential of these dynamic control strategies




Figure 1.5: Conventional greedy control focuses on individual wind turbine performance, controlling them to alignment with the inflow. Wake steering makes use of an intentional yaw misalignment γ with the free-stream wind direction to redirect the wake around downstream turbines. For wind directions with a lot of wake interaction, a small sacrifice in terms of power on misaligned turbines yields a net gain in wind farm production.

was demonstrated in wind tunnel experiments [24, 25], but not yet tested in the field.

The most mature of the three flow control techniques is wake steering, which uses intentional misalignment of the rotor with the free-stream flow to redirect the wake away from downstream turbines. When operating under yaw misalignment, the wake from the wind turbine is deflected laterally. This lateral movement induces the formation of a counter-rotating vortex pair, which produces a characteristic ‘kidney-shaped’ cross-section of the wake deficit [26–30], as illustrated in Figure 1.6.

Effective use of wake redirection has been demonstrated to increase wind farm power production and reduce fatigue loads on downstream turbines in simulation studies [31, 32] and wind tunnel experiments [33–35]. A series of field experiments have shown situational gains in power production from wake steering [36–41], although estimated improvements in annual energy production were inconclusive until demonstrated by Howland *et al.* [42]. Siemens-Gamesa have even implemented the technique in a commercially available product, *Wake Adapt*, claiming up to 1 % improvement in annual energy production [43], which supports confidence in the strategy. The added economic value from wake steering control may even be higher than the increase in energy production during low wind speed situations when electricity prices tend to be higher [44].




Figure 1.6: Operation under yaw misalignment produces a counter-rotating vortex pair that yields a characteristic ‘kidney-shaped’ cross section of the wake. The free-vortex wake model (Chapter 2) captures these curled wake dynamics and models the wake deflection that occurs from operation under yaw misalignment.

1.5 Control for wake steering

The control strategies for wind farm flow control can be roughly divided in two groups: model-based and model-free control. The model-free approaches to control use measurements of flow and wind turbine outputs to identify optimal control strategies. A closed-loop, model-free yaw controller [33] and extremum-seeking control [45] have been demonstrated to produce power gains from wake steering under steady flow conditions. Recently, a closed-loop controller based on reinforcement learning achieved gains in wind farm power production from wake steering under steady conditions [46]. These fully data-driven techniques have only been demonstrated under quasi-steady flow conditions. They generalise poorly to other operating conditions and the rate of convergence is generally slow, limiting implementation in field experiments.

On the other hand, model-based approaches make use of prior knowledge of the underlying physics and empirical engineering approximations to model wake effects and optimise for optimal control set-points. The most common approach to wake steering uses steady-state engineering wake models, such as those in the FLORIS toolbox [47], e.g. the Gaussian wake model [48] or the curled wake model [49]. These steady-state models provide an approximation of the time-averaged wake deficits to pre-optimize yaw angle for quasi-steady inflow conditions. Look-up tables with steady-state optimal set-points are then used for implementation in wind farms [31, 35, 42, 50, 51].

These strategies based on look-up tables are efficient from a computational perspective and predictable for analysis of applied control signals. However, wind directions and wind speeds vary over time and inclusion of wake propagation dynamics have a significant impact on performance of the wind farm [52]. Control optimisation needs to consider the dynamics of realistic wind direction variations [51]. Additionally, higher wind direction variability at lower wind speeds – where wake steering is most valuable – has contributed to poor performance of wake steering control [41].

A challenging time-varying control problem has to be solved for wake steering under realistic inflow conditions [53]. Given the efficiency and widespread adoption of the engineering wake models, one approach to accommodate dynamics in model-based control is the extension of steady-state models to incorporate wake propagation over time [54–56].

On the other hand, physics-based numerical models of wind farm flows may also be applied for control optimisation. These often naturally include the time-varying dynamics that are required to account for wake propagation through the wind farm. Large-eddy simulation has been used for wind farm flow control, making use of the adjoint formulation to enable gradient-based control optimisation [57–59]. Although the potential gains from flow control were demonstrated,

this was far from real-time control application. Recent work has reduced the computational cost of LES-based optimisation towards real-time performance by coarsening the mesh and adjusting control parameters [60].

Approaches such as WFSim [61] and FRED [62–64] have attempted to use two-dimensional fluid dynamics to reduce model complexity for control optimisation, which has been shown to be effective for induction-based active power control [65] and for power maximisation through induction control with the derivation of the adjoints [66]. However, these models are unsuitable for wake steering as the wake dynamics of wake steering are fundamentally three-dimensional [63].

The models based on conventional computational fluid dynamics calculate flow on a dense grid, where not all states are controllable. Free-vortex wake (FVW) methods, on the contrary, simulate flow dynamics based on vortex elements convected as Lagrangian particles, potentially reducing the number of states simulated to a skeletal representation of the wake – only those states that are actually affected by the wind turbine rotor are simulated. These FVW methods are based on the vorticity formulation of the Navier-Stokes equations and require the assumption of inviscid, potential flow [67, 68].

The FVW method has been initially applied to wake modelling for helicopter rotors with a focus on tip vortices [69]. From there, it was adapted to application for unsteady aerodynamics in wakes of wind turbine rotors [70] and demonstrated in several studies of rotor aerodynamics and wake stability [71–73]. The flexibility of a meshless formulation is especially relevant for studies on floating wind turbines, where the rotor is not stationary [74–76]. Of particular note for this thesis is the demonstration of the formation of a curled wake with a simple actuator-disc model simulated using the FVW [77], which has the potential to be implemented as a control-oriented model of the wind turbine wake.

The control-oriented wind farm flow model is a core element of an economic model-predictive control (EMPC) framework as illustrated in Figure 1.7. In this framework, the wind farm model provides predictions of the impact of control signals on future wind farm performance. These predictions are used to find optimal controls on a finite prediction horizon and implemented in a receding horizon setting – the first (set of) control(s) is then implemented, after which the problem is shifted in time and solved again [78]. Economic model-predictive control considers an objective such as power maximisation, where the extremum is not known a priori. On the other hand, conventional model-predictive control aims to drive an objective functional to zero, such as for tracking a reference signal for active power control.

The framework illustrated in Figure 1.7 is closed-loop; it shows how measurements from the wind farm feed back into the controller. This is essential for implementation of such a control strategy, as the model state may diverge from the




Figure 1.7: The closed-loop economic model-predictive control framework. The wind farm model is used for control optimisation to maximise power production in the wind farm. State estimation provides updates the model state and parameters based on measurements.

real wind farm state if left uncorrected. A state estimation strategy assimilates the data from wind farm measurements and provides corrections to the model state to avoid this divergence. Additionally, model parameters may be adapted to adjust for changing atmospheric conditions or variations in plant performance. This closed-loop implementation may be achieved by means of an Ensemble Kalman filter [50, 79, 80] or other data assimilation techniques.

1.6 This thesis

In the current state of model-based wind farm flow control, the implementation of yaw-based wake steering based on steady-state models has demonstrated potential for improving wind farm power production. However, for realistic, time-varying wind directions, the dynamics of wake propagation may impact the effectiveness of wake redirection. In order to further advance the power gains that may be achieved with wake steering control, we formulate the following objective:

Thesis objective: Develop an economic model-predictive wind farm flow control strategy and assess the potential for improved power production from wake steering in wind farms under time-varying conditions.

At the core of such a model-based control strategy is a control-oriented model of the wind farm flow. Physics-based models may naturally include the temporal dynamics of wake propagation. The free-vortex wake is a simplified, skeletal model for wake simulation that has been shown to capture curled wake dynamics with a limited number of states. Given the potential of the model formulation and the previous results on adjoint optimisation for wind farm flow control, the first contribution concerns development of a control-oriented wake model:

Contribution I: Formulation of a free-vortex wake model for gradient-based control optimisation and demonstration in an economic model-predictive control setting.

This control-oriented wake model should provide power predictions in the mid to far wake at a low computational cost to enable real-time control optimisation in wind farms. However, free-vortex methods have mostly been applied with a focus on rotor aerodynamics and flow in the near wake. Therefore, this thesis contributes to the validation of the free-vortex wake model for its intended purpose of control optimisation for power maximisation:

Contribution II: Assessment of the accuracy of power predictions with the dynamic free-vortex wake model in the mid to far wake of wind turbines operating under yaw misalignment.

Finally, assessing accuracy does not necessarily prove suitability for control optimisation. In the end, a demonstration in a realistic wind farm scenario against a proper reference controller is necessary to evaluate performance of the model in a controller. As initial results with the free-vortex wake consider individual wakes, this requires a strategy for scaling the model-predictive control optimisation to large wind farms. This leads to the final contribution of this thesis – advancing wake steering under realistic conditions:

Contribution III: Development of a distributed economic model-predictive control strategy with a network of free-vortex wake models to provide a scalable solution for wake steering in large wind farms under time-varying conditions.

1.6.1 Outline

This chapter has introduced the overarching objective for this thesis and the contributions made within, which will be presented in the subsequent chapters. The contents of the chapters have been written as independent articles and, therefore, each contains its own individual introduction and conclusions.

Firstly, **Chapter 2** develops a control-oriented model of the wind turbine wake based on free-vortex methods. The associated derivation of the adjoint equations enables gradient-based optimisation, which provides the basis for using the model in an economic model-predictive control strategy. Two test cases demonstrate the potential of this approach for generating dynamic control solutions and realising wake steering subject to a changing wind direction.

Secondly, **Chapter 3** evaluates the suitability of the level of model simplifications for power predictions downstream from wind turbines operating under yaw misalignment. Three different models of the wind turbine rotor are implemented using free-vortex wake methods to simulate the wake, one of which is the actuator-disc model from Chapter 2. The accuracy of these models is validated with three sets of data from wind tunnel experiments for steady conditions and with dynamic control variations.

Lastly, **Chapter 4** describes the construction of a distributed strategy for dynamic wake steering control. Individual free-vortex wake models are connected into a network to represent the total wind farm flow. The distribution of the optimisation problem over all turbines in the farm enables parallel optimisation of wind turbine controls. The novel control strategy is tested in a large-eddy simulation environment and compared against the state-of-the-art for wake steering.

Finally, **Chapter 5** gathers the general conclusions that can be drawn from the preceding chapters and provides recommendations for future research.

2

Free-vortex wake model

Wind farm flow control aims to improve wind turbine performance by reducing aerodynamic wake interaction between turbines. Dynamic, physics-based models of wind farm flows have been essential for exploring control strategies such as wake redirection and dynamic induction control. Free-vortex methods can provide a computationally efficient way to model wind turbine wake dynamics for control optimisation. We present a control-oriented free-vortex wake model of a 2D and 3D actuator disc to represent wind turbine wakes. The novel derivation of the discrete adjoint equations allows efficient gradient evaluation for gradient-based optimisation in an economic model-predictive control algorithm. Initial results are presented for mean power maximisation in a two-turbine case study. An induction control signal is found using the 2D model that is roughly periodic and supports previous results on dynamic induction control to stimulate wake mixing. The 3D model formulation effectively models a curled wake under yaw misalignment. Under time-varying wind direction, the optimisation finds solutions demonstrating both wake steering and a smooth transition to greedy control. The free-vortex wake model with gradient information shows potential for efficient optimisation and provides a promising way to further explore dynamic wind farm flow control.

Chapter contents

Ch. 2

2.1	Introduction	17
2.2	Control-oriented free-vortex wake model	18
2.2.1	General formulation	19
2.2.2	Two-dimensional model specifics	22
2.2.3	Three-dimensional model specifics	23
2.3	Optimisation for power maximisation	24
2.3.1	Economic model-predictive control	24
2.3.2	Objective function definition	24
2.3.3	Discrete adjoint method for constructing the gradient	25
2.3.4	Gradient-based optimisation methods	27
2.4	Results and discussion	27
2.4.1	Steady-state operation	27
2.4.2	Induction control in two-dimensional flow	30
2.4.3	Yaw control with wind direction variation	34
2.4.4	Finite horizon effects in economic MPC	37
2.5	Conclusions	39
2.A	Partial derivatives	40
2.A.1	State update	40
2.A.2	Output and objective function	47
2.B	Convergence study for numerical methods	48

This chapter is based on the following publication:

[81] M.J. van den Broek, D. De Tavernier, B. Sanderse, and J.W. van Wingerden, *Adjoint optimisation for wind farm flow control with a free-vortex wake model*, Renewable Energy, <https://doi.org/10.1016/j.renene.2022.10.120>, 2022.

The software developed for this chapter is openly available at <https://doi.org/10.4121/20278620> or on GitHub at <https://github.com/TUDeft-DataDrivenControl/vortexwake> [82]. Additionally, the data presented in this chapter is openly available at <https://doi.org/10.4121/20278590> [83].

2.1 Introduction

Large, densely spaced wind farms are designed and constructed to make use of limited offshore parcels. Within these farms, aerodynamic interaction between wind turbines reduces power production and increases fatigue loading as turbulent, low-energy wakes travel through the farm and negatively affect downstream turbines. Wind farm topology is designed to minimise these interactions, but is inflexible to cope with dynamic, varying atmospheric conditions [15]. The purpose of wind farm control is to minimise the detrimental effects of aerodynamic interaction between wind turbines in a wind farm.

Control strategies for wind farm control can be roughly divided in three categories: wake redirection by yaw misalignment, induction control, and wake mixing strategies [17]. First, the use of yaw misalignment with respect to the free-stream wind direction to redirect wakes downstream has been shown to effectively improve performance under steady conditions in both wind tunnel experiments [33–35] and field studies [36, 38–41]. Second, sinusoidal thrust variations, and consequent induction variations, through collective pitch control have been found to improve wake recovery in an LES study [22] and in wind tunnel experiments [21]. Finally, recent developments in stimulating wake mixing have shown the potential to improve upon collective pitch variations with the helix approach, an individual pitch control strategy [23].

Control-oriented models are often at the core of wind farm control algorithms. Steady-state engineering wake models, such as those that have been implemented in FLORIS [47], are the current industry standard. These include, for example, the Gaussian model [48] or a steady representation of curled wake dynamics [49]. As steady-state wake representations are limited in realistic time-varying conditions, dynamic effects have been added to these engineering wake models to improve upon the steady-state results by including dynamic wake meandering [84] or using Lagrangian particles to incorporate wake dynamics [56].

Several studies have also developed physics-based dynamic models for wind farm flow control, especially using the adjoint method to efficiently calculate gradient information for a scalar objective function with a large number of parameters. The patterns found through optimal control studies with adjoint large-eddy simulations [57, 58] provided the basis for dynamic induction control methods, although these simulations are too computationally expensive for real-time control applications [22, 59]. WFSim provides a 2D Navier-Stokes based wind-farm flow model for control [61], which has then been used for adjoint optimisation of induction control [66]. FRED [64] builds on the results from WFSim to simulate wind farm performance with the adjoint for gradient calculation [62, 63]. However, the 2D physics inherent in this model lack the curled wake dynamics of a wind turbine

under yaw misalignment [26–29] and could not accurately model the effects of wake redirection [63].

In contrast to conventional computational fluid dynamic approaches, free-vortex methods use the vorticity formulation of the Navier-Stokes equations to model wind turbine wakes with Lagrangian elements [68]. Within the field of wind energy, free-vortex wake models have been used to study floating wind turbines and wake dynamics [75] and to study dynamic wake control methods and analyse wake stability [73]. The latter uses the CACTUS code which has been shown to be mostly accurate for near-wake regions [85]. Even though vortex methods are generally more accurate in near-wake regions, a free-vortex ring method has been used to model far wake dynamics for both fixed-bottom and floating wind turbines [74]. Additionally, an actuator-disc model based on discretised vortex rings has been shown to capture the 3D dynamics of the kidney-shaped wake under yaw misalignment [77].

In this chapter, we propose the use of the free-vortex wake method as a computationally efficient, physics-based wake model for control optimisation, especially coupled with the adjoint for efficient evaluation of the gradient. This work aims to extend the possibilities for optimisation of induction and yaw signals for dynamic wind farm flow control. For that purpose, the contribution of this chapter is three-fold: (i) a control-oriented free-vortex wake model of an actuator disc in 2D and 3D with the discrete adjoint for gradient computation, (ii) an economic model-predictive control implementation for dynamic wind farm flow control, and (iii) initial results that demonstrate dynamic induction control and yaw control under time-varying wind direction.

The remainder of the chapter is structured as follows. A 2D and 3D free-vortex model of an actuator disc to represent a wind turbine wake is presented in Section 2.2. The non-linear optimisation problem for economic model-predictive control is formulated in Section 2.3 together with the discrete adjoint method for calculating the gradient. Results are discussed in Section 2.4, which provides an overview of operation under steady conditions followed by receding horizon control optimisation of time-varying axial induction and yaw signals. Finally, conclusions are presented in Section 2.5.

2.2 Control-oriented free-vortex wake model

The general formulation for the control-oriented free-vortex wake (FVW) representation is described in Section 2.2.1. Aspects specific to the 2D and 3D implementations are then defined in Section 2.2.2 and Section 2.2.3, respectively. The convergence and validation of the method for the numerical parameters used in this chapter is provided in 2.B.

2.2.1 General formulation

An actuator-disc representation of a wind turbine is implemented with the free-vortex method in both a two-dimensional (2D) and three-dimensional (3D) formulation. The free-vortex method is based on Lagrangian particles that advect downstream. These particles induce a velocity based on their associated vorticity. The resultant flow velocity may be calculated at any position based on the free-stream velocity and the sum of induced velocities. For a further description of the fundamentals, the reader is referred to aerodynamic literature, such as [68].

The use of the free-vortex wake method requires the assumption of inviscid and incompressible flow. The actuator disc is assumed to be uniformly loaded so it only releases vorticity along its edge [68]. For the 2D model, the wake is modelled by releasing pairs of vortex points at the edge of the actuator disc at every simulation time-step. The 3D code is based on the simulation of discretised vortex rings with vortex filaments, adapted from the model described by Berdowski *et al.* [77]. For convenience, all units have been non-dimensionalised by the rotor diameter and inflow speed.

A system with fixed dimensionality is preferred for control optimisation, therefore the wake models are set up with n_e elements per vortex ring and a fixed number of vortex rings n_r . The number of points to define the vortex elements n_p equals n_e in 2D and $n_e + 1$ in 3D. The spatial dimension of the simulation is n_d , which equals either two or three. The number of turbines modelled is n_t and the number of control parameters per turbine is n_c . For example, the total number of states is $n_s = 2n_r n_p n_d + n_r n_e + n_t n_c$ for a single wake modelled with the FVW, where additional virtual turbines are evaluated using the flow velocity without including their effect on the wake.

We set up the model as a non-linear state-space system in discrete time,

$$\mathbf{q}_{k+1} = f(\mathbf{q}_k, \mathbf{m}_k), \quad (2.1)$$

$$\mathbf{y}_k = g(\mathbf{q}_k, \mathbf{m}_k), \quad (2.2)$$

where for every discrete time step k the updated state $\mathbf{q}_{k+1} \in \mathbb{R}^{n_s}$ and the output vector $\mathbf{y}_k \in \mathbb{R}^{n_t}$ are a function of the current state $\mathbf{q}_k \in \mathbb{R}^{n_s}$ and the control inputs $\mathbf{m}_k \in \mathbb{R}^{n_t n_c}$. The state vector is built up as

$$\mathbf{q} = \begin{bmatrix} X \\ \Gamma \\ U \\ M \end{bmatrix}, \quad (2.3)$$

from the vortex element positions $X \in \mathbb{R}^{n_r n_p n_d}$, the vortex element circulations $\Gamma \in \mathbb{R}^{n_r n_e}$, the stored free-stream velocity $U \in \mathbb{R}^{n_r n_p n_d}$, and the control inputs from

the previous time step $\mathbf{M} \in \mathbb{R}^{n_t n_c}$. The full control vector $\mathbf{m}$ is defined as

$$\mathbf{m} = \begin{bmatrix} a_0 \\ \psi_0 \\ a_1 \\ \psi_1 \end{bmatrix}, \quad (2.4)$$

for a two-turbine configuration with axial induction a and turbine yaw angle ψ .

States corresponding to a ring are indicated with a subscript, rings are indexed with a superscript starting from 0. This allows, for example, the convenient relation of a point $\mathbf{x}_i^{(b)} \in \mathbb{R}^{n_d}$ to the point in the same position in the previous ring $\mathbf{x}_i^{(b-1)}$, or all points in a ring $\mathbf{X}^{(b)} \in \mathbb{R}^{n_p n_d}$ to all points in the previous ring $\mathbf{X}^{(b-1)}$.

For all rings except the first ($b \geq 1$), the position update is calculated from the position of the previous ring with simulation time step h , the stored inflow velocity $\mathbf{u}_\infty \in \mathbb{R}^{n_d}$, and the total induced velocity $\mathbf{u}_{\text{ind}} \in \mathbb{R}^{n_d}$,

$$\mathbf{x}_i^{(b)} \Big|_{k+1} = \mathbf{x}_i^{(b-1)} \Big|_k + h \left(\mathbf{u}_{\infty, i}^{(b-1)} + \mathbf{u}_{\text{ind}}(\mathbf{x}_i^{(b-1)}, \mathbf{q}) \right) \Big|_k. \quad (2.5)$$

The velocity $\mathbf{u}_{\text{ind}}$ induced at any point $\mathbf{x}$ is the sum of the contribution from all vortex elements in the system

$$\mathbf{u}_{\text{ind}}(\mathbf{x}, \mathbf{q}) = \sum_{b=0}^{n_r-1} \sum_{j=1}^{n_e} \mathbf{u}_i^{(b)}, \quad (2.6)$$

where $\mathbf{u}_i \in \mathbb{R}^{n_d}$ is the velocity induced by a single vortex element. The generation of new vortex elements in the first ring $\mathbf{X}^{(0)}$ and the velocity induced by a single vortex element $\mathbf{u}_i$ is defined for 2D and 3D in Section 2.2.2 and Section 2.2.3, respectively.

The vector Γ contains the vortex strength Γ for all elements in all rings. The vortex strength of the first ring is given according to

$$\Gamma_i^{(0)}(\mathbf{q}, \mathbf{m}) = \frac{d\Gamma}{dt} h = c'_t(a) \frac{1}{2} (\mathbf{u}_r \cdot \mathbf{n}(\psi))^2 h \quad \text{for } i = 1, 2, \dots, n_e. \quad (2.7)$$

In this expression, $\mathbf{u}_r$ is the average wind speed at the rotor. The vector $\mathbf{n} \in \mathbb{R}^{n_d}$ is a unit vector orthogonal to the rotor disc, pointing in the downstream direction, with the rotation matrix $\mathbf{R}_z \in \mathbb{R}^{n_d \times n_d}$ and axis-aligned unit vector $\mathbf{e}_x \in \mathbb{R}^{n_d}$,

$$\mathbf{n}(\psi) = \mathbf{R}_z(\psi) \mathbf{e}_x. \quad (2.8)$$

The local thrust coefficient c'_t , is calculated from the axial induction a as

$$c'_t(a) = \begin{cases} \frac{4a(1-a)}{(1-a)^2} = \frac{4a}{1-a} & \text{if } a \leq a_t, \\ \frac{c_{t1} - 4(\sqrt{c_{t1}} - 1)(1-a)}{(1-a)^2} & \text{if } a > a_t, \end{cases} \quad (2.9)$$

where the induction a_t at the transition point is

$$a_t = 1 - \frac{1}{2}\sqrt{c_{t1}}, \quad (2.10)$$

and the parameter $c_{t1} = 2.3$. The thrust coefficient calculation is based on momentum theory with a transition to a linear approximation for high induction values that is an empirical correction based on the Glauert correction [86]. Vortex strength of subsequent rings is inherited downstream,

$$\Gamma_i^{(b)} \Big|_{k+1} = \Gamma_i^{(b-1)} \Big|_k \quad \text{for } i = 1, 2, \dots, n_e \text{ and } b = 1, 2, \dots, n_r - 1. \quad (2.11)$$

Ring zero is initialised at the turbine position with the free-stream velocity, which may vary over space and simulation time,

$$\mathbf{U}^{(0)} \Big|_{k+1} = \mathbf{u}_\infty(\mathbf{X}^{(0)}, k). \quad (2.12)$$

The inflow velocity is then propagated downstream with the state update

$$\mathbf{U}^{(b)} \Big|_{k+1} = \mathbf{U}^{(b-1)} \Big|_k \quad \text{for } b = 1, 2, \dots, n_r - 1. \quad (2.13)$$

The vector $\mathbf{M}$ is an augmentation of the system state to store controls for power calculation at the next time-step,

$$\mathbf{M} \Big|_{k+1} = \mathbf{m}_k. \quad (2.14)$$

This avoids a direct feed-through of control actions to the output function.

The output vector $\mathbf{y}$ contains the power of all turbines as

$$\mathbf{y} = \begin{bmatrix} P_0 \\ P_1 \end{bmatrix}, \quad (2.15)$$

for a two-turbine case. The power P at turbine i is calculated as

$$P_i = \frac{1}{2} c'_p(a) A_r (\mathbf{u}_r \cdot \mathbf{n}(\psi))^3, \quad (2.16)$$

with the local power coefficient c'_p , rotor area A_r , the disc-averaged velocity $\mathbf{u}_r \in \mathbb{R}^{n_d}$, and the yaw angle ψ . The local power coefficient is calculated with the induction factor a as

$$c'_p(a) = \frac{4a(1-a)^2}{(1-a)^3} = \frac{4a}{1-a}. \quad (2.17)$$

For the disc-averaged velocity, we distribute n_u points over a disc representing the turbine according to an equal-area distribution [87] and rotate the disc over the yaw angle. The rotor-disc averaged velocity is then

$$\mathbf{u}_r = \frac{1}{n_u} \sum_{i=1}^{n_u} (\mathbf{u}_\infty(\mathbf{x}_i, \mathbf{q}) + \mathbf{u}_{\text{ind}}(\mathbf{x}_i, \mathbf{q})) , \quad (2.18)$$

where the local free-stream flow is calculated as an average from neighbouring points weighted by distance,

$$\mathbf{u}_\infty(\mathbf{x}, \mathbf{q}) = \sum_{i=0}^{n_p} \sum_{b=0}^{n_r} \tilde{w}_i^{(b)} \mathbf{u}_{\infty,i}^{(b)} , \quad (2.19)$$

with normalised weights $\tilde{w}_i^{(b)}$,

$$w_i^{(b)} = \exp(-10 \|\mathbf{x} - \mathbf{x}_i^{(b)}\|) , \quad (2.20)$$

$$\tilde{w}_i^{(b)} = \frac{w_i^{(b)}}{\sum_{i=0}^{n_p} \sum_{b=0}^{n_r-1} w_i^{(b)}} . \quad (2.21)$$

For the calculation of the power of a virtual turbine – one that does not act on the flow simulation, but is included for the purposes of optimisation – we lower the disc-averaged velocity by the induction factor

$$\mathbf{u}_r^* = (1 - a) \mathbf{u}_r . \quad (2.22)$$

2.2.2 Two-dimensional model specifics

The $n_e = 2$ vortex elements of the first ring are initiated at the edge of the rotor disc with radius r

$$\mathbf{x}_0^{(0)}(\psi) \Big|_{k+1} = \mathbf{R}_z(\psi_k) \begin{bmatrix} 0 \\ r \end{bmatrix} , \quad \mathbf{x}_1^{(0)}(\psi) \Big|_{k+1} = \mathbf{R}_z(\psi_k) \begin{bmatrix} 0 \\ -r \end{bmatrix} , \quad (2.23)$$

where $\mathbf{R}_z(\psi)$ is the rotation matrix for a rotation of an angle ψ around the z -axis,

$$\mathbf{R}_z(\psi) = \begin{bmatrix} \cos \psi & \sin \psi \\ -\sin \psi & \cos \psi \end{bmatrix} . \quad (2.24)$$

The velocity $\mathbf{u}_i$ induced at point $\mathbf{x}_0$ by a single vortex element located at $\mathbf{x}_1$ in 2D is calculated with the Biot-Savart law as

$$\mathbf{u}_i(\mathbf{x}_0, \mathbf{x}_1) = \begin{bmatrix} -r_y \\ r_x \end{bmatrix} \left(\frac{\Gamma}{2\pi} \frac{1}{\|\mathbf{r}\|^2} \right) \left(1 - \exp \left(-\frac{\|\mathbf{r}\|^2}{\sigma^2} \right) \right) , \quad (2.25)$$

where the relative position $\mathbf{r}$ is

$$\mathbf{r} = \mathbf{x}_1 - \mathbf{x}_0 . \quad (2.26)$$

A Gaussian core with core size σ is included to regularise singular behaviour of the induced velocity close to the vortex element.

2.2.3 Three-dimensional model specifics

At every time-step, the vortex filaments that make up a new vortex ring discretised with n_e elements are distributed over a circle with radius r , with yaw angle ψ ,

$$\mathbf{x}_i^{(0)}(\psi) \Big|_{k+1} = \mathbf{R}_z(\psi_k) \begin{bmatrix} 0 \\ r \cos(2\pi \frac{i}{n_e}) \\ r \sin(2\pi \frac{i}{n_e}) \end{bmatrix} \quad \text{for } i = 0, 1, \dots, n_e , \quad (2.27)$$

where $\mathbf{R}_z(\psi)$ is the rotation matrix for a rotation of an angle ψ around the z -axis,

$$\mathbf{R}_z(\psi) = \begin{bmatrix} \cos \psi & \sin \psi & 0 \\ -\sin \psi & \cos \psi & 0 \\ 0 & 0 & 1 \end{bmatrix} . \quad (2.28)$$

The induced velocity $\mathbf{u}_i$ at a point $\mathbf{x}_0$ is calculated with the Biot-Savart law from a single vortex element starting at $\mathbf{x}_1$ and ending at $\mathbf{x}_2$, with vortex strength Γ ,

$$\mathbf{u}_i(\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2) = \left(\frac{\Gamma}{4\pi} \frac{\mathbf{r}_1 \times \mathbf{r}_2}{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2} \right) \left(\mathbf{r}_0 \cdot \left(\frac{\mathbf{r}_1}{\|\mathbf{r}_1\|} - \frac{\mathbf{r}_2}{\|\mathbf{r}_2\|} \right) \right) \left(1 - \exp \left(-\frac{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2}{\sigma^2 \|\mathbf{r}_0\|^2} \right) \right) , \quad (2.29)$$

where the relative positions $\mathbf{r}$ are defined as

$$\mathbf{r}_0 = \mathbf{x}_2 - \mathbf{x}_1 , \quad (2.30)$$

$$\mathbf{r}_1 = \mathbf{x}_1 - \mathbf{x}_0 , \quad (2.31)$$

$$\mathbf{r}_2 = \mathbf{x}_2 - \mathbf{x}_0 . \quad (2.32)$$

A Gaussian core with core size σ is included to regularise singular behaviour of the induced velocity close to the vortex filament.

2.3 Optimisation for power maximisation

The free-vortex wake model described in the previous section is implemented as a novel control-oriented model for dynamic wind farm flow control. Wind turbine power maximisation is introduced in Section 2.3.1 in an economic model-predictive control setting. The associated non-linear optimisation problem is formulated in Section 2.3.2. The derivation of the discrete adjoint for calculation of the gradient is described in Section 2.3.3, followed by the choice of a gradient-based optimisation method to solve the non-linear problem in Section 2.3.4.

2.3.1 Economic model-predictive control

The conventional model-predictive control (MPC) approach is a model-based optimisation of control signals to drive an objective functional to zero, for example for optimal tracking of a reference signal. However, for maximisation of wind farm power production, the optimal objective value is not known a priori, leading to an economic problem formulation. The economic MPC (EMPC) approach considers optimisation of an objective to an unknown extremum. This optimisation problem is conventionally solved in a receding horizon setting with a finite prediction horizon. After optimisation, the first (set of) control(s) is implemented and the problem is shifted and solved again up to the new horizon [78].

One problem with optimisation to a finite horizon is that the optimisation considers the prediction horizon as the end of time. Therefore, control actions that prioritise gain within the horizon may be optimal, although they would have undesired consequences post-horizon. This is known as the turnpike effect [88], where a solution stays close to the optimal trajectory for most of the window but diverges towards the horizon. These finite horizon effects may be treated by terminal constraints or terminal conditions [78]. For example, the control signal has been kept constant towards the end of the horizon to limit undesired effects in wind farm control [66] or a terminal condition on rotor kinetic energy has been used to regularise optimisation results for wind turbine control [89]. Given a sufficiently long prediction horizon, EMPC has been shown to also converge without terminal constraints [90].

In this chapter, the turnpike effects are treated by considering sufficiently long prediction horizons within the receding horizon setting, so as not to require terminal constraints. The control problem is formulated in a non-linear EMPC setting without terminal conditions with the goal of maximising mean power production over time.

2.3.2 Objective function definition

A non-linear minimisation problem with a scalar objective function J is constructed to find the set of optimal controls $\mathbf{m}_{k_0+i} \in \mathbb{R}^{n_m}$, with $n_m \leq n_t n_c$ the number of free

controls and $i = 0, 1, \dots, N_h$. The objective is the total power output over the prediction horizon of N_h steps from the current step k_0 ,

$$\min_{\mathbf{m}_k} \sum_{k=k_0}^{k_0+N_h} J(\mathbf{q}_k, \mathbf{m}_k) = \min_{\mathbf{m}_k} \sum_{k=k_0}^{k_0+N_h} Qy_k(\mathbf{q}_k, \mathbf{m}_k) + \Delta \mathbf{m}_k^T \mathbf{R} \Delta \mathbf{m}_k, \quad (2.33)$$

where y_k contains the power of modelled and virtual turbines, $\Delta \mathbf{m}_k = \mathbf{m}_k - \mathbf{m}_{k-1}$ is the change in control value between time steps, and $Q \in \mathbb{R}^{1 \times n_t}$ and $R \in \mathbb{R}^{n_m \times n_m}$ are weights to balance power output and actuation cost. A linear sum of power is chosen for mean power maximisation because power is already a positive objective function. A quadratic functional would more heavily weight peaks in power production and be suboptimal for maximisation of mean power. The output weight is chosen negative ($Q < 0$) so that power is maximised for minimisation of the objective. The input weight R functions as a regularisation term and aids convergence to suitable control solutions by smoothing the optimisation landscape.

2.3.3 Discrete adjoint method for constructing the gradient

The gradient of the objective function is calculated following the discrete adjoint method [91] because the method scales well for a large number of input sensitivities. We take the non-linear state-space system in (2.1) and define the objective function $J_k = J(\mathbf{q}_k, \mathbf{m}_k)$ at time-step k , such that the total objective function J_{total} is accumulated over a number of steps N_h ,

$$J_{\text{total}} = J_{N_h} + \sum_{i=0}^{N_h-1} J_i, \quad (2.34)$$

where $i = 0$ at the current time-step $k = k_0$. This is the total objective function to be minimised in the optimisation problem in (2.33).

To derive the adjoint system, we extend the objective function with adjoint states and system constraint,

$$\tilde{J}_{\text{total}} = J_{N_h} + \sum_{i=0}^{N_h-1} (J_i + \boldsymbol{\lambda}_{i+1}^T (f_i - \mathbf{q}_{i+1})), \quad (2.35)$$

where the adjoint states $\boldsymbol{\lambda}$ can be chosen freely because $f_i - \mathbf{q}_{i+1} = 0$. Since $J_k = J(\mathbf{q}_k, \mathbf{m}_k)$, a differential change $\delta \tilde{J}_{\text{total}}$ can be expanded in terms of changes in $\mathbf{q}$

and $\mathbf{m}$ as:

$$\begin{aligned} \delta \bar{J}_{\text{total}} = & \left(\frac{\partial J_{N_h}}{\partial \mathbf{q}_{N_h}} - \boldsymbol{\lambda}_{N_h}^T \right) \delta \mathbf{q}_{N_h} + \frac{\partial J_{N_h}}{\partial \mathbf{m}_{N_h}} \delta \mathbf{m}_{N_h} \\ & + \sum_{i=0}^{N_h-1} \left(\left(\frac{\partial J_i}{\partial \mathbf{q}_i} + \boldsymbol{\lambda}_{i+1}^T \frac{\partial f_i}{\partial \mathbf{q}_i} - \boldsymbol{\lambda}_i^T \right) \delta \mathbf{q}_i + \left(\frac{\partial J_i}{\partial \mathbf{m}_i} + \boldsymbol{\lambda}_{i+1}^T \frac{\partial f_i}{\partial \mathbf{m}_i} \right) \delta \mathbf{m}_i \right), \end{aligned} \quad (2.36)$$

We then choose the adjoint states to be

$$\boldsymbol{\lambda}_{N_h}^T = \frac{\partial J_{N_h}}{\partial \mathbf{q}_{N_h}}, \quad \boldsymbol{\lambda}_i^T = \frac{\partial J_i}{\partial \mathbf{q}_i} + \boldsymbol{\lambda}_{i+1}^T \frac{\partial f_i}{\partial \mathbf{q}_i}, \quad \boldsymbol{\lambda}_0^T = \mathbf{0}, \quad (2.37)$$

such that the variations due to $\mathbf{q}$ in (2.36) are cancelled out. The adjoint states are solved for by propagation backwards in time, starting from the final adjoint state. The gradient of the objective function parts J_k to the input can then be calculated from these adjoint states

$$\nabla J_{N_h} = \frac{\partial J_{N_h}}{\partial \mathbf{m}_{N_h}}, \quad \nabla J_i = \frac{\partial J_i}{\partial \mathbf{m}_i} + \left(\frac{\partial f_i}{\partial \mathbf{m}_i} \right)^T \boldsymbol{\lambda}_{i+1}. \quad (2.38)$$

The total gradient $\nabla \bar{J}_{\text{total}}$ with respect to all control parameters $\mathbf{m}_i$ is then constructed as

$$\nabla \bar{J}_{\text{total}} = \sum_{i=0}^{N_h} \nabla J_i. \quad (2.39)$$

The partial derivatives of the state update and output function with respect to the full model state and controls are stored in memory during the forward simulation of the model. These partial derivatives are provided for the given model and objective function in 2.A.

The evaluation of the gradient thus requires a single forward simulation with evaluation of the partial derivatives and a single backward pass to solve for the adjoint states and construct the gradient. In that sense, this method of gradient evaluation is considerably more efficient than finite difference methods as the computational cost of the discrete adjoint increases only minimally with the number of control parameters for which the derivative is required. The computational cost of gradient evaluation with the discrete adjoint primarily scales with the expense of the forward simulation and the associated partial derivatives.

2.3.4 Gradient-based optimisation methods

The availability of the gradient allows the use of gradient-based optimisation techniques for control optimisation. Exploration of the objective function shows that it is non-linear and non-convex, with almost flat regions and numerous local minima. Initial experiments were run with L-BFGS-B optimisation [92] as was also used in the work by Munters and Meyers [59]. However, this optimiser appeared sensitive to initialisation at local maxima and to convergence to local minima.

The Adam optimiser [93] is a gradient-based method often used in machine learning for optimisation of neural network weights, where it is applied for gradient descent with noisy gradients in complex optimisation landscapes. It uses a momentum approach to accelerate gradient descent and has proven to be less sensitive to the choice of initial guess and local minima. Within this work, we use the Adam optimiser with the default parameters; a maximum step size $\alpha = 0.001$ and the default decay rates $\beta_1 = 0.9$ and $\beta_2 = 0.999$. Tuning of these parameters may still improve performance. The yaw angle is on a different order of magnitude than axial induction. Therefore, it is scaled by a factor 10^{-2} in the optimisation, so that the step size covers a similar range of the allowable range of induction value and yaw angle.

2.4 Results and discussion

A brief overview of the 2D and 3D FVW under steady conditions is given in Section 2.4.1 to illustrate the test case configuration and provide a steady baseline for control performance. This is followed by two example cases to demonstrate the use of the FVW as a novel dynamic model for control optimisation in the receding horizon setting described previously; a 2D case with induction control is provided in Section 2.4.2 and a 3D case for yaw control under time-varying wind direction in Section 2.4.3. Finally, Section 2.4.4 discusses finite horizon effects in EMPC for wind farm control.

2.4.1 Steady-state operation

We define a two-turbine case for evaluating control optimisation with the FVW, starting with steady-state control characteristics. The two turbines are spaced $5D$ apart, where D is the rotor diameter, aligned with the uniform unit inflow. The upstream turbine is modelled with the FVW and the virtual downstream turbine performance is evaluated using the flow velocity over the rotor area at the downstream position. The parameters for the FVW are provided in Table 2.1. An exploration of parameter sensitivity is supplied in 2.B.

Figure 2.1 shows a 2D FVW simulation with an induction factor $a = 0.33$ and without yaw misalignment. The pairs of vortex points provide the basis for the

Table 2.1: FVW parameters for 2D and 3D case

		2D	3D
time-step	h	0.2	0.3
core size	σ	0.1	0.16
number of rings	n_r	60	40
elements per ring	n_e	2	16




Figure 2.1: Illustration of the 2D FVW for uniform unit inflow without yaw misalignment and with a constant induction factor $a = 0.33$. The pairs of vortex points (top) can be used to calculate the velocity at any point, allowing visualisation of a dense velocity field (bottom). The figure illustrates the two-turbine case where the second turbine performance is calculated from the flow velocity 5 D downstream.

simulation and allow calculation of a dense velocity field. Disturbances in the far wake are the result of numerical instabilities. It is also notable that the wake is quite wide as is expected for planar flow.

The 3D FVW produces a vortex ring structure as shown in Figure 2.2 for a simulation with yaw heading of $\psi = 30^\circ$ and induction factor $a = 0.33$. The figure shows the dense velocity field with the wake deficit calculated from the skeleton of vortex filaments. A kidney-shaped wake appears from the pair of counter-rotating vortices that are generated by a turbine operating under yaw misalignment, as shown in [77].

The model response to control signal variation is verified by examining power production in steady state. First, Figure 2.3 shows the 2D and 3D FVW power curve for a variation in axial induction from $a = 0$ to $a = 0.5$. The maximum individual




Figure 2.2: The 3D FVW models the wake from a series of vortex rings discretised into vortex elements (top), allowing calculation of a velocity field showing the wake deficit (bottom). The kidney-shaped wake appears as a pair of counter-rotating vortices is formed under yaw misalignment. Simulation under uniform inflow with a yaw heading of $\psi = 30^\circ$ and induction $a = 0.33$. The figure illustrates the case where the upstream turbine is modelled with the 3D FVW and the downstream turbine performance is calculated from the flow velocity over a rotor disc 5 D downstream.

turbine power matches the expectation from momentum theory for the chosen parameters at the theoretical optimum induction of $a = 0.33$. Steady under-induction provides a power gain of 3.6 % over greedy control. The 2D and 3D FVW show remarkably similar behaviour in terms of power production for varying induction factor on the upstream turbine. The similarity in the power estimate shows an opportunity for doing induction control in 2D model studies. Additionally, 2D wind farm flow models have already been used for studies of induction control in a wind farm setting [65, 66].

Second, a yaw sweep from $\psi = -45^\circ$ to $\psi = 45^\circ$ is illustrated in Figure 2.4. This steady sweep shows a demonstrable lack of power gain from yaw misalignment in the 2D FVW. However, in 3D, yaw misalignment does lead to wake redirection




Figure 2.3: Power production in steady state for varying induction on the upstream turbine, in 2D (left) and 3D (right), with turbine configuration as in Figures 2.1 and 2.2, respectively. Total power is the sum of the powers from turbine 0 and turbine 1. Maximum greedy power production occurs for $a = 0.33$. Within this model, lowering the induction on the upstream turbine to $a = 0.27$ leads to a 3.6 % gain in total power.

and the maximum power achieved in steady state is 0.313 for a yaw angle of 34° , providing a gain of 26.1 % over greedy control.

The 3D FVW shows the formation of a kidney-shaped wake from a counter-rotating vortex pair when the turbine is operated under yaw misalignment. The subsequent deflection of the turbine wake leads to an increase of the combined power production. These dynamics are not present in 2D, which may explain the lack of wake redirection. This supports previous results that found 2D flow modelling ineffective in capturing the essential effects of wake steering [63].

The model is currently symmetric, which means there is no difference between positive or negative yaw misalignment. Experimental studies have found wake steering to be asymmetric due to the rotation induced by the rotor [27–29]. A normal actuator disc was chosen for simplicity, but a root vortex could be included to model the turbine as a rotating actuator disc to model the asymmetric aspect of wake redirection.

2.4.2 Induction control in two-dimensional flow

Given the similarity between 2D and 3D in power curves for variation of axial induction, an optimisation case for induction control is set up in 2D with a configuration as in Figure 2.1 and parameters as in Table 2.1. Both turbines are aligned with the inflow wind direction and set to a $\psi = 0^\circ$ yaw angle. The downstream turbine is assumed to be performing at its greedy optimum with an induction factor $a = 0.33$, whereas the induction control signal of the upstream turbine is to be found by solving the optimisation problem.




Figure 2.4: Power production in steady state for varying yaw on the upstream turbine, in 2D (left) and 3D (right), with turbine configuration as in Figures 2.1 and 2.2, respectively. The downstream turbine power is calculated from the flow velocity 5 D downstream from the first turbine. Total power is the sum of the powers from turbine 0 and turbine 1. The 2D FVW does not include the dynamics to model wake steering effectively. In the 3D FVW, a 26.1 % gain in total power is observed for a yaw angle of $\psi = 34^\circ$ compared to greedy control where $\psi = 0^\circ$.

The objective function (2.33) is constructed with the control signal $\mathbf{m}_k = [a_k]$, over a prediction horizon of $N_h = 100$ samples. The output contains the power of both turbines, $\mathbf{y}_k = [p_0 \ p_1]^T$. The objective function weights are set to $\mathbf{Q} = [-1 \ -1]$ and $\mathbf{R} = [10]$. This choice of input weight resulted in an adequate balance between input action and power production in an exploratory parameter sweep. At every time step, the optimisation is run for 50 iterations, after which the first value of the control signal from the optimisation is implemented in the receding horizon control scheme. Further iterations have not lead to consistently better performance in terms of objective function value given the current optimiser configuration. The starting state for the optimisation case is the result of a steady simulation under greedy operating conditions to remove transient effects. This initial condition is shown in the first frame of Figure 2.6.

The control signal produced in this economic MPC framework is illustrated in Figure 2.5. The optimisation converges to a roughly periodic excitation with a dominant frequency of approximately $f = 0.20$. The signal features sharp downward peaks where the induction is lowered, thus reducing thrust and allowing more flow to pass through the rotor disc. In addition to the periodic excitation, the mean induction factor is lowered to $\bar{a} = 0.30$ below the greedy optimum of $a = 0.33$. The mean power produced in the final two-thirds of the simulation ($t > 20$) is 0.283, which is an increase of 6.0 % over the maximum power achieved with steady induction control.

A series of snapshots of the wake produced with this control signal is shown




Figure 2.5: Receding horizon optimisation of the axial induction control signal for total power yields dynamic behaviour that stimulates wake breakdown and improves time-average power production by 6.0% over steady under-induction. Control signal for turbine 0 is optimised whilst turbine 1 is virtually modelled to be operating at its greedy optimum, positioned 5 D downstream in fully waked conditions. Total power is the sum of power from turbine 0 and turbine 1. Snapshots of the flow field at times t_0 to t_7 are illustrated in Figure 2.6.

in Figure 2.6. The effects of periodic induction excitation are apparent in the flow field of the wake as coherent structures are formed that travel downstream.

We observe that optimisation with the 2D FVW leads to induction control signals that combine static under-induction with a strong dynamic component. Within the current model, this combination of periodic excitation and lowering of the mean induction factor outperforms a simple steady induction decrease in terms of mean power production. The sharp downward peaks in induction signal appear to stimulate breakdown of the wake and mixing with the free-stream flow. Note that mixing here does not refer to turbulent mixing as no turbulence is present in the FVW model.

The use of purely static induction control was previously shown not to be a very effective solution for improving wind farm power production [18, 20]. These results are supported as, within the 2D FVW, the use of under-induction on its own is less effective than the dynamic induction signal acquired through optimisation. Further study will need to find out whether the combination of slight under-induction and periodic excitation is effective in a realistic wind farm scenario.

The periodic aspect of this control signal resembles the sinusoidal thrust sig-




Figure 2.6: Snapshots of the flow field from the 2D free-vortex wake simulation under uniform inflow in positive x direction with the induction signal applied to the turbine as shown in Figure 2.5. The effects of the periodic induction excitation can be seen in the structure of the wake and appear to enhance wake breakdown.

nals that Munters and Meyers [22] found to improve wake mixing by stimulating shedding of vortex rings from the wind turbine. They find sinusoidal actuation at a non-dimensionalised frequency of $f = 0.25$ with a mean local thrust coefficient $c'_t = 2.0$ and amplitude $A = 1.5$ to be optimal for turbines operating in a small farm at 5 D spacing. The signal found in the current work has a slightly lower frequency at $f = 0.20$, and the induction signal corresponds to a lower mean thrust coefficient $c'_t = 1.75$ with amplitude $A = 0.87$. Their optimisation in a 3D LES environment with turbulent inflow is considerably more complex and more expensive than the 2D FVW, which runs well on a regular laptop. It is interesting to note that both studies consider a non-rotating actuator-disc wind turbine model. The differences between the two signals are worth exploring further and will be investigated in future work.

2.4.3 Yaw control with wind direction variation

Optimisation for yaw control requires the 3D FVW model because it captures the dynamics of the curled wake and therefore shows demonstrable power gain from wake steering, as shown in Figure 2.4. The set-up for the optimisation case is as illustrated in Figure 2.2 with the parameters listed in Table 2.1. A smooth wind direction change with unit velocity magnitude from 0° to -20° is implemented to test the capabilities for yaw control under time-varying conditions. The downstream turbine is assumed to be performing at its greedy optimum with an induction factor $a = 0.33$ and a yaw angle that perfectly tracks the inflow direction. The yaw control signal of the upstream turbine is found as the solution of the optimisation problem.

The objective function (2.33) is constructed with the control signal $\mathbf{m}_k = [\psi_k]$, over a prediction horizon of $N_h = 60$ samples. The output contains the power of both turbines, $\mathbf{y}_k = [p_0 \ p_1]^T$. The objective function weights are set to $\mathbf{Q} = [-1 \ -1]$ and $\mathbf{R} = [0.025]$. This choice of input weight resulted in an adequate balance between input action and power production in an exploratory parameter sweep. It differs from the 2D case because the yaw control signal has a different magnitude than the induction signal. At every time step, the optimisation is run for 10 iterations, after which the first value of the control signal is implemented in the receding horizon control scheme. Given the slower variations in yaw angle compared to induction control, fewer iterations were required before further iterations no longer yielded consistent improvement in objective function with the current optimiser. The starting state for the yaw optimisation case is the result of a steady simulation with a 30° yaw heading on the upstream turbine to reach steady conditions with wake redirection. This initial condition is illustrated in the first frame of Figure 2.8.

The control signal implemented in the receding horizon control is shown in




Figure 2.7: Optimisation of the yaw control signal for total power maximisation finds a solution that utilises yaw misalignment to redirect the wake away from the downstream turbine. As the wind direction changes, the turbine is aligned with the flow to perform at its greedy optimum as the wake no longer impinges on the downstream turbine. Control signal for turbine 0 is optimised whilst turbine 1 is virtually modelled to be operating at its greedy optimum, positioned 5 D downstream in fully waked conditions. Total power is the sum of power from turbine 0 and turbine 1. Snapshots of the flow field at times t_0 to t_5 are illustrated in Figure 2.8.

Figure 2.7 together with the inflow and the associated power production for both turbines. The result shows wake redirection through yaw misalignment is maintained for the first section where the wind direction has not yet changed. The turbine is slowly aligned with the inflow in anticipation of the wind direction change. As the wind direction changes, the upstream turbine is rotated into the wind until it is aligned with the free-stream inflow direction. The wake no longer interacts with the downstream turbine which makes its greedy optimum a good control solution.

A series of snapshots of the flow field averaged over the rotor height are shown in Figure 2.8. The snapshots illustrate how the change in wind direction propagates through the wake. It is visible that the turbine control transitions from wake steering with yaw misalignment to greedy control and alignment with the new wind direction.

The use of the 3D FVW as a novel dynamic model for online control optimisation adds to previous work under time-varying wind directions, where Campagnolo *et al.* [35] applied pre-optimised set-points in wind tunnel experiments




Figure 2.8: A series of snapshots of the flow field averaged over rotor height for 3D free-vortex wake simulation under a changing inflow direction. The black arrow indicates the wind direction and the yaw signal applied to the turbine is shown in Figure 2.7. This case illustrates a transition from wake steering to greedy control as the optimal operating point changes with inflow variation.

and Doekemeijer *et al.* [50] used FLORIS to generate steady-state optimal yaw set-points in an online closed-loop controller. Howland *et al.* [79] presented dynamic yaw control using another, more simplified, physics-based model – a lifting line model with a Gaussian wake. They operated under unsteady inflow, but with an invariant mean wind direction. Similarly, unsteady flow without direction changes was considered in the model-free yaw control work by Ciri *et al.* [94].

Especially for yaw control, the initial guess is critical for attaining good results with the optimisation algorithm. The optimiser struggles to find good solutions with an initial guess at zero misalignment from a configuration with greedy yaw control and two turbines with full wake interaction. This seems to be the result of a rather flat optimisation landscape in that configuration. Given some misalignment in the initial guess for the control signal, the optimiser will tend to find a wake steering solution. However, when the turbine is currently misaligned to one side, while the other is more effective, the controller is unlikely to switch because the gradient-based optimisation does not cover a large enough search space. A multi-start optimisation may be a solution to avoid having to predetermine which side to initialise.

Further work is required to validate the effectiveness of the dynamic yaw control results under realistic conditions and investigate whether these solutions improve upon wind farm control strategies with steady control-oriented models. Additionally, a combination of yaw-based control and over-induction has been shown to improve wake steering results [95], which could be further explored in optimisation with the 3D FVW.

2.4.4 Finite horizon effects in economic MPC

In addition to the EMPC results in Sections 2.4.2 and 2.4.3, we illustrate the intermediate results that show the effects of optimisation on a finite horizon. Figure 2.9 illustrates the optimisation result for a single optimisation window for both the induction control and the yaw control case. It illustrates the executed signal, the result of optimisation with a finite horizon at the current time step, and the actual future control signal that is applied in the receding horizon setting. Towards the finite horizon, the induction returns towards the greedy optimum as the effects no longer reach the downstream turbine. Right at the horizon, a final peak in induction occurs as the optimiser tries to generate a little more power with over-inductive behaviour. The effect of optimisation on a finite horizon for yaw control means that the turbine will rotate back to alignment with the inflow when the wake effects no longer propagate to the downstream turbine within the optimisation window.

For the cases presented in this chapter, the horizon length is long enough that these end-of-window effects do not affect the control solution as the receding hori-




Figure 2.9: The control signal for the finite horizon is found through optimisation at every time-step. The executed signal is the past control input leading up to the current state. An optimisation result is found up to the prediction horizon. The first step of the optimised signal is implemented, after which it is shifted and re-used at the next time-step to start re-optimisation. The future signal is the actual control input that ends up being applied in this receding horizon approach.

zon control scheme prevents these finite horizon effects from being implemented. However, especially in optimisation cases where the window is relatively short and input cost is high, the horizon effect needs to be properly treated.

Horizon length for single wind turbine optimisation should be long enough to push turnpike effects away from the signal to be executed. In a multiple-turbine setting, horizon length should be chosen at least long enough such that the effects of control signal variation are observed at the downstream turbine for long enough such that the optimiser finds a balanced control solution. Shorter horizons converge to greedy solutions because the effects do not propagate to downstream turbines within the optimisation window.

2.5 Conclusions

This work presents a control-oriented free-vortex wake (FVW) model of an actuator disc in 2D and 3D to represent a wind turbine wake. A main novelty in our work is the derivation of the discrete adjoint equations associated with the FVW model, which allows efficient gradient evaluation for use in gradient-based optimisation methods. The FVW model is computationally efficient enough that the experiments in this chapter could be run on a regular laptop, without requiring high performance computing clusters. The evaluation of the gradient requires only a single forward simulation and backward integration of the adjoint states, which is on the order of ten times slower than a simple simulation in the current implementation due to the calculation and storage of all partial derivatives.

The FVW model is implemented as a novel dynamic model for gradient-based control optimisation in an economic model-predictive control setting for maximising mean power production by reducing the negative effects of wind turbine wake interaction. This implementation allows generation of optimal control solutions and exploration of dynamical wake behaviour.

In a 2D simulation with receding horizon optimisation, dynamic induction control signals are found that combine slight under-induction with a roughly periodic excitation at a frequency of $f = 0.20$. This results in a 6.0 % gain over the maximum power generated with optimised steady induction control. The FVW provides a new, efficient model for exploring dynamic induction control, which has previously been studied in comparatively expensive LES studies.

The 3D free-vortex wake model exhibits curled wake dynamics under yaw misalignment and is therefore suitable as a novel physics-based control-oriented model for dynamic wake steering control. The economic model-predictive control strategy finds yaw signals under time-varying wind direction that show both wake redirection and a smooth return to greedy control. This adds to the existing literature that bases yaw control on steady model results or considers dynamic control with unsteady flow, but invariant mean wind directions.

Future work will include validation of the model response, exploration of the dynamics of the optimisation problem, and further experiments with the optimisation of induction control and wake redirection.

2.A Partial derivatives

This section presents the partial derivatives required for the calculation of the gradient with the discrete adjoint. Derivatives of the model state update from Section 2.2 are provided in 2.A.1 and those of the output and objective function from Section 2.3 are provided in 2.A.2.

2.A.1 State update

The Jacobian matrix of the state update with respect to the previous state can be written as

$$\frac{\partial f(\mathbf{q}_k, \mathbf{m}_k)}{\partial \mathbf{q}_k} = \frac{\partial \mathbf{q}_{k+1}}{\partial \mathbf{q}_k} = \frac{\partial}{\partial \mathbf{q}_k} \begin{bmatrix} X \\ \Gamma \\ U \\ M \end{bmatrix}_{k+1} = \begin{bmatrix} \frac{\partial X_{k+1}}{\partial X_k} & \frac{\partial X_{k+1}}{\partial \Gamma_k} & \frac{\partial X_{k+1}}{\partial U_k} & 0 \\ \frac{\partial \Gamma_{k+1}}{\partial X_k} & \frac{\partial \Gamma_{k+1}}{\partial \Gamma_k} & \frac{\partial \Gamma_{k+1}}{\partial U_k} & 0 \\ 0 & 0 & \frac{\partial U_{k+1}}{\partial U_k} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad (2.40)$$

where those derivatives that are zero have been removed, leaving a number of sub-matrices to be constructed.

First, take the partial derivatives of the induced velocity of a single vortex element $\mathbf{u}_i$ in 2D, which is divided in three parts (u_0 , u_1 and u_2) to simplify calculation,

$$\mathbf{u}_i(\mathbf{x}_0, \mathbf{x}_1) = \underbrace{\begin{bmatrix} -r_y \\ r_x \end{bmatrix}}_{u_0} \underbrace{\left(\frac{\Gamma}{2\pi} \frac{1}{\|\mathbf{r}\|^2} \right)}_{u_1} \underbrace{\left(1 - \exp\left(-\frac{\|\mathbf{r}\|^2}{\sigma^2}\right) \right)}_{u_2}. \quad (2.41)$$

For clarity, this section will refer to $\mathbf{u}_i$ as $\mathbf{u}$.

The derivative of induced velocity with respect to positions is

$$\frac{\partial \mathbf{u}}{\partial \mathbf{x}_0} = -\frac{\partial \mathbf{u}}{\partial \mathbf{r}}, \quad \frac{\partial \mathbf{u}}{\partial \mathbf{x}_1} = \frac{\partial \mathbf{u}}{\partial \mathbf{r}}, \quad (2.42)$$

where the derivative to the relative position can be expanded as

$$\frac{\partial \mathbf{u}}{\partial \mathbf{r}} = \frac{\partial \mathbf{u}_0}{\partial \mathbf{r}} u_1 u_2 + \mathbf{u}_0 \frac{\partial u_1}{\partial \mathbf{r}} u_2 + \mathbf{u}_0 u_1 \frac{\partial u_2}{\partial \mathbf{r}}. \quad (2.43)$$

The required partial derivatives are

$$\frac{\partial \mathbf{u}_0}{\partial \mathbf{r}} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}, \quad (2.44)$$

$$\frac{\partial u_1}{\partial \mathbf{r}} = -\frac{\Gamma}{\pi} \left(\frac{\mathbf{r}^T}{\|\mathbf{r}\|^4} \right), \quad (2.45)$$

$$\frac{\partial u_2}{\partial \mathbf{r}} = \frac{2\mathbf{r}^T}{\sigma^2} \exp\left(-\frac{\|\mathbf{r}\|^2}{\sigma^2}\right). \quad (2.46)$$

For all rings $b > 0$, the Jacobian matrix from element positions to positions is filled with the partial derivatives as

$$\left[\frac{\partial X^{(b)}}{\partial X^{(a)}} \right]_{ij} = h \frac{\partial \mathbf{u}_i^{(b)}}{\partial \mathbf{x}_j^{(a)}} \left(\mathbf{x}_i^{(b)}, \mathbf{x}_j^{(a)} \right) \quad (2.47)$$

for $i, j = 0, 1$, $a = 1, 2, \dots, n_r$ and $b = 1, 2, \dots, n_r - 1$.

If $a + 1 = b$, then the diagonal elements are substituted as

$$\left[\frac{\partial X_k^{(b)}}{\partial X_k^{(b-1)}} \right]_{ii} \leftarrow 1 + h \sum_{m=1}^{n_r-1} \sum_{n=0}^1 \frac{\partial \mathbf{u}_i^{(b-1)}}{\partial \mathbf{x}_j^{(b-1)}} \left(\mathbf{x}_i^{(b-1)}, \mathbf{x}_n^{(m)} \right). \quad (2.48)$$

The induced velocity $\mathbf{u}_i$ in 3D is similarly divided into three parts $\mathbf{u}_0$, $\mathbf{u}_1$, and $\mathbf{u}_2$,

$$\mathbf{u}_i(\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2) = \underbrace{\left(\frac{\Gamma}{4\pi} \frac{\mathbf{r}_1 \times \mathbf{r}_2}{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2} \right)}_{\mathbf{u}_0} \underbrace{\left(\mathbf{r}_0 \cdot \left(\frac{\mathbf{r}_1}{\|\mathbf{r}_1\|} - \frac{\mathbf{r}_2}{\|\mathbf{r}_2\|} \right) \right)}_{\mathbf{u}_1} \underbrace{\left(1 - \exp \left(-\frac{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2}{\sigma^2 \|\mathbf{r}_0\|^2} \right) \right)}_{\mathbf{u}_2}. \quad (2.49)$$

The required full derivatives of induced velocity with respect to the position inputs can be expanded as

$$\frac{\partial \mathbf{u}}{\partial \mathbf{x}_0} = -\frac{\partial \mathbf{u}}{\partial \mathbf{r}_1} - \frac{\partial \mathbf{u}}{\partial \mathbf{r}_2}, \quad (2.50)$$

$$\frac{\partial \mathbf{u}}{\partial \mathbf{x}_1} = \frac{\partial \mathbf{u}}{\partial \mathbf{r}_1} - \frac{\partial \mathbf{u}}{\partial \mathbf{r}_0}, \quad (2.51)$$

$$\frac{\partial \mathbf{u}}{\partial \mathbf{x}_2} = \frac{\partial \mathbf{u}}{\partial \mathbf{r}_2} + \frac{\partial \mathbf{u}}{\partial \mathbf{r}_0}. \quad (2.52)$$

We expand the required derivatives according to the product rule into manageable parts,

$$\frac{\partial \mathbf{u}}{\partial \mathbf{r}_0} = \frac{\partial \mathbf{u}_0}{\partial \mathbf{r}_0} u_1 u_2 + \mathbf{u}_0 \frac{\partial u_1}{\partial \mathbf{r}_0} u_2 + \mathbf{u}_0 u_1 \frac{\partial u_2}{\partial \mathbf{r}_0}, \quad (2.53)$$

$$\frac{\partial \mathbf{u}}{\partial \mathbf{r}_1} = \frac{\partial \mathbf{u}_0}{\partial \mathbf{r}_1} u_1 u_2 + \mathbf{u}_0 \frac{\partial u_1}{\partial \mathbf{r}_1} u_2 + \mathbf{u}_0 u_1 \frac{\partial u_2}{\partial \mathbf{r}_1}, \quad (2.54)$$

$$\frac{\partial \mathbf{u}}{\partial \mathbf{r}_2} = \frac{\partial \mathbf{u}_0}{\partial \mathbf{r}_2} u_1 u_2 + \mathbf{u}_0 \frac{\partial u_1}{\partial \mathbf{r}_2} u_2 + \mathbf{u}_0 u_1 \frac{\partial u_2}{\partial \mathbf{r}_2}. \quad (2.55)$$

The partial derivatives of $\mathbf{u}_0$ are

$$\frac{\partial \mathbf{u}_0}{\partial \mathbf{r}_0} = \mathbf{0}, \quad (2.56)$$

$$\frac{\partial \mathbf{u}_0}{\partial \mathbf{r}_1} = \frac{\Gamma}{4\pi} \left(\frac{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2 \mathbf{I}_3 - 2(\mathbf{r}_1 \times \mathbf{r}_2)(\mathbf{r}_1 \times \mathbf{r}_2)^T}{\|\mathbf{r}_1 \times \mathbf{r}_2\|^4} \right) [\mathbf{r}_2]_{\times}^T, \quad (2.57)$$

$$\frac{\partial \mathbf{u}_0}{\partial \mathbf{r}_2} = \frac{\Gamma}{4\pi} \left(\frac{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2 \mathbf{I}_3 - 2(\mathbf{r}_1 \times \mathbf{r}_2)(\mathbf{r}_1 \times \mathbf{r}_2)^T}{\|\mathbf{r}_1 \times \mathbf{r}_2\|^4} \right) [\mathbf{r}_1]_{\times}. \quad (2.58)$$

Here, $[\mathbf{a}]_{\times}$ indicates the skew-symmetric constructed from a vector $\mathbf{a}$, such that the cross product can be written in the form

$$\mathbf{a} \times \mathbf{b} = [\mathbf{a}]_{\times} \mathbf{b} = \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}, \quad (2.59)$$

allowing compact notation of the derivative

$$\frac{\partial \mathbf{a} \times \mathbf{b}}{\partial \mathbf{b}} = [\mathbf{a}]_{\times}. \quad (2.60)$$

The partial derivatives of u_1 are

$$\frac{\partial u_1}{\partial \mathbf{r}_0} = \left(\frac{\mathbf{r}_1}{\|\mathbf{r}_1\|} - \frac{\mathbf{r}_2}{\|\mathbf{r}_2\|} \right)^T, \quad (2.61)$$

$$\frac{\partial u_1}{\partial \mathbf{r}_1} = \frac{\|\mathbf{r}_1\|^2 \mathbf{r}_0^T - (\mathbf{r}_0 \cdot \mathbf{r}_1) \mathbf{r}_1^T}{\|\mathbf{r}_1\|^3}, \quad (2.62)$$

$$\frac{\partial u_1}{\partial \mathbf{r}_2} = -\frac{\|\mathbf{r}_2\|^2 \mathbf{r}_0^T - (\mathbf{r}_0 \cdot \mathbf{r}_2) \mathbf{r}_2^T}{\|\mathbf{r}_2\|^3}. \quad (2.63)$$

Finally, the partial derivatives of u_2 are

$$\frac{\partial u_2}{\partial \mathbf{r}_0} = -\exp \left(-\frac{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2}{\sigma^2 \|\mathbf{r}_0\|^2} \right) \left(2 \frac{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2}{\sigma^2 \|\mathbf{r}_0\|^4} \right) \mathbf{r}_0^T, \quad (2.64)$$

$$\frac{\partial u_2}{\partial \mathbf{r}_1} = \exp \left(-\frac{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2}{\sigma^2 \|\mathbf{r}_0\|^2} \right) \left(\frac{2(\mathbf{r}_1 \times \mathbf{r}_2)^T}{\sigma^2 \|\mathbf{r}_0\|^2} \right) [\mathbf{r}_2]_{\times}^T, \quad (2.65)$$

$$\frac{\partial u_2}{\partial \mathbf{r}_2} = \exp \left(-\frac{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2}{\sigma^2 \|\mathbf{r}_0\|^2} \right) \left(\frac{2(\mathbf{r}_1 \times \mathbf{r}_2)^T}{\sigma^2 \|\mathbf{r}_0\|^2} \right) [\mathbf{r}_1]_{\times}. \quad (2.66)$$

For all rings $b > 0$, the Jacobian matrix is filled with the partial derivatives as

$$\left[\frac{\partial X^{(b)}}{\partial X^{(a)}} \right]_{ij} = \begin{cases} h \frac{\partial u_i^{(b)}}{\partial x_j^{(a)}} \left(\mathbf{x}_i^{(b)}, \mathbf{x}_j^{(a)}, \mathbf{x}_{j+1}^{(a)} \right) & \text{if } j = 0, \\ h \frac{\partial u_i^{(b)}}{\partial x_j^{(a)}} \left(\mathbf{x}_i^{(b)}, \mathbf{x}_{j-1}^{(a)}, \mathbf{x}_j^{(a)} \right) & \text{if } j = n_e, \\ h \frac{\partial u_i^{(b)}}{\partial x_j^{(a)}} \left(\mathbf{x}_i^{(b)}, \mathbf{x}_{j-1}^{(a)}, \mathbf{x}_j^{(a)} \right) + h \frac{\partial u_i^{(b)}}{\partial x_j^{(a)}} \left(\mathbf{x}_i^{(b)}, \mathbf{x}_j^{(a)}, \mathbf{x}_{j+1}^{(a)} \right) & \text{otherwise,} \end{cases} \quad (2.67)$$

for $i, j = 0, 1, \dots, n_e$, $a = 0, 1, \dots, n_r - 1$ and $b = 1, 2, \dots, n_r - 1$.

If $a + 1 = b$, then the diagonal elements are substituted as

$$\left[\frac{\partial X^{(b)}}{\partial X_k^{(b-1)}} \right]_{ii} \leftarrow 1 + h \sum_{m=1}^{n_r-1} \sum_{n=1}^{n_e} \frac{\partial u_i^{(b-1)}}{\partial x_j^{(b-1)}} \left(\mathbf{x}_i^{(b-1)}, \mathbf{x}_{n-1}^{(m)}, \mathbf{x}_n^{(m)} \right). \quad (2.68)$$

The position of the first ring depends only on the yaw angle, therefore

$$\frac{\partial X^{(0)}}{\partial X^{(a)}} = \mathbf{0}. \quad (2.69)$$

Further derivatives are mostly independent from the dimension of the FVW model used.

The position update for all rings $b > 0$ depends on the vortex strength of all elements as

$$\left[\frac{\partial X_{k+1}^{(b)}}{\partial \Gamma_k^{(a)}} \right]_{ij} = \frac{\partial x_i^{(b)}}{\partial u_i^{(b-1)}} \frac{\partial u_i^{(b-1)}}{\partial \Gamma_n^{(a)}} = h \frac{\partial u_i^{(b-1)}}{\partial \Gamma_j^{(a)}} \left(\mathbf{x}_i^{(b-1)}, \mathbf{x}_{j-1}^{(a)}, \mathbf{x}_j^{(a)} \right) \quad (2.70)$$

for $i = 0, 1, \dots, n_e$, $j = 1, 2, \dots, n_e$ and $a = 0, 1, \dots, n_r - 1$,

where the partial derivative of induced velocity to vortex strength is

$$\frac{\partial u_i}{\partial \Gamma_j} = \frac{1}{\Gamma_j} u_i. \quad (2.71)$$

The points defining the vortex filament travel based on their stored free-stream velocity, so for all rings $b > 0$,

$$\frac{\partial X_{k+1}^{(b)}}{\partial U_k^{(b-1)}} = h I_{n_p}. \quad (2.72)$$

Vortex strength and free-stream velocity are passed from one ring to the next for all rings $b > 0$,

Ch. 2

$$\frac{\partial \Gamma_{k+1}^{(b)}}{\partial \Gamma_k^{(b-1)}} = I_{n_e}, \quad (2.73)$$

$$\frac{\partial U_{k+1}^{(b)}}{\partial U_k^{(b-1)}} = I_{n_p}. \quad (2.74)$$

The initialisation of a new ring depends on the entire state of the model through the rotor velocity,

$$\frac{\partial \Gamma_{k+1}^{(0)}}{\partial X_k} = \frac{\partial \Gamma^{(0)}}{\partial \mathbf{u}_r} \frac{\partial \mathbf{u}_r}{\partial X}, \quad (2.75)$$

$$\frac{\partial \Gamma_{k+1}^{(0)}}{\partial \Gamma_k} = \frac{\partial \Gamma^{(0)}}{\partial \mathbf{u}_r} \frac{\partial \mathbf{u}_r}{\partial \Gamma}, \quad (2.76)$$

$$\frac{\partial \Gamma_{k+1}^{(0)}}{\partial U_k} = \frac{\partial \Gamma^{(0)}}{\partial \mathbf{u}_r} \frac{\partial \mathbf{u}_r}{\partial U}, \quad (2.77)$$

where

$$\frac{\partial \Gamma}{\partial \mathbf{u}_r} = h c'_t(a) \mathbf{n}(\psi) (\mathbf{u}_r^T \mathbf{n}(\psi)). \quad (2.78)$$

The Jacobian of rotor speed with respect to all element positions is calculated as

$$\frac{\partial \mathbf{u}_r}{\partial X_k} = \frac{1}{n_u} \sum_{i=0}^{n_u-1} \left(\frac{\partial}{\partial X_k} (\mathbf{u}_\infty(\mathbf{x}_i, \mathbf{q})) + \frac{\partial}{\partial X_k} (\mathbf{u}_{\text{ind}}(\mathbf{x}_i, \mathbf{q})) \right), \quad (2.79)$$

with

$$\frac{\partial \mathbf{u}_\infty(\mathbf{x}_i, \mathbf{q})}{\partial X} = 20(\mathbf{x} - \mathbf{x}_i^{(b)}) \bar{\mathbf{w}}_i^{(b)} - \bar{\mathbf{w}}_i^{(b)} \sum_{i=0}^{n_u-1} \sum_{b=0}^{n_r} 20(\mathbf{x} - \mathbf{x}_i^{(b)}) \bar{\mathbf{w}}_i^{(b)}, \quad (2.80)$$

with the partial derivatives placed in 2D as

$$\left[\frac{\partial \mathbf{u}_{\text{ind}}(\mathbf{x}_i, \mathbf{q})}{\partial X^{(a)}} \right]_{0j} = \frac{\partial \mathbf{u}}{\partial \mathbf{x}_j^{(a)}} (\mathbf{x}_i, \mathbf{x}_j^{(a)}) \quad (2.81)$$

for $j = 0, 1, a = 0, 1, \dots, n_r - 1$,

and for the 3D FVW as

$$\left[\frac{\partial \mathbf{u}_{\text{ind}}(\mathbf{x}_i, \mathbf{q})}{\partial \mathbf{X}^{(a)}} \right]_{0j} = \begin{cases} \frac{\partial \mathbf{u}}{\partial \mathbf{x}_j^{(a)}}(\mathbf{x}_i, \mathbf{x}_j^{(a)}, \mathbf{x}_{j+1}^{(a)}) & \text{if } j = 0, \\ \frac{\partial \mathbf{u}}{\partial \mathbf{x}_j^{(a)}}(\mathbf{x}_i, \mathbf{x}_{j-1}^{(a)}, \mathbf{x}_j^{(a)}) & \text{if } j = n_e, \\ \frac{\partial \mathbf{u}}{\partial \mathbf{x}_j^{(a)}}(\mathbf{x}_i, \mathbf{x}_{j-1}^{(a)}, \mathbf{x}_j^{(a)}) + \frac{\partial \mathbf{u}}{\partial \mathbf{x}_j^{(a)}}(\mathbf{x}_i, \mathbf{x}_j^{(a)}, \mathbf{x}_{j+1}^{(a)}) & \text{otherwise,} \end{cases} \quad (2.82)$$

for $j = 0, 1, \dots, n_e$, $a = 0, 1, \dots, n_r - 1$.

The Jacobian with respect to all vortex strengths is calculated as

$$\frac{\partial \mathbf{u}_r}{\partial \Gamma_k} = \frac{1}{n_u} \sum_{i=0}^{n_u-1} \frac{\partial}{\partial \Gamma_k} (\mathbf{u}_{\text{ind}}(\mathbf{x}_i, \mathbf{q})), \quad (2.83)$$

with the partial derivatives

$$\left[\frac{\partial \mathbf{u}_{\text{ind}}(\mathbf{x}_i, \mathbf{q})}{\partial \Gamma_k^{(a)}} \right]_{1j} = \frac{\partial \mathbf{u}}{\partial \Gamma_j^{(a)}}(\mathbf{x}_i, \mathbf{x}_{j-1}^{(a)}, \mathbf{x}_j^{(a)}) \quad (2.84)$$

for $j = 1, 2, \dots, n_e$ and $a = 0, 1, \dots, n_r - 1$.

The derivative of the free-stream contribution to the disc-averaged velocity over the rotor is

$$\frac{\partial \mathbf{u}_r}{\partial \mathbf{u}_{\infty, i}^{(b)}} = \frac{1}{n_u} \sum_{i=0}^{n_u-1} \tilde{\mathbf{w}}_i^{(b)}(\mathbf{x}_i, \mathbf{q}). \quad (2.85)$$

The Jacobian matrix of the state update with respect to the inputs is defined as

$$\frac{\partial \mathbf{q}_{k+1}}{\partial \mathbf{m}_k} = \frac{\partial}{\partial \mathbf{m}_k} \begin{bmatrix} X \\ \Gamma \\ U \\ M \end{bmatrix}_{k+1} = \begin{bmatrix} \frac{\partial X_{k+1}}{\partial \mathbf{m}_k} \\ \frac{\partial \Gamma_{k+1}}{\partial \mathbf{m}_k} \\ \mathbf{0} \\ \frac{\partial M_{k+1}}{\partial \mathbf{m}_k} \end{bmatrix}. \quad (2.86)$$

The position of the first ring is directly controlled with the yaw angle

$$\left[\frac{\partial \mathbf{X}^{(0)}}{\partial \psi} \right]_i = \frac{\partial \mathbf{x}_i^{(0)}}{\partial \psi} \quad \text{for } i = 0, 1, \dots, n_e, \quad (2.87)$$

which in 2D yields

Ch. 2

$$\frac{\partial \mathbf{x}_0^{(0)}}{\partial \psi} = \frac{\partial \mathbf{R}_z(\psi_k)}{\partial \psi} \begin{bmatrix} 0 \\ r \end{bmatrix}, \quad \frac{\partial \mathbf{x}_1^{(0)}}{\partial \psi} = \frac{\partial \mathbf{R}_z(\psi_k)}{\partial \psi}(\psi_k) \begin{bmatrix} 0 \\ -r \end{bmatrix}, \quad (2.88)$$

and in 3D yields

$$\frac{\partial \mathbf{x}_i^{(0)}(\psi)}{\partial \psi} = \frac{\partial \mathbf{R}_z(\psi_k)}{\partial \psi} \begin{bmatrix} 0 \\ r \cos(2\pi \frac{i}{n_e}) \\ r \sin(2\pi \frac{i}{n_e}) \end{bmatrix} \quad \text{for } i = 0, 1, \dots, n_e, \quad (2.89)$$

given the derivative of the rotation matrix in the relevant dimension,

$$\frac{\partial \mathbf{R}_z(\psi)}{\partial \psi} = \begin{bmatrix} -\sin \psi & \cos \psi \\ -\cos \psi & -\sin \psi \end{bmatrix}, \quad \frac{\partial \mathbf{R}_z(\psi)}{\partial \psi} = \begin{bmatrix} -\sin \psi & \cos \psi & 0 \\ -\cos \psi & -\sin \psi & 0 \\ 0 & 0 & 0 \end{bmatrix}. \quad (2.90)$$

The vortex strength of the first ring depends on the yaw angle as

$$\left[\frac{\partial \Gamma^{(0)}}{\partial \psi} \right]_i = \frac{\partial \Gamma_i^{(0)}}{\partial \psi}(a, \psi) \quad \text{for } i = 1, 2, \dots, n_e, \quad (2.91)$$

where

$$\frac{\partial \Gamma_i^{(0)}}{\partial \psi} = h c_t'(a)(\mathbf{u}_r^T \mathbf{n}(\psi)) \left(\mathbf{u}_r^T \frac{\partial \mathbf{n}(\psi)}{\partial \psi} \right). \quad (2.92)$$

The vortex strength of the first ring depends on the thrust coefficient as

$$\left[\frac{\partial \Gamma^{(0)}}{\partial a} \right]_i = \frac{\partial \Gamma_i^{(0)}}{\partial a}(a, \psi) \quad \text{for } i = 1, 2, \dots, n_e, \quad (2.93)$$

where

$$\frac{\partial \Gamma_i^{(0)}}{\partial a} = \left(h \frac{1}{2} (\mathbf{u}_r \cdot \mathbf{n}(\psi))^2 \right) \frac{\partial c_t'}{\partial a}, \quad (2.94)$$

$$\frac{\partial c_t'}{\partial a} = \begin{cases} \frac{4}{(1-a)^2} & \text{if } a \leq a_t, \\ \frac{2c_{t1}}{(1-a)^3} - \frac{4(\sqrt{c_{t1}}-1)}{(1-a)^2} & \text{if } a > a_t. \end{cases} \quad (2.95)$$

The derivative of saved controls is identity,

$$\frac{\partial \mathbf{M}_{k+1}}{\partial \mathbf{m}_k} = \mathbf{I}_{n_e}. \quad (2.96)$$

2.A.2 Output and objective function

The input sensitivity of the objective function is calculated as

$$\frac{\partial J_k}{\partial \mathbf{q}_k} = \mathbf{Q} \frac{\partial y_k}{\partial \mathbf{q}_k} - 2\Delta \mathbf{m}_k^T \mathbf{R}, \quad (2.97)$$

$$\frac{\partial J_k}{\partial \mathbf{m}_k} = \mathbf{Q} \frac{\partial y_k}{\partial \mathbf{m}_k} + 2\Delta \mathbf{m}_k^T \mathbf{R}. \quad (2.98)$$

The partial derivative of the output to the controls is $\partial y_k / \partial \mathbf{m}_k = 0$ because the power is calculated with controls saved in the state vector $\mathbf{q}_k$. The Jacobian of the output to the states is given as

$$\frac{\partial y_k}{\partial \mathbf{q}_k} = \begin{bmatrix} \frac{\partial y_k}{\partial \mathbf{u}_r} \end{bmatrix} \begin{bmatrix} \frac{\partial \mathbf{u}_r}{\partial X_k} & \frac{\partial \mathbf{u}_r}{\partial \Gamma_k} & \frac{\partial \mathbf{u}_r}{\partial U_k} & \frac{\partial \mathbf{u}_r}{\partial M_k} \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & \frac{\partial y_k}{\partial M_k} \end{bmatrix}.$$

The Jacobian of the power output to the saved control inputs for the two-turbine case presented in this chapter is constructed as

$$\frac{\partial y_k}{\partial \mathbf{M}_k} = \begin{bmatrix} \frac{\partial P_0}{\partial a_0} & \frac{\partial P_0}{\partial \psi_0} & 0 & 0 \\ 0 & 0 & \frac{\partial P_1}{\partial a_1} & \frac{\partial P_1}{\partial \psi_1} \end{bmatrix}, \quad (2.99)$$

with the partial derivatives with respect to the axial induction

$$\frac{\partial P_0}{\partial a_0} = \frac{1}{2} A_r(\mathbf{u}_{r_0} \cdot \mathbf{n}(\psi_0))^3 \frac{4}{(1 - a_0)^2}, \quad (2.100)$$

$$\frac{\partial P_1}{\partial a_1} = \frac{1}{2} A_r(\mathbf{u}_{r_1} \cdot \mathbf{n}(\psi_1))^3 (-8a_1(1 - a_1) + 4(1 - a_1)^2), \quad (2.101)$$

and the partial derivatives with respect to the yaw angle

$$\frac{\partial P_0}{\partial \psi_0} = \frac{3}{2} c'_P(a_0) A_r(\mathbf{u}_{r_0} \cdot \mathbf{n}(\psi_0))^2 (\mathbf{u}_{r_0}^T \frac{\partial \mathbf{n}(\psi_0)}{\partial \psi_0}), \quad (2.102)$$

$$\frac{\partial P_1}{\partial \psi_1} = \frac{3}{2} c'_P(a_1) A_r(\mathbf{u}_{r_1}^* \cdot \mathbf{n}(\psi_1))^2 (\mathbf{u}_{r_1}^{*T} \frac{\partial \mathbf{n}(\psi_1)}{\partial \psi_1}). \quad (2.103)$$

Taking the derivative of the output with respect to the rotor velocity yields

$$\begin{bmatrix} \frac{\partial y_k}{\partial \mathbf{u}_r} \end{bmatrix} = \begin{bmatrix} \frac{\partial P_0}{\partial \mathbf{u}_{r_0}} & 0 \\ 0 & \frac{\partial P_1}{\partial \mathbf{u}_{r_1}} \end{bmatrix}, \quad (2.104)$$

$$\frac{\partial P_0}{\partial \mathbf{u}_{r_0}} = \frac{3}{2} c'_P(a_0) A_r(\mathbf{u}_{r_0} \cdot \mathbf{n}(\psi_0))^2 \mathbf{n}(\psi_0)^T, \quad (2.105)$$

$$\frac{\partial P_1}{\partial \mathbf{u}_{r_1}} = \frac{3}{2} c'_P(a_1) A_r(\mathbf{u}_{r_1} \cdot \mathbf{n}(\psi_1))^2 \mathbf{n}(\psi_1)^T (1 - a_1). \quad (2.106)$$

2.B Convergence study for numerical methods

A variation of model parameters for the numerical methods was performed to show convergence for the model configuration chosen in this chapter as presented in Section 2.4.1. The power output under steady conditions is compared to the reference power from momentum theory for validation, calculated as

$$p_{0,\text{ref}} = c'_p(a_0)A_T(u_\infty(1 - a_0))^3, \quad (2.107)$$

$$p_{1,\text{ref}} = c'_p(a_1)A_T(u_\infty(1 - 2a_0)(1 - a_1))^3, \quad (2.108)$$

and normalised as $\bar{p} = p/p_{\text{ref}}$. Turbine configuration is chosen with a 5 D spacing as used in this chapter and illustrated in Figure 2.1.

The normalised power for a parameter sweep around the operating point from Table 2.1 is shown in Figure 2.10, evaluating combinations of time step h , vortex core size σ , and number of rings n_r . The chosen parameters appear to be a good match for the normalised power of the upstream turbine, $\bar{p}_0$. There is a dependency between time step and core size, so they need to be chosen together. The combination of number of rings to model the wake needs to be chosen together with the time-step such that the wake covers the downstream turbine. As long as the wake adequately covers the downstream turbine position, the power estimate is not very sensitive to variations in the parameters observed here. The 2D FVW does not exactly represent momentum theory, so it is unsurprising that downstream power is overestimated.

A similar sweep with the 3D FVW is shown in Figure 2.11. The trends in the results are similar to those shown in the parameter variations for the 2D FVW. For this configuration without yaw misalignment, results appeared not very sensitive to the number of elements used in the discretisation of the vortex rings.

The number of elements in ring discretisation is important for dynamics of wake redirection under yaw misalignment because finer discretisations better represent the curled wake. Figure 2.12 shows that power gain from yaw misalignment is observed even for coarse discretisations of the vortex filament rings. In the trade-off between computational limits and better convergence of the dynamics, a number of elements $n_e = 16$ was chosen for use in this chapter.




Figure 2.10: Variation of power for 2D FVW normalised by power estimate from momentum theory, for upstream (top) and downstream (bottom) turbine in a configuration as in Figure 2.1. The cross indicates the parameter choice used in this chapter. The base settings for this parameter sweep are: axial induction $a_0 = 0.3$ and $a_1 = 0.33$, number of rings $n_r = 60$, time step $h = 0.2$, and core size $\sigma = 0.1$.




Figure 2.11: Variation of power for 3D FVW normalised by power estimate from momentum theory, for upstream (top) and downstream (bottom) turbine in a configuration as in Figure 2.2. The cross indicates parameter choice used in this chapter. The base settings for this parameter sweep are: axial induction $a_0 = 0.3$ and $a_1 = 0.33$, number of rings $n_r = 40$, number of elements per ring $n_e = 16$, time step $h = 0.3$, and core size $\sigma = 0.16$.




Figure 2.12: Total power estimate with the 3D FVW under yaw misalignment for varying number of vortex filaments in vortex ring discretisation. The dashed line indicates the number of elements used in this chapter. Configuration as in Figure 2.2.

3

Validation with wind tunnel data

Near-wake effects of wind turbine models using the free-vortex wake have been studied extensively, but there is a lack of validation for such predictions in the mid to far wake. This chapter presents a novel validation study using three free-vortex wake models of increasing complexity: an actuator disc, an actuator disc with rotation, and a lifting-line model. We emphasise the application for dynamic wind farm flow control optimisation with a focus on wake redirection using yaw misalignment. For this purpose, wake models should provide sufficiently accurate power predictions at low computational expense to enable real-time control optimisation. Three sets of wind tunnel data are used for validation: flow measurements under steady yaw misalignment, time-resolved flow measurements for a step change in yaw, and turbine output measurements with yaw control and simulated wind direction variation. Results indicate that the actuator-disc model provides the best balance of computational cost and accuracy in power predictions for the mid to far wake, which is not significantly improved upon by the addition of rotation. In the near wake, the added complexity of the lifting-line model may provide value as it models blade loading and individual tip vortices. Altogether, this study provides important validation for further studies into optimisation of wake steering under time-varying conditions and suggests that the actuator-disc model is a suitable candidate for use in a model-predictive wind farm flow control framework.

Chapter contents

3.1	Introduction	53
3.2	Free-vortex wake models	55
3.2.1	Free-vortex wake	55
3.2.2	Wind turbine models	56
3.3	Parameter study and convergence	60
3.3.1	Streamwise spatial discretisation	61
3.3.2	Time discretisation	62
3.3.3	Azimuthal spatial discretisation	64
3.3.4	Vortex core size	65
3.4	Validation with wind tunnel data	66
3.4.1	Wind tunnel experiments	66
3.4.2	Performance measures	68
3.4.3	Experiment replication	70
3.5	Results and discussion	71
3.5.1	Steady yaw misalignment – WTA	71
3.5.2	Step change in yaw – WTB	74
3.5.3	Wind direction variation – WTC	77
3.6	Conclusions	80

This chapter is based on the following publication:

[96] M.J. van den Broek, D. De Tavernier, P. Hulsman, D. van der Hoek, B. Sanderse, and J.W. van Wingerden, *Free-vortex models for wind turbine wakes under yaw misalignment – a validation study on far-wake effects*, Wind Energy Science, <https://doi.org/10.5194/wes-8-1909-2023>, 2023.

Model code supporting the work in this chapter is available at <https://doi.org/10.4121/e32a9868-c5ea-43d3-8969-b1908662b2b2> [97]. Experimental data from WTA is available at <https://doi.org/10.5281/zenodo.5734877> [98], from WTB at <https://doi.org/10.4121/70ae7f4c-f87f-45f1-8360-f4958a60249f.v1> [99], and from WTC will be made available with [100].

3.1 Introduction

The limited availability of offshore and onshore parcels for wind energy production means that large, densely spaced wind farms are commonly used. However, in these farms, wake effects can lead to a significant decrease in power production and an increase in fatigue loading. While farm topology is typically optimised to minimise aerodynamic interaction, it lacks flexibility for time-varying wind conditions [15]. Under those conditions, wind farm flow control uses existing control degrees of freedom to reduce aerodynamic interaction between wind turbines with methods such as wake redirection through yaw misalignment, dynamic induction control with collective pitch control, and wake mixing strategies with individual pitch control [17].

This chapter focuses on the use of yaw misalignment for wake steering, where an intentional misalignment in the yaw angle with respect to the dominant wind direction is used to deflect the low-energy, turbulent wake behind the turbine. After demonstrations of effectiveness in both simulation and wind tunnel experiments, wake steering has been shown to yield power gains in wind farms for predefined yaw-angle offsets under steady conditions in field studies [36, 38–41]. An important aspect of the wind turbine wake under yaw misalignment is the formation of a counter-rotating vortex pair which generates a curled, or kidney-shaped, wake [26–30].

Wind turbine wake models are essential tools for developing and implementing wake steering control strategies. Accurate predictions of wake behaviour allow for optimisation of wind turbine controls for objectives such as power production and reduction of fatigue loading. Current control strategies are mostly based on look-up tables generated by steady-state optimisation with engineering wake models, such as those in the FLORIS toolbox [47]. This includes, for example, the curled wake model [49], which has been extended with dynamics [101] as the steady-state models are limited for use in time-varying conditions. Another approach is the use of Lagrangian particle methods to use the wake models within FLORIS for dynamic wake prediction [55, 56].

Instead of implementing the dynamics into steady-state models, physics-based approaches attempt to simplify first principles to reduce complexity while maintaining essential dynamics. Studies with large-eddy simulation have been successful in control optimisation [102] and recent work has moved towards real-time control by coarsening mesh resolution and adjusting control parameters [60]. On the other hand, simplified flow models based on two-dimensional Navier-Stokes equations, such as WFSim [61] and FRED [62], were developed in an attempt to provide computationally efficient flow estimates for control. However, these have been shown to be unsuitable for yaw control as the wake dynamics under yaw misalignment are fundamentally three-dimensional [63]. A physics-based model for

efficient control optimisation was introduced by van den Broek *et al.* [81] (Chapter 2), modelling the wind turbine wake with an actuator-disc model based on free-vortex methods and representing the curled wake dynamics.

Free-vortex wake (FVW) methods are meshless methods, using Lagrangian elements to model flow dynamics based on the vorticity formulation of the Navier-Stokes equations [67, 68]. This leads to an efficient, skeletal representation of the wind turbine wake. The FVW method was initially applied to wake modelling for helicopter rotors with a focus on tip vortices in studies [69]. From there, it was adapted to application for unsteady aerodynamics in wakes of wind turbine rotors [70].

Several studies have applied FVW methods to model wind turbine wakes, such as a lattice method [71], axisymmetric vortex rings for the near wake under yaw misalignment [72], and the study of the effectiveness of dynamic induction control for near-wake breakdown [73]. The latter model has been shown to only be accurate within the near wake in a comparison with data from large-eddy simulations [85]. Other studies utilise the flexibility of the meshless formulation for the study of wind turbines on floating platforms. These consider, for example, the effects of platform motion on rotor induction [103], unsteady aerodynamics in the near wake [104], rotor performance on a moving platform [74], wake dynamics for specific motions [75], and control optimisation for a coupled aero- and hydro-dynamic model [105].

Most FVW models focus on wake dynamics close to the rotor and, to the best of our knowledge, little validation has been done for the mid to far wake. We define the mid wake from $1D$ to $4D$ and the far wake beyond $4D$ downstream from the rotor, where D is the rotor diameter. Therefore, this study aims to validate the applicability of FVW models for dynamic wind farm flow control optimisation, especially focusing on wake steering through yaw misalignment. For this purpose, we consider three different turbine model formulations: the three-dimensional actuator-disc model, an extension of the actuator disc with a root vortex to model wake rotation, and an extension to a lifting-line model [106]. This simplified lifting-line model is also known as the Joukowsky rotor.

The validity of these models for wake predictions under yaw misalignment is evaluated with three sets of experimental data from wind tunnel measurements; first, a set of data that consists of lidar measurements of wind turbine wakes under steady yaw misalignment [30]; second, a set of time-resolved particle-image velocimetry measurements of the wake following a step change in yaw angle [25]; third, turbine output measurements from an experiment for yaw-based wake steering with wind direction variation [100].

The contribution of this chapter is twofold: (i) an analysis of model parameter choice and suitable levels of simplification of the turbine representation for

modelling the wind turbine wake, and (ii) a validation of free-vortex wake models for mid- to far-wake power predictions with wind tunnel data, in light of control optimisation for yaw control.

The remainder of this chapter is structured as follows. Section 3.2 presents the free-vortex method and develops the three different models of the wind turbine wake, followed by a study of parameter sensitivity and convergence in Section 3.3. The data sets from the wind tunnel experiments and methods for validation are presented in Section 3.4. The results are discussed in Section 3.5, and, finally, the conclusions are shown in Section 3.6.

3.2 Free-vortex wake models

First, we construct the models of the wind turbine wake that are studied in this chapter. The free-vortex methods and straight-line vortex filament definition are introduced in Section 3.2.1. Using these filaments as building blocks, the three wind turbine representations for wake modelling for control optimisation are then described in Section 3.2.2.

3.2.1 Free-vortex wake

The basis of the vortex methods is the vorticity formulation of the Navier-Stokes equations. The FVW method is based on Lagrangian particles that advect downstream. These particles induce a velocity based on their associated circulation strength. The resultant flow velocity may be calculated at any position based on the free-stream velocity and the sum of induced velocities. The vorticity formulation requires the assumption of inviscid and incompressible flow, although diffusion may be approximated. For a further description of the fundamentals, the reader is referred to aerodynamic literature, such as Leishman [67] or Katz and Plotkin [68].

Vortex filaments

The three-dimensional model formulations in this study are based on straight-line vortex filaments. The induced velocity $\mathbf{u}_i \in \mathbb{R}^3$ at a point $\mathbf{x}_0 \in \mathbb{R}^3$ is calculated with the Biot-Savart law from a single vortex filament starting at $\mathbf{x}_1 \in \mathbb{R}^3$ and ending at $\mathbf{x}_2 \in \mathbb{R}^3$, with vortex strength Γ ,

$$\mathbf{u}_i(\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2) = \frac{\Gamma}{4\pi} \frac{\mathbf{r}_1 \times \mathbf{r}_2}{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2} \mathbf{r}_0 \cdot \left(\frac{\mathbf{r}_1}{\|\mathbf{r}_1\|} - \frac{\mathbf{r}_2}{\|\mathbf{r}_2\|} \right), \quad (3.1)$$

where the relative positions $\mathbf{r} \in \mathbb{R}^3$ are defined as

$$\mathbf{r}_0 = \mathbf{x}_2 - \mathbf{x}_1, \quad (3.2)$$

$$\mathbf{r}_1 = \mathbf{x}_1 - \mathbf{x}_0, \quad (3.3)$$

$$\mathbf{r}_2 = \mathbf{x}_2 - \mathbf{x}_0. \quad (3.4)$$

A Gaussian core with core size σ is included to regularise singular behaviour of the induced velocity close to the vortex filament.

$$\mathbf{u}_{i,\sigma}(\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2) = \mathbf{u}_i \left(1 - \exp \left(- \frac{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2}{\sigma^2 \|\mathbf{r}_0\|^2} \right) \right). \quad (3.5)$$

Convection of vortex filaments

Vortex filaments are convected over time according to the combination of the free-stream velocity $\mathbf{u}_\infty \in \mathbb{R}^3$ and the total velocity induced by all filaments $\mathbf{u}_{\text{ind}} \in \mathbb{R}^3$ at the vortex position $\mathbf{x} \in \mathbb{R}^3$,

$$\dot{\mathbf{x}} = \mathbf{u}_{\text{ind}}(\mathbf{x}) + \mathbf{u}_\infty(\mathbf{x}), \quad (3.6)$$

where $\dot{\mathbf{x}} \in \mathbb{R}^3$ is the time derivative of the vortex position. At fixed intervals, a new set of vortex filaments is released from the rotor according to the wind turbine model definition. The oldest set of vortex filaments is then discarded from the simulation, such that a fixed number of sets of filaments n_f is maintained.

Modelling viscous diffusion

Turbulence is not explicitly accounted for when using the FVW to construct models of wind turbine wakes. However, growth of the vortex core may be used to approximate the effects of turbulent and viscous diffusion as

$$\sigma_{k+1} = \sqrt{4\alpha\delta\nu\Delta t + \sigma_k^2}, \quad (3.7)$$

which is Squire's modification of the diffusive growth of the Lamb-Oseen vortex core [107], with the discrete time step k , constant $\alpha = 1.25643$, effective turbulent viscosity coefficient δ to tune the core growth, kinematic viscosity $\nu = 1.5 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$, and time step Δt .

3.2.2 Wind turbine models

The wind turbine models used for this study are the three-dimensional actuator disc as used by van den Broek *et al.* [81] (Chapter 2), an extension with rotation, and a lifting-line model. These three concepts are illustrated in Figure 3.1. Note that a two-dimensional actuator disc could be considered as a further simplification of the wind turbine wake under axisymmetric conditions. It is however not

considered in the current study because it has already been shown to be ineffective in modelling the wind turbine wake under yaw misalignment due to a lack of axisymmetry [81] (Chapter 2).

Coefficients for wind turbine modelling

Ch. 3

The turbine thrust T is calculated according to the magnitude of the free-stream inflow velocity u_∞ ,

$$T = c_t \cdot \frac{1}{2} \rho A_r u_\infty^2 \cos^{\beta_t}(\gamma), \quad (3.8)$$

with thrust coefficient c_t , air density ρ , and rotor swept area A_r . The thrust is assumed to vary over yaw misalignment γ with a cosine exponent β_t which may be adjusted to represent experimental data. Similarly, aerodynamic power p_a is calculated with the power coefficient c_p as

$$p_a = c_p \cdot \frac{1}{2} \rho A_r u_\infty^3 \cos^{\beta_p}(\gamma), \quad (3.9)$$

where the exponent β_p can be used to tune power variation over yaw misalignment, as seen in Hulsman *et al.* [108]. The thrust and power coefficient are assumed to be a function of the axial induction factor a normal to the rotor based on momentum theory,

$$c_t(a) = \begin{cases} 4a(1-a) & \text{if } a \leq a_t, \\ c_{t1} - 4(\sqrt{c_{t1}} - 1)(1-a) & \text{if } a > a_t, \end{cases} \quad (3.10)$$

$$c_p(a) = 4a(1-a)^2, \quad (3.11)$$

with the parameter $c_{t1} = 2.3$ and the induction at the transition point $a_t = 1 - \frac{1}{2}\sqrt{c_{t1}}$. The thrust coefficient calculation is based on momentum theory with a transition to a linear approximation for high induction values, which is an empirical correction based on the Glauert correction [86]. The induction factor is used as a control input to determine thrust and power; it is assumed uniform over the rotor disc.

Actuator disc model (ADM)

An actuator-disc representation of a wind turbine is implemented with the free-vortex method and illustrated in Figure 3.1(a). The actuator disc is assumed to be uniformly loaded so it only sheds vorticity along its edge [68].

At fixed time intervals Δt , a vortex ring discretised in n_e vortex filaments is generated at the edge of the rotor. The vorticity Γ generated along the edge of




Figure 3.1: The vortex filament structures and the direction of circulation for the three different free-vortex wind turbine representations under consideration: (a) the actuator disc model (ADM) builds up a wake with discretised vortex rings, (b) the actuator disc model with rotation (ADMR) adds a root vortex along the wake centre to model swirl due to turbine rotation, and (c) the lifting-line model (LLM) models individual blade circulation, tip vortices, and a combined root vortex. In (c), the vortex structure from a single blade is highlighted for clarity.

an actuator disc is directly related to the pressure differential generated by the disc [106],

$$\Gamma = \Delta t \frac{\partial \Gamma}{\partial t} = \Delta t \frac{1}{\rho} \frac{T}{A_r}, \quad (3.12)$$

where this pressure jump is due to the thrust force. The vortex system is a finite cylinder because a finite number of vortex rings n_r are released. This approximates the theoretical semi-infinite vortex system. The numerical parameters for the model are listed in Table 3.1.

The vorticity generated by the ADM is purely azimuthal as long as the turbine is yaw-aligned with the free-stream wind direction. Under yaw misalignment, the vortex rings deform into the characteristic curled shape of the wake as a counter-rotating vortex pair is formed.

Actuator disc model with rotation (ADMR)

An extension of the ADM is the actuator disc with rotation (ADMR). A root vortex is released along the centre-line of the wake as shown in Figure 3.1(b). Note that the associated distributed vorticity over the disc and the wake boundary is neglected. This root vortex models the swirl in the wake induced by the rotation of the wind turbine rotor. The inclusion of rotation may contribute to modelling asymmetry in wake steering.

Assuming again that the rotor is uniformly loaded, the thrust force is equally distributed over each of the blades,

$$L = \frac{T}{n_b}, \quad (3.13)$$

where L is the individual blade loading and n_b is the number of blades. The bound vorticity at the blade Γ_b is then calculated according to the Kutta-Joukowski theorem,

$$\frac{\partial L}{\partial r} = -\rho u_{\text{rel}}(r) \Gamma_b = -\rho u_\infty \omega r \Gamma_b, \quad (3.14)$$

with the relative velocity u_{rel} along the blade r with rotational velocity ω . Integration over the blade length then yields

$$L = \rho \lambda u_\infty \frac{1}{2} R \Gamma_b, \quad (3.15)$$

where λ is the tip-speed ratio and R the rotor radius. The bound vortex strength of a single blade is then

$$\Gamma_b = \frac{c_t u_\infty \pi R}{n_b \lambda}. \quad (3.16)$$

Combining circulation of the bound vortices of each blade yields the root vortex strength Γ_r ,

$$\Gamma_r = n_b \Gamma_b. \quad (3.17)$$

Lifting-line model (LLM)

The Joukowski rotor model is a lifting-line model (LLM) that assumes uniform blade loading, forming a rotating horseshoe vortex system for each blade [106]. The vortex filament structure is shown in Figure 3.1(c), where the vortex system from a single blade is highlighted. Each blade is modelled with a bound vorticity Γ_b as in Eq. (3.16). The tip vortices coming off from each of the blades have the same vorticity, $\Gamma_t = \Gamma_b$. The root vortex is the combination of the bound vorticity of each of the blades, $\Gamma_r = n_b \Gamma_b$, which is equivalent to the one previously introduced in the ADMR.

3.3 Parameter study and convergence

An important aspect of the wind turbine models from Section 3.2.2 is the sensitivity to parameter changes. This section explores the convergence behaviour of the ADM with the aim of finding a set of suitable parameters in terms of numerical convergence. These parameters are then used for the comparison with wind tunnel data in Section 3.4. For brevity, the convergence behaviour of the ADMR and LLM is performed but not included, as the results are similar to the ADM. The reference parameter values for all three models are listed in Table 3.1. The study is split into four parts: first, the effect of the streamwise spatial discretisation in Section 3.3.1, second, time discretisation in Section 3.3.2, third, the sensitivity to the azimuthal spatial discretisation of the vortex rings in Section 3.3.3, and, fourth, the effect of the core size in Section 3.3.4.

Table 3.1: Numerical parameters for the FVW models as used for this validation study.

		ADM	ADMR	LLM
time step	$\Delta t \cdot u_\infty / D$	0.3	0.3	0.033
number of rings	n_r	40	40	360
elements per ring	n_e	16	17	7
initial core size	σ / D	0.16	0.16	0.16
turbulent growth	δ	100	100	100




Figure 3.2: Relative error in the velocity deficit at the rotor plane $x/D = 0$ and at a downstream distance $x/D = 3$. The spacing between the vortex rings is varied and circulation strength of the filaments is adjusted accordingly. The reference solution for spacing $\Delta x/D = 5 \times 10^{-4}$. The crosses '+' mark the approximate streamwise spacing produced for $\Delta t \cdot u_\infty/D = 0.3$ as used throughout this chapter.

3.3.1 Streamwise spatial discretisation

The streamwise spatial discretisation of the ADM is studied by constructing a cylindrical vortex tube of length $12D$ from discretised vortex rings, approximating the ADM wake. The spacing between vortex rings is varied to study the effect on the wake deficit without the effects of temporal evolution. The number of rings is adjusted accordingly to maintain a constant wake length and the circulation of the vortex filaments is adjusted to maintain the same distribution of total circulation. The velocity error ε_u is defined as

$$\varepsilon_u(x) = |u(x) - u_{\text{ref}}(x)|, \quad (3.18)$$

where u is the induced velocity on the wake centre line and u_{ref} is a reference value, which is generated for a spacing of $\Delta x/D = 5 \times 10^{-4}$.

The convergence behaviour of the velocity deficit with an increasing number of rings is first order as is illustrated in Figure 3.2. The variation in error over filament spacing within the wake, at $x/D = 3$, is small compared to the variation in error at the entry of the tube, which corresponds to the rotor plane, $x/D = 0$. The sharp increase in error for $x/D = 3$ for the coarsest spacing is caused by an insufficiently large core size, which produces an oscillating velocity profile in the wake. This is not an issue as the core size $\sigma/D = 0.16$ used in the rest of this chapter produces a smooth velocity deficit profile for the chosen numerical settings, i.e.,

there are no significant oscillations in velocity magnitude along the wake centre streamline. The inflow at the rotor disc varies a lot for different filament spacing values, which limits the consistency of using local velocity measurements at the rotor plane.

It is important to note that streamwise spatial discretisation is directly connected to the time discretisation and the computational complexity. The largest possible time step is such that a vortex ring is released at every time step. High spatial resolution is thus only possible for small time steps. Additionally, the large number of elements required to generate a wake of sufficient length with high streamwise resolution leads to large increases in the computational cost of the induced velocity calculation; the cost of the induced velocity evaluation increases quadratically with the number of vortex filaments. Small time steps and expensive induced velocity calculation both contribute to a significant increase in computational cost for a given prediction horizon. Therefore, a relatively large time step and coarse spatial resolution are chosen for the purpose of efficient optimisation of wind turbine controls.

The ADMR introduces a single extra vortex filament per time step compared to the ADM, which makes it about $1.1\times$ more expensive with the current numerical settings. The LLM requires four times as many filaments as the ADM for a wake of the same length, which makes a single time step sixteen times more expensive. Accounting for the smaller time steps, simulating a given time with the LLM is theoretically about $140\times$ more expensive than the ADM.

A small benchmark is run on a regular laptop running Windows 10 on an i7-8650 CPU at 1.90 GHz with 8 GB RAM. The benchmark is run in Julia 1.8.0 using the BenchmarkTools module. For comparison to real-time flow, a rotor diameter of $D = 200$ m and inflow wind speed $u_\infty = 8$ m s⁻¹ are used. The ADM simulates a single wake for 600 s of real-time flow in 0.9 s, or $670\times$ faster than real-time. The same simulation takes 1.1 s with the ADMR and 85 s with the LLM, approximately $550\times$ and $7\times$ faster than real-time, respectively.

To put these numbers into perspective, evaluating a single wake in FLORIS with the cumulative-curl model [49] takes about 2 ms on the same laptop. However, this is a steady-state, time-averaged engineering approximation of the wake – it includes no dynamics. A large-eddy simulation of the wake would include dynamics, but requires at least several hours on a computing cluster and is infeasible to run on a regular laptop.

3.3.2 Time discretisation

The time discretisation of the FVW is studied by examining convergence for a first- and second-order integration scheme. In order to perform this convergence experiment, it is necessary to decouple streamwise spatial discretisation and time




Figure 3.3: Relative error in the position of the vortex filaments for varying time discretisation, comparing the first-order explicit Euler method and the second-order explicit Heun method in simulating the wake of a yaw-misaligned rotor. The reference solution for time step $\Delta t \cdot u_\infty / D = 1.5 \times 10^{-3}$. The cross ‘+’ marks $\Delta t \cdot u_\infty / D = 0.3$ as used throughout this chapter.

discretisation. We reformulate the problem such that a number of sub-steps may be taken between releasing vortex rings.

The largest time step considered is $\Delta t \cdot u_\infty / D = 3 \times 10^{-1}$, where one set of vortex filaments is released at every step. From there, the time step is reduced to 5×10^{-3} and a reference solution is generated with a step size of $\Delta t \cdot u_\infty / D = 1.5 \times 10^{-3}$. The convergence is then quantified using the mean position error ε_x of all vortex filaments with respect to the reference solution,

$$\varepsilon_x = \text{mean} \left(\|\mathbf{x}_i - \mathbf{x}_{\text{ref},i}\|_2 \right) \quad \text{for } i = \{1, 2, \dots, n_p\}, \quad (3.19)$$

where $\mathbf{x}_i$ are the n_p coordinates defining the positions of the vortex filaments, and $\mathbf{x}_{\text{ref},i}$ is the reference solution.

Figure 3.3 shows the convergence of time integration of the wake from a yaw-misaligned rotor for decreasing step size with the first-order explicit Euler method as used in this chapter and with the second-order explicit Heun method for comparison. The convergence for a yaw-aligned rotor exhibits similar trends and has been omitted for conciseness. For the numerical parameters presented here, the methods converge as expected. The chosen time step $\Delta t \cdot u_\infty / D = 0.3$ is rather large because of the emphasis on computational efficiency for control optimisation. This is also the reason for choosing explicit Euler, as it requires only a single function evaluation per time step. If a higher degree of convergence is required




Figure 3.4: Relative error in the cross-stream rotor-averaged velocity profile at 5D downstream for the ADM under a yaw misalignment of $\gamma = 30^\circ$. The azimuthal discretisation is $360^\circ/n_e$ and the reference is at 0.7° for $n_e = 512$ elements. The cross ‘+’ marks 22.5° for $n_e = 16$ as used throughout this chapter.

from the time integration, a change of integration method is more efficient than a reduction in time step.

3.3.3 Azimuthal spatial discretisation

The convergence of azimuthal discretisation is tested by varying the number of elements in the vortex rings. A simulation with a yaw misalignment of 30° under steady inflow is performed for the different discretisation steps. When the initial transient of the developing wake is passed, a cross-stream profile of rotor-averaged velocity is recorded. The rotor-averaged quantity is chosen because of the intended application for power predictions, where downstream turbines inherently average flow quantities over the rotor area in the power output. The error norm of this deficit profile ε_d is

$$\varepsilon_d = \text{mean}(|u_r(y) - u_{r,\text{ref}}(y)|) \quad \text{for } -2 < y/D < 2, \quad (3.20)$$

where $u_r(y)$ is the rotor-averaged velocity at cross-stream position y and the reference solution is generated for $n_e = 512$ elements in the ring discretisation.

Figure 3.4 shows that the velocity deficit profile converges for increasing number of vortex filaments in the vortex ring discretisation. The azimuthal discretisation in the current study is for $n_e = 16$.

The time discretisation of the LLM is chosen such that it achieves the same azimuthal resolution, which is for $\Delta t \cdot u_\infty/D = 0.033$. A time step that is nine times smaller implies nine times as many vortex rings – the set of vortex filaments released at one time step – are necessary to model a wake of the same length as the ADM. The LLM thus combines a smaller time step and a more expensive




Figure 3.5: Illustration of vortex particle/filament trajectories for varying initial core sizes. Larger core sizes produce more stable, i.e. less unstable, trajectories. The results in this chapter are produced for an initial core size $\sigma/D = 0.16$.

velocity calculation due to the larger number of vortex filaments, which makes it less attractive for control optimisation for long wakes.

3.3.4 Vortex core size

The choice of vortex core size σ plays a major role in the stability of the FVW models as illustrated for the ADM in Figure 3.5 for different sizes of the Gaussian vortex core. For small constant core sizes, the wake structure transitions into instability, leading to chaotic development of the wake downstream. For larger core sizes, this disturbance growth is smoothed out and the wake structure appears more stable. The size of the vortex core needs to be tuned to the streamwise spatial resolution of the wake. It should at least be large enough to guarantee a smooth velocity profile between vortex filaments, avoiding oscillations in the wake deficit. On the other hand, it should be small enough not to lose information.

Variation in the vortex core size has very little influence on the initial wake




Figure 3.6: Introducing vortex core growth using the Lamb-Oseen model allows tuning of the diffusion to approximately match turbulent mixing in the wake. Slices at $x/D = 5$ downstream for turbulent growth parameter δ from 1 to 500, with an initial core size $\sigma/D = 0.16$.

depth. However, the wake recovery can be tuned with vortex growth, implemented with Eq. (3.7). Figure 3.6 shows how increasing the turbulent core growth parameter δ impacts the recovery of the wake and allows for tuning the representation of turbulent mixing. The validation simulations run in this chapter are for $\delta = 100$, which is chosen to model some wake recovery.

3.4 Validation with wind tunnel data

Following the study of numerical model parameters, it is essential to validate the wake flow and power predictions of the FVW model for yaw control optimisation. This section first presents the available data from the three wind tunnel experiments in Section 3.4.1. The performance measures used to quantify performance are introduced in Section 3.4.2, followed by details on the replication of the experiments with the FVW models in Section 3.4.3.

3.4.1 Wind tunnel experiments

Three sets of experimental data are used in this chapter for the model validation study. The first is a set of steady-state flow measurements for the wind turbine wake under yaw misalignment [30]. The second is a dynamic experiment with

high temporal resolution of a step change in yaw angle [25]. The third is a longer set of turbine output measurements for wake redirection with wind direction variation [100].

The wind direction θ is defined clockwise positive, with 0° along the positive x -axis pointing downstream. The yaw angle ψ is clockwise positive with a 180° shift such that the rotor is fully aligned with the wind direction if $\psi = \theta$. The yaw misalignment $\gamma = \theta - \psi$ such that a counter-clockwise misalignment is positive.

All experiments used the MoWiTO 0.6 turbine with a rotor diameter $D = 0.58$ m [109]. The model turbine has pitch control and the generator can be used for torque control to regulate rotor speed.

Steady yaw misalignment – WTA

The first set of experimental data (labelled WTA) was recorded in the wind tunnel at ForWind, University of Oldenburg and has been published in [30]. The wind tunnel has a $3\text{ m} \times 3\text{ m}$ test section with an active grid, used passively, to control inflow turbulence and boundary layer profiles. It may be operated as an open jet, but for this experiment three movable test sections are installed to form a closed section 18 m in length. The flow circulates in a closed circuit. The experimental data provides measurements of wind turbine wakes under yaw misalignment with steady flow conditions.

The data contains measurements of wakes for yaw misalignment angles $\gamma = \{-30^\circ, 0^\circ, 30^\circ\}$. The turbines were mounted at a hub height of 0.77 m. Operation was at a tip-speed ratio of $\lambda = 5.7$ for aligned flow and $\lambda = 5.3$ when misaligned. The cross-sectional averaged flow measurements were obtained with a WindScanner lidar performing a Lissajous scan within a $3D \times 3D$ area, for uniform and sheared inflow with a turbulence intensity of around 1 %. Wind speed was 7.5 m s^{-1} at hub height. Only the uniform inflow data is analysed in the current study, primarily the vertical measurement planes at downstream distances from $x/D = 1$ to 7.

Step change in yaw – WTB

The second set of experimental data (labelled WTB) was recorded in the Open Jet Facility (OJF) at the TU Delft, with the same set-up as used in [25]. The OJF is a wind tunnel with an open jet that runs in a closed circuit with turbulence intensities between 0.5 % and 2 %. A step change in yaw was measured at high temporal resolution using tomographic particle-image velocimetry (PIV).

A PIV set-up with four cameras was used to measure the flow velocity at downstream distances $x/D = \{1, 2, 3, 4\}$ using helium-filled soap bubbles as flow tracers. The seeding rake for these bubbles was placed at the outlet of the open jet. The measurements were recorded at 500 Hz for a 5 s duration. The step change was initiated 1 s after a trigger signal. This trigger also initiated the PIV measurements

and was used to synchronise the data at the four downstream distances. The turbine yawed from a $\gamma = 0^\circ$ to a $\gamma = 20^\circ$ misalignment with the inflow, at a yaw rate of approximately $\dot{\gamma} = 16^\circ \text{ s}^{-1}$. The turbine operated at a tip-speed ratio of $\lambda = 5.5$ for an inflow velocity $u_\infty = 4.7 \text{ m s}^{-1}$.

Ch. 3

Wind direction variation – WTC

The third set of experimental data (labelled WTC) was again recorded in the wind tunnel at ForWind, University of Oldenburg [100]. Instead of steady yaw misalignment, it provides turbine output measurements of a dynamic yaw experiment. Two turbines were used to test open-loop yaw control strategies under a simulated time-varying wind direction.

The upstream turbine was stationary and yaw-controlled to achieve wake steering. The second turbine was placed $2.66 D$ downstream with an increased hub height of $0.16 D$. The downstream turbine is mounted on an x - y traverse system which models wind direction variation by translating the downstream turbine along a circular path around the upstream turbine. Effectively, this rotates the wind farm layout in the constant wind tunnel inflow, thus changing the wind direction relative to the turbine alignment. The active grid was used passively to generate inflow with a turbulence intensity $TI < 1\%$ and a shear profile with shear exponent $\alpha = 0.28$, where the wind speed was 7.3 m s^{-1} at hub height for the upstream turbine.

The yaw setpoints for the upstream turbine were stored in a look-up table and applied differently for each control experiment. The experiments are labelled BW30, BW60, BW120, BW300, and BW600 based on the length of the wind direction averaging window used in the controller, with shorter windows leading to more frequent yaw variations. Figure 3.7 illustrates the variation in yaw angle and relative wind direction for the BW30 experiment. Each controller experiment yielded 10 min of turbine data, such as generator power, torque, and rotor speed, recorded with a 5 kHz sampling frequency. The raw data has been filtered with a low-pass filter with a cut-off frequency at 20 Hz for noise reduction before use in the current study.

3.4.2 Performance measures

The performance measures in this study reflect the purpose of this model. It is oriented towards control for power maximisation and therefore the predictive qualities for wake deflection and downstream aerodynamic power availability are important aspects to measure.

Wake deflection is determined according to the wake centre position, which is defined as the cross-stream position where aerodynamic power available for a virtual rotor at hub height is minimal. The potential power follows quite directly




Figure 3.7: Time series data from WTC showing yaw misalignment on the upstream turbine and the wind direction variation for the BW30 experiment. The changes in wind direction are achieved by translation of the downstream turbine on an x - y traverse system.

from the measured or simulated flow field and is directly related to downstream turbine performance as available aerodynamic power p^* is calculated from the rotor-averaged velocity $\mathbf{u}_r$,

$$p^* = c_p^* \frac{1}{2} \rho A_r (\mathbf{n} \cdot \mathbf{u}_r)^3 \quad (3.21)$$

where c_p^* is the theoretical maximum power coefficient and $\mathbf{n}$ is the unit vector orthogonal to the rotor plane. This is similar to the potential power method introduced in [110].

For statistical analysis, the fits of the power predictions are evaluated with the variance accounted for (VAF),

$$\text{VAF} = \left(1 - \frac{\text{var}(p - \hat{p})}{\text{var}(p)} \right) \cdot 100 \%, \quad (3.22)$$

and the normalised mean absolute error (NMAE),

$$\text{NMAE} = \frac{\text{mean}(|p - \hat{p}|)}{\max(p)} \cdot 100 \%, \quad (3.23)$$

where $\hat{p}$ is the predicted power from the FVW model, and p is the measured power from the wind tunnel. Note that VAF values closer to 100 % indicate better performance, whereas NMAE values closer to 0 % represent a close fit.

A total comparison of power at different yaw angles is performed by binning the results according to wind direction and yaw angle bins and calculating the mean and standard deviation of the power in each bin. The analysis is the same for both model and experimental data, thus allowing equivalent comparison.

3.4.3 Experiment replication

A visual, qualitative comparison of the available flow measurements provides a general overview of the strengths and weaknesses of each of the models. These are provided for steady-state measurements from WTA and WTB. The cross-stream planes of the flow are considered to be more important than hub height planes because of the three-dimensional nature of wind turbine wakes under yaw misalignment.

A quantitative analysis of the steady-state wake deflection from WTA is performed by analysis of the flow cross-sections for the cross-stream position where potential power is minimal. This is considered to be the wake centre and a measure for predictive power for wake deflection under steady-state conditions.

The WTB experiment is replicated with all three FVW models to analyse the temporal dynamics of the model at high time resolution. The time series of potential power production provide insight into the propagation of yaw effects downstream through the wake.

Finally, the dynamic experiment in WTC is fully replicated with the ADM free-vortex wake model. The upstream turbine is set to the specified yaw angle as recorded in the experiment data and operated under a constant thrust coefficient. The downstream turbine performance is evaluated from the rotor-averaged velocity over a rotor disc at the downstream turbine position. This position varies over time as the turbine is translated to track the specified wind direction from the experiment. For both upstream and downstream turbines, the rotor-averaged velocity is recorded. The downstream velocity is increased by 2.5 % to account for the increased velocity due to the higher hub height in the shear profile of the inflow.

The wake model provides an estimate of the velocity, but the experimental data records generator power. A simple turbine model is specified to account for inertial dynamics,

$$\omega_{k+1} = \omega_k + \frac{\Delta t}{J} (p_a - \tau \omega_k) , \quad (3.24)$$

where ω_k is the angular velocity of the rotor at time step k , Δt is the time step size, J is the rotor inertia, p_a is the available aerodynamic power and τ is the generator torque. The generator torque τ is calculated from the angular velocity of the rotor

ω as

$$\tau = k_1 \omega + k_2 \omega^2, \quad (3.25)$$

with the gains k_1 and k_2 . The form of this control law is based on the turbine controller used in the experiment, which was developed by Petrović *et al.* [111]. Generator torque and angular velocity are multiplied by an efficiency η to obtain a generator power estimate,

$$p_g = \eta \tau \omega. \quad (3.26)$$

The efficiency term is there to capture all inefficiencies in converting aerodynamic power to electrical power, such as a suboptimal power coefficient and drive train losses. The parameters for the turbine model and controller polynomial are found through a least-squares estimate. The controller gains are first estimated using measured rotor speed and generator torque. The rotor inertia and power conversion efficiency are then estimated using the torque controller, modelled wind speed, and measured power at the upstream turbine. The yaw dependency of thrust and power is tuned with respective cosine exponents $\beta_t = 1$ and $\beta_p = 3$. The yaw-aligned thrust coefficient for all experiments is set to $c_t = 0.91$ for $a = 0.33$. The final power estimate is filtered with the same low-pass filter that is applied to the experimental data.

3.5 Results and discussion

This section presents the core results of this chapter, the comparison of the model predictions with wind tunnel data for validation. Yaw misalignment under steady conditions is discussed in Section 3.5.1, followed by the time-resolved step change in yaw in Section 3.5.2. Finally, analysis of the wake steering experiment with wind direction variation is provided in Section 3.5.3.

3.5.1 Steady yaw misalignment – WTA

A visual comparison of cross-stream wake velocity profiles is provided in Figure 3.8, illustrating wake development under yaw misalignment angles $\gamma = \{-30^\circ, 0^\circ, 30^\circ\}$ for experimental data from WTA and simulation results with the ADM, ADMR, and LLM.

For yaw-aligned flow ($\gamma = 0^\circ$), the wake is almost axisymmetric, both in the wind tunnel and in the FVW models. The wakes in the FVW are stable with a consistent wake deficit. There is some underestimation of wake depth in the near wake with the current numerical parameter choice as listed in Table 3.1. Some recovery is modelled through the growth of the Gaussian core, as described in Section 3.3.4, to represent the turbulent mixing that is visible in the wind tunnel measurements.




Figure 3.8: A comparison of normalised streamwise velocity for wind turbines under yaw misalignment. The experimental data from WTA are compared to model results with the ADM, ADMR, and LLM with slices at $x/D = \{1, 3, 5, 7\}$ for yaw misalignment angles $\gamma = \{-30^\circ, 0^\circ, 30^\circ\}$. All three models represent the curled shape of the wake under yaw misalignment. The inclusion of rotation (ADMR, LLM) improves the qualitative representation of the asymmetry and vertical displacement that is observed in the wind tunnel. The deformation of the vortex filament structure with the LLM beyond $x/D = 5$ becomes too large to provide useful predictions.

Both positive and negative yaw misalignment angles result in the generation of a counter-rotating vortex pair and subsequently a curled wake shape, which becomes apparent from $3 D$ downstream. The ADM produces wakes that are symmetric between positive and negative yaw misalignment, as expected. The inclusion of the root vortex in the ADMR models some of the asymmetry in wake shape that is also present in the experimental data. The LLM produces a similar asymmetric deformation of the wake.

A large deformation of the wake is visible from $5 D$ onwards for the wake under yaw misalignment. In the FVW models, this leads to stretching of the vortex filaments and unstable wake structures. The wake of the LLM breaks down beyond $5 D$ downstream because of the large number of vortex filaments in close proximity. The ADMR still demonstrates similar stability in wake structure as the ADM at $7 D$ downstream, although resemblance of the wind tunnel data is reduced.

Finally, some of the details of the experimental data are not represented in the FVW models. The effect of the wake from the turbine tower on power predictions is assumed to be minor and therefore it is not considered in the FVW models, although it is present in the wind tunnel measurements. Furthermore, the models assume the rotor to be uniformly loaded, which is not the case in the experiment. The inclusion of non-uniform rotor loading would increase the model complexity but might improve the modelling of wake deflection [112]. Additionally, the effect of the ground is neglected, whereas the experimental data shows a thin boundary layer near the bottom of the wind tunnel. Ground effects in FVW models may be modelled using a mirrored vortex structure [67]. An initial experiment showed that some asymmetry in wake deflection may be achieved this way. However, for the sake of limiting computational complexity, this option is not presented in this study.

The displacement of the wake centre is evaluated according to the cross-stream position where available aerodynamic power is minimal. Figure 3.9 compares the wake deflection for the data from WTA and the three FVW models. This corresponds to the wakes shown in Figure 3.8.

The ADM appears to have the best fit to the experimental data over the measured downstream distances. The ADMR and LLM show a similar deflection profile, which only shows good agreement with the experimental data up to $x/D = 5$, which matches the visual analysis of the wake structure. Towards $7 D$ downstream in particular, the wake centre predictions diverge. This is related to the large deformation of the wake under these high yaw angles and the numerical stability of the FVW wake structures. The lack of rotation in the ADM allows for a more stable wake structure and a better prediction of the wake deflection further downstream.




Figure 3.9: Wake centre deflection over downstream distance calculated from the flow slices illustrated in Figure 3.8, where the wake centre position is defined as the cross-stream position with minimal available aerodynamic power. The ADM provides the best estimate of wake centre deflection, whereas the rotation in the ADMR and LLM models leads to divergence of the deflection estimate by $x/D = 7$.

3.5.2 Step change in yaw – WTB

Experiment WTB recorded wake development at downstream distances $x/D = \{1, 2, 3, 4\}$ for a turbine yawing from $\gamma = 0^\circ$ to $\gamma = 20^\circ$. The cross-stream wake velocity profiles for $x/D = \{2, 3, 4\}$ are shown in Figure 3.10 when the wake has settled after the step change. The comparison with the ADM supports the qualitative correspondence under steady-state conditions that was found in comparison with WTA. The good quality of the fit follows expectations because the wind tunnel has uniform inflow and a low turbulence intensity.

More importantly, this experiment can provide insight into the wake dynamics for changes in yaw misalignment with a high temporal resolution. The rotor-averaged wind speed for a virtual rotor is evaluated based on the PIV snapshots and is shown in Figure 3.11, together with the realised yaw angle. The actual yaw signal is used to replicate the experiment with the three dynamic FVW models. Additionally, a look-up table is constructed with steady results from the ADM to illustrate the added value of including time-resolved wake dynamics. The downstream available power in steady state is recorded for a yaw misalignment varying in 1° increments and is linearly interpolated to produce the illustrated results. This approach yields results similar to the steady-state engineering models in the FLORIS toolbox.

The value of VAF and NMAE for the three dynamic model simulations is listed




Figure 3.10: Flow slices showing streamwise velocity for yaw misalignment $\gamma = 20^\circ$, comparing experimental measurements from WTB with ADM results.

in Table 3.2, as well as the steady results with the ADM look-up table. The use of precomputed steady results with the ADM performs considerably worse than the dynamic model in terms of accounting for wake dynamics, especially further downstream where the delays for transients propagating through the wake grow larger. The comparison highlights the motivation for implementing a dynamic wake model to be able to account for such wake dynamics in a model-predictive control strategy.

The dynamic ADM and ADMR perform to a similar level of accuracy in this mid-wake region. They are marginally outperformed by the LLM, which is considerably more computationally expensive. These results support the findings from Section 3.5.1 that the inclusion of rotation may improve the qualitative flow representation in the mid to far wake, but the added complexity does not appear necessary for control purposes.

The replication of the yaw step experiment with the ADM in Figure 3.11 shows that the ADM estimates a potential power improvement of similar magnitude as a result of the wake deflection by yaw misalignment. The low-frequency changes show the delays of control effects propagating downstream through the wake. These slow dynamics are also well represented in the dynamic model estimate,




Figure 3.11: The realisation of the $\gamma = 20^\circ$ step in yaw angle measured in WTB is shown in the top graph. The rotor-averaged potential power is calculated directly downstream of the turbine for $x/D = \{1, 2, 3, 4\}$, comparing the experimental data and a simulation with the ADM. The model estimates a power improvement of similar magnitude, although it lacks turbulent disturbances. The transient effects of the change in yaw angle propagating downstream through the wake are well represented in the dynamic FVW implementation of the ADM. These are not accounted for when using a steady-state look-up table with ADM results, which assumes a fully developed wake at every instant.

Table 3.2: Fit quality of the time series of potential power estimates from the replication of the WTB experiment, listing VAF (V) and NMAE (N). The experimental data and ADM estimate, both using steady-state assumptions and dynamic simulation, are illustrated in Figure 3.11.

x	ADM (steady)		ADM (dynamic)		ADMR		LLM	
	V (%)	N (%)	V (%)	N (%)	V (%)	N (%)	V (%)	N (%)
1 D	90.7	3.9	90.8	4.3	90.7	4.4	92.0	3.8
2 D	86.8	6.8	92.1	5.6	92.2	5.6	94.7	4.3
3 D	74.9	9.2	91.5	6.3	91.6	6.5	95.3	4.2
4 D	70.8	10.1	92.8	5.2	92.6	5.5	94.4	4.4

whereas they are not accounted for in a steady-state wake modelling approach. Some turbulence develops in the wake in the wind tunnel that causes variations in the velocity deficit which are not accounted for in the FVW. Considering yaw control is quite slow to actuate, it is more important that the slower dynamics are properly represented than the resolution of turbulence at smaller timescales.

3.5.3 Wind direction variation – WTC

The WTC dataset provides many performance measurements for varying yaw angle and wind direction. It is replicated only with the ADM, as it appears to perform similarly to the ADMR and the LLM yields only minor improvements at a computational expense which is prohibitive for control optimisation in the mid to far wake.

The replication of the experiment with the BW30 controller settings with the ADM yields a power estimate for which the time series is shown in Figure 3.12. The conversion from available aerodynamic power to generator power is performed with a first-order rotor model for which the inertia and controller settings are estimated based on the upstream turbine measurements. The downstream power estimate is then computed using the same turbine model. The turbine model fit has an inertia $J = 0.6 \text{ kg m}^2$. The gains for the torque controller polynomial are estimated to be $k_1 = -3.7 \times 10^{-5} \text{ Nm s rad}^{-1}$ and $k_2 = 6.8 \times 10^{-6} \text{ Nm s}^2 \text{ rad}^{-2}$ for the upstream turbine and $k_1 = -2.1 \times 10^{-5} \text{ Nm s rad}^{-1}$ and $k_2 = 5.5 \times 10^{-6} \text{ Nm s}^2 \text{ rad}^{-2}$ for the downstream turbine. The estimated efficiency in converting potential aerodynamic power to generator power is $\eta = 54 \%$.

The upstream turbine power is estimated with a NMAE of 1.7 % and a VAF of 92.1 %. However, the downstream turbine power estimate is primarily relevant for evaluation of the performance of the wake model. The downstream power series estimate achieves a NMAE of 8.8 % and VAF of 93.3 %, which indicates most of the wake dynamics are accounted for. A similar fit quality is achieved for the power




Figure 3.12: Power predictions with the ADM are compared with experimental data from WTC for the BW30 control setting. The time series fit VAF is 92.1 % for the upstream turbine and 93.3 % for the downstream turbine. The power curves for wind direction and yaw misalignment show the mean and standard deviation in 3° bins. The power variation of the upstream turbine is completely due to variation in yaw misalignment from the open-loop yaw controller. The downstream power is affected both by yaw control on the upstream turbine and by translation for modelling wind direction variation. For larger wind direction magnitudes, the downstream turbine is no longer in waked conditions. The downstream effects of yaw misalignment are clouded by the variation in wind direction and therefore the power-yaw curves filtered for alignment with the wind direction are presented in Figure 3.13.

estimates with the other control settings from [100], which are listed in Table 3.3. The first-order rotor model appears to be adequate in accounting for the delays due to rotor dynamics and the efficiency term captures most of the losses in power conversion.

Figure 3.12 also presents a statistical analysis of these power signals, where the mean and standard deviation are illustrated based on 3° wind direction bins. As expected from the time series, there is a close fit for the upstream power estimate. The variation in power over wind direction is due to changes in yaw angle based on the control strategy. The downstream power estimate matches well for wind directions close to 0° . For wind directions away from 0° , where the downstream turbine is translated out of the waked conditions, the FVW underpredicts the power production on the downstream turbine. It is noteworthy that, in the wind tunnel, the downstream power exceeds the upstream power generated by about 6 % for wind directions where there is no aerodynamic wake interaction.

Besides possible differences between the model turbines themselves, this increase in power is likely due to a combination of two effects. First, the inflow has a shear profile with a power-law exponent of 0.28 and the downstream turbine has

Table 3.3: Fit quality of the time series of generator power estimates from replication of the WTC experiment with the ADM, listing VAF (V) and NMAE (N). The experimental data and power estimate for the BW30 control strategy are illustrated in Figure 3.12.

	upstream		downstream	
control	V (%)	N (%)	V (%)	N (%)
BW600	92.7	3.7	95.2	7.6
BW300	92.5	1.7	93.5	8.7
BW120	93.0	1.5	93.7	8.7
BW60	92.6	1.7	92.9	9.1
BW30	92.1	1.7	93.3	8.8

a higher hub height by $0.16 D$. Consequently, it experiences a higher rotor-average velocity in unwaked conditions. This shear layer is unaccounted for in the FVW which assumes uniform inflow. Second, the wind tunnel has a limited $3 \text{ m} \times 3 \text{ m}$ cross-section, which means there are blockage and speed-up effects due to the presence of the upstream turbine. These effects are especially amplified when the modelled wind direction $|\theta| > 20^\circ$. The lateral flow component due to translation of the downstream turbine is neglected.

The impact of wake steering through yaw misalignment on downstream power production is clouded in Figure 3.12 because of the variation in wind direction. In order to analyse the model representation of wakes under yaw misalignment, the power signals are filtered for wind directions with fully waked conditions, i.e. $-3^\circ < \theta < 3^\circ$. The mean and standard deviation of power are calculated for 3° yaw angle bins and are shown in Figure 3.13.

The power-yaw curve of the upstream turbine is slightly asymmetric in the wind tunnel experiment due to the operation under sheared inflow conditions. This asymmetric power profile is not represented in the ADM because of the model symmetry and uniform inflow. The downstream expected power production matches well given the model simplicity. The benefit of wake redirection is slightly underestimated for large misalignment angles.

The important aspect of the ADM estimate of power for these two turbines is that the trends are captured well. In the end, what matters for control optimisation of wake steering is accuracy in representing the optimal operating point more so than exactness in the predicted power. The presented data indicate that there is a considerable correspondence between the model and experiment, but model implementation in a control strategy will have to point out whether that is sufficient. Additional error integration and state estimation could always be included if required, such as for tracking a power reference.




Figure 3.13: Wind turbine power curve for the modelled yaw misalignment distribution, showing the mean and standard deviation of power in 3° yaw angle bins for the power data shown in Figure 3.12. Data has been filtered to only consider those sections where the turbines are aligned with the wind direction, i.e. $|\theta| < 3^\circ$. The wind tunnel power-yaw curve of the upstream turbine is slightly asymmetric due to sheared inflow. The predicted downstream power matches quite well for these aligned conditions.

3.6 Conclusions

Three free-vortex wake wind turbine models (ADM, ADMR, LLM) are presented in this work for the prediction of wake dynamics under yaw misalignment for control optimisation. The two highlights in this work are (i) a study of parameter sensitivity and convergence and (ii) a comparison with three sets of experimental data from wind tunnel measurements in order to validate the power predictions in the far wake.

The parameter and convergence study indicated that the best results for mid-to far-wake predictions are achieved with the ADM, i.e. when the wake model has minimum complexity. The addition of rotation does improve qualitative agreement of flow fields with experimental data but does not necessarily improve power predictions under yaw misalignment. The LLM may generate a slightly more accurate response, but the computational cost is prohibitive for use in online control optimisation for wake redirection.

The comparison with experimental data illustrates to what extent the FVW models can provide predictions for available power when utilising wake steering control. Even under the assumptions of uniform inflow and uniform rotor loading, there is considerable agreement with experimental data in terms of steady-state wake deflection, dynamic response to yaw change, and power estimates with yaw control and wind direction variation. However, implementation in a control framework will have to point out whether the accuracy is sufficient for the intended purpose of yaw control for power maximisation.

In conclusion, the ADM appears to be a suitable candidate for efficiently predicting the dynamic mid- to far-wake effects of wake steering, a range from approximately $1 D$ to $7 D$. As such, it could play a central role in the development of novel model-based strategies for wind farm flow control. These new controllers could further improve wind farm energy yield as more accurate wake dynamics are included in control optimisation for wake redirection.

For near-wake stability and rotor-plane effects, the LLM has added value as it models individual blades and tip vortices. Further downstream, large wake deformations under yaw misalignment limit the usefulness of the vortex filament approach. A transition to vortex particles or engineering wake models may be a suitable option to continue wake predictions further downstream.

4

Distributed wind farm control

A novel dynamic economic model-predictive control strategy is presented that improves wind farm power production and reduces the additional demands of wake steering on yaw actuation when compared to an industry state-of-the-art reference controller. The novel controller takes a distributed approach to yaw control optimisation using a free-vortex wake model. An actuator-disc representation of the wind turbine is employed and adapted to the wind-farm scale by modelling secondary effects of wake steering and connecting individual turbines through a directed graph network. The economic model-predictive control problem is solved on a receding horizon using gradient-based optimisation, demonstrating sufficient performance for realising real-time control. The novel controller is tested in a large-eddy simulation environment and compared against a state-of-the-art look-up table approach based on steady-state model optimisation and an extension with wind direction preview. Under realistic variations in wind direction and wind speed, the preview-enabled look-up table controller yielded the largest gains in power production. The novel controller based on the free-vortex wake produced smaller gains in these conditions, while yielding more power under large changes in wind direction. Additionally, the novel controller demonstrated potential for a substantial reduction in yaw actuator usage.

Chapter contents

4.1	Introduction	85
4.2	Model development	87
4.2.1	Wake model for control optimisation	87
4.2.2	Modelling secondary steering	90
4.2.3	Directed graph network	91
4.3	Controller synthesis	94
4.3.1	Basis functions for control signal	94
4.3.2	Distributed optimisation	95
4.4	Methods for controller validation	98
4.4.1	Turbine yaw controller	99
4.4.2	Reference wind farm controllers	99
4.4.3	FVW controller settings	100
4.4.4	Wind farm definitions	101
4.4.5	Real-world wind signal	101
4.4.6	Simulation environment	101
4.5	Results and discussion	103
4.5.1	Three-turbine wind farm	103
4.5.2	Ten-turbine subset of HKN	105
4.5.3	Realistic wind variations	109
4.5.4	On wake length and the prediction horizon	113
4.5.5	Closing the loop	116
4.5.6	Towards real-time control	117
4.5.7	Preview-enabled look-up table control	118
4.6	Conclusions	119

This chapter is based on the following publication:

[113] M.J. van den Broek, M. Becker, B. Sanderse, and J.W. van Wingerden, *Dynamic wind farm flow control using free-vortex wake models*, Wind Energy Science Discussions, <https://doi.org/10.5194/wes-2023-119>, in review, 2023.

Data and code for this chapter are available at <https://doi.org/10.4121/50138917-cf01-4780-9d1d-443593b7e974> [114].

4.1 Introduction

Wind farm flow control aims to improve wind turbine performance by reducing aerodynamic wake interaction between turbines which are often placed in large, densely spaced wind farms to effectively make use of limited available space [15]. Strategies such as wake redirection through yaw misalignment and dynamic induction control with blade pitch variations have been shown to achieve improvements in power production and reductions in fatigue loading [17].

Wake redirection makes use of intentional yaw misalignment to steer wakes away from downstream turbines. When effectively applied, a small power loss is incurred on the upstream wind turbine which results in a larger power gain on the downstream turbine. This has been demonstrated in wind tunnel experiments [33–35] and several field studies [36, 38–42].

The control strategies to apply wake steering in wind farms may be roughly divided into model-based and model-free approaches. The latter attempts to synthesise control signals directly from measurements of the wind farm. In wind tunnel experiments, a closed-loop, model-free yaw controller [33] and extremum-seeking control [45] have been demonstrated to produce power gains from wake steering under steady flow conditions. Extremum-seeking control has also been demonstrated in large-eddy simulation (LES) [94]. These data-driven methods have not been tested under realistic time-varying wind direction variations. To improve interpretability of these methods, Sengers *et al.* [115] introduces a purely data-driven wake model with physically explainable parameters. However, it still requires wake measurements which are not generally available in the field.

Recent work on wake steering uses a model-based approach that embeds prior knowledge and allows better generalisation to different operating conditions. The steady-state models in the FLORIS toolbox [47], such as the cumulative curl [49] and Gauss-curl hybrid [116] models, provide an approximation for the time-averaged velocity profiles in the wake. These models allow efficient optimisation of steady-state optimal yaw angles for wake steering to generate look-up tables (LUT) with yaw offsets for varying wind directions. These LUT approaches have been used, for example, for yaw control under steady conditions in LES [31], in a wind tunnel setting with simulated wind direction changes [35], and in a closed-loop control framework with model adaptation under time-varying inflow in LES [50]. Howland *et al.* [42] most recently demonstrates the effective use of a tuned steady-state model for wake steering in a field experiment.

However, the validity of steady-state models may be limited under realistic, time-varying inflow conditions. The inclusion of wake dynamics is essential for active power control in wind farms [52] and the dynamics of realistic wind direction variations need to be accounted for in control optimisation [51]. For that purpose, some studies have adapted the steady-state engineering wake models to include

dynamics [55, 56, 101] or investigated wind direction preview to account for the dynamics of wake propagation [117, 118]. On the other hand, a physics-based approach may naturally include the dynamics of wake propagation. The use of LES for control optimisation showed promising results [59] and recent work has approached real-time control by coarsening mesh resolution and adjusting control parameters [60]. An approximation of wind farm flow using two-dimensional computational fluid dynamics [61, 62] has been attempted and proven useful for induction control [65, 66], but inherently lacks the wake dynamics required to capture the wake deflection under yaw misalignment [63].

A dynamic, control-oriented free-vortex wake model (FVW) of the wind turbine wake was developed for gradient-based control optimisation and shown to capture sufficient wake flow dynamics to model wake deflection for control [81] (Chapter 2). The economic model-predictive control implementation yielded promising results for wake steering under time-varying inflow conditions. The model formulation based on Lagrangian particles allows greater flexibility compared to mesh-based flow calculations [76]. Additionally, the model has been validated for power predictions for wind turbines operating under yaw misalignment [96] (Chapter 3). Despite its flexibility, the optimisation with the FVW is currently limited to single wakes by the stability of the free-vortex methods and the exponential increase in computational complexity with larger numbers of vortex elements.

To extend economic model-predictive control with the FVW to larger wind farms, this chapter develops a distributed approach to control optimisation for wake steering under time-varying inflow conditions. The performance of the novel control strategy will then be evaluated in LES against the greedy control baseline, and, more importantly, a reference controller based on the industry state-of-the-art use of a LUT with steady-state optimised yaw offsets as well as an extension with wind direction preview. In addition to synthetic wind signals, a set of measured wind direction and wind speed variations will be used to evaluate performance in a simulated section of the Hollandse Kust Noord (HKN) wind farm.

The contribution of this chapter is twofold: (i) development of a distributed approach to dynamic economic model-predictive control for wake steering with a free-vortex wake model, (ii) validation of the control strategy under realistic, turbulent inflow conditions with wind direction and wind speed variation.

The remainder of this chapter is structured as follows. Section 4.2 introduces the FVW model for the wind turbine wake and the coupling to facilitate farm-scale optimisation. The model-predictive control strategy is developed in Section 4.3. The reference controllers and simulation test cases for validation are defined in Section 4.4. The results are then discussed in Section 4.5 and, finally, the conclusions are shown in Section 4.6.

Table 4.1: Numerical parameters for the FVW actuator-disc model.

time step	$\Delta t \cdot u_\infty / D$	0.3
number of rings	n_r	40
elements per ring	n_e	16
initial core size	σ / D	0.16
turbulent growth	δ	100
yaw exponent - thrust	β_t	1
yaw exponent - power	β_p	3

4.2 Model development

The core of the novel dynamic model-predictive control strategy is the FVW model, briefly described in Section 4.2.1. In order to implement this model in a farm-scale controller, Section 4.2.2 presents a strategy for incorporating secondary steering effects when a turbine operates in the wake of a yaw-misaligned turbine. Section 4.2.3 then illustrates the strategy for connecting wind turbines into wind farms by constructing a directed graph connecting upstream and downstream neighbouring turbines.

4.2.1 Wake model for control optimisation

The wake model used for yaw control optimisation is an actuator-disc representation of the wind turbine modelled with the free-vortex wake as developed in [81] (Chapter 2) and validated for power predictions for wake steering control with yaw misalignment [96] (Chapter 3), which yielded the current model parameters listed in Table 4.1. The model, illustrated in Figure 4.1, assumes a uniformly loaded actuator disc that sheds vorticity from its edge. These rings of vorticity are discretised in straight-line vortex filaments and advected downstream as Lagrangian particles, forming a skeletal representation of the wind turbine wake.

A non-linear state-space system is defined for the model dynamics which updates the model state vector $\mathbf{q}_k \in \mathbb{R}^{n_s}$, with the number of states n_s , at discrete time-step k as

$$\mathbf{q}_{k+1} = f(\mathbf{q}_k, \psi_k, a_k, \mathbf{u}_\infty), \quad (4.1)$$

where the state vector contains the start and end points, vorticity, and core size for all vortex filaments. The turbine yaw heading ψ_k and the induction factor a_k are control inputs and $\mathbf{u}_\infty$ is the free-stream velocity. The yaw misalignment angle $\gamma = \theta - \psi$ is the difference between turbine heading ψ and wind direction θ .

At fixed time intervals Δt , a vortex ring discretised in n_e vortex filaments is generated at the edge of the rotor. At the same time, a vortex ring at the end of the




Figure 4.1: Free-vortex wake model of the wind turbine wake. Rings of vorticity discretised in straight-line vortex filaments are shed from the wind turbine rotor modelled as an actuator disc forming a skeletal representation of the wake. The wake develops the characteristic curled shape for turbine operation under yaw misalignment.

wake is removed to maintain a finite wake with n_r rings. The vorticity Γ generated along the edge of an actuator disc is directly related to the pressure differential generated by the disc [106],

$$\Gamma = \Delta t \frac{\partial \Gamma}{\partial t} = \Delta t \frac{1}{\rho} \frac{T}{A_r}, \quad (4.2)$$

where ρ is the air density, A_r is the area swept by the rotor, and T is the thrust force. The vortex filaments are convected over time with a rate $\dot{\mathbf{x}} \in \mathbb{R}^3$

$$\dot{\mathbf{x}} = \mathbf{u}_{\text{ind}}(\mathbf{x}) + \mathbf{u}_{\infty}(\mathbf{x}), \quad (4.3)$$

which is the combination of the free-stream velocity $\mathbf{u}_{\infty} \in \mathbb{R}^3$ and the total velocity induced by all filaments $\mathbf{u}_{\text{ind}} \in \mathbb{R}^3$ at the vortex position $\mathbf{x} \in \mathbb{R}^3$.

The induced velocity of an individual vortex filament $\mathbf{u}_i \in \mathbb{R}^3$ at a point $\mathbf{x}_0 \in \mathbb{R}^3$ according to the Biot-Savart law [67, 68],

$$\mathbf{u}_i(\mathbf{x}_0) = \frac{\Gamma}{4\pi} \frac{\mathbf{r}_1 \times \mathbf{r}_2}{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2} \mathbf{r}_0 \cdot \left(\frac{\mathbf{r}_1}{\|\mathbf{r}_1\|} - \frac{\mathbf{r}_2}{\|\mathbf{r}_2\|} \right), \quad (4.4)$$

where the relative positions $\mathbf{r} \in \mathbb{R}^3$ for a vortex filament starting at $\mathbf{x}_1 \in \mathbb{R}^3$ and ending at $\mathbf{x}_2 \in \mathbb{R}^3$, with vortex strength Γ , are defined as

$$\mathbf{r}_0 = \mathbf{x}_2 - \mathbf{x}_1, \quad (4.5)$$

$$\mathbf{r}_1 = \mathbf{x}_1 - \mathbf{x}_0, \quad (4.6)$$

$$\mathbf{r}_2 = \mathbf{x}_2 - \mathbf{x}_0. \quad (4.7)$$

A Gaussian core with core size σ is included to regularise singular behaviour of the induced velocity close to the vortex filament,

$$\mathbf{u}_{i,\sigma}(\mathbf{x}_0) = \mathbf{u}_i(\mathbf{x}_0) \left(1 - \exp \left(-\frac{\|\mathbf{r}_1 \times \mathbf{r}_2\|^2}{\sigma^2 \|\mathbf{r}_0\|^2} \right) \right). \quad (4.8)$$

The effects of turbulent and viscous diffusion are approximated using the growth of the vortex core

$$\sigma_{k+1} = \sqrt{4\alpha\delta\nu\Delta t + \sigma_k^2}, \quad (4.9)$$

which is Squire's modification of the diffusive growth of the Lamb-Oseen vortex core [107], with the discrete time step k , the constant $\alpha = 1.25643$, effective turbulent viscosity coefficient δ to tune core growth, kinematic viscosity $\nu = 1.5 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$, and time step Δt .

The thrust T and aerodynamic power p_a for the rotor model are calculated as

$$T = c_t \cdot \frac{1}{2} \rho A_r u_\infty^2 \cos^{\beta_t}(\gamma), \quad (4.10)$$

$$p_a = c_p \cdot \frac{1}{2} \rho A_r u_\infty^3 \cos^{\beta_p}(\gamma), \quad (4.11)$$

where c_t and c_p are, respectively, the thrust and power coefficient and u_∞ is the magnitude of the free-stream inflow velocity. For performance evaluation in terms of available power for downstream turbines, the free-stream velocity u_∞ in (4.11) is replaced by the rotor-averaged velocity u_r at the position of the downstream rotor, which includes the velocity deficit from the aerodynamic wake. This rotor-averaged velocity is calculated as

$$u_r = \left\| \frac{1}{n_u} \sum_{i=1}^{n_u} \mathbf{u}_\infty(\mathbf{p}_i) + \mathbf{u}_{\text{ind}}(\mathbf{p}_i) \right\|_2, \quad (4.12)$$

where n_u sampling points $\mathbf{p}_i \in \mathbb{R}^3$ are evenly distributed over the rotor area.

The yaw dependence of the coefficients can be tuned with the cosine exponents β_t and β_p for thrust and power, respectively, such as seen in Hulsman *et al.* [108]. The current values for these exponents are based on work by van den Broek *et al.* [96] (Chapter 3), although they may differ in reality [79, 119] and, thus, may require tuning for different turbine types or atmospheric conditions. Additionally, a dependence on thrust force [120] or on wind field heterogeneity [121] is not included in the current work.

The induction factor is used to calculate the thrust coefficient and power coefficient for the model as

$$c_t(a) = \begin{cases} 4a(1-a) & \text{if } a \leq a_t, \\ c_{t1} - 4(\sqrt{c_{t1}} - 1)(1-a) & \text{if } a > a_t, \end{cases} \quad (4.13)$$

$$c_p(a) = 4a(1-a)^2, \quad (4.14)$$

with parameter $c_{t1} = 2.3$ and the induction at the transition point $a_t = 1 - \frac{1}{2}\sqrt{c_{t1}}$ [86]. In the current study, the induction factor is fixed to the optimum value known from momentum theory, $a = 0.33$, however it may also be used as a degree of freedom for induction control or to adapt the model to above-rated operating conditions.

4.2.2 Modelling secondary steering

One important effect that is not immediately accounted for in the FVW is the cumulative effect of wake steering. Wind turbines in the wake of a yaw-misaligned turbine need to yaw less to achieve the same wake deflection as an isolated turbine, as shown in simulation [122] and wind tunnel experiments [34]. This cumulative effect of wake deflection is attributed to cross-flow on the waked rotor and the trailing vortices from the yaw-misaligned turbine. The secondary steering effects have been accounted for in a control-oriented model in FLORIS by calculation of an effective yaw angle [116].

A simulation study is used to develop a method for incorporating these secondary steering effects in the current wake model. The study is performed with LES using settings as described in Section 4.4.6. The turbulent inflow has an average speed of 9 m s^{-1} . The effects of yaw misalignment are measured for 1, 2, 3, and 5 turbines with a $5D$ inter-turbine spacing, where D is the rotor diameter. The layout is aligned with the wind direction. For three, and fewer, turbines, the domain size is $4 \text{ km} \times 2 \text{ km} \times 1 \text{ km}$. The five-turbine test is performed on a $6 \text{ km} \times 2 \text{ km} \times 1 \text{ km}$ domain. Cross-stream flow slices are recorded at $1D$ intervals downstream from the first turbine. The wake deflection, illustrated in Figure 4.2, is calculated based on the average flow over the final 1500 s of the 2000 s simulations.

Based on these simulation results, we present a method for calculation of an induced yaw angle which is used to propagate the effects of secondary steering to downstream turbines with minimal additional complexity. It differs from the solution proposed by King *et al.* [116] because the induced yaw effects are calculated directly from the sampled velocity.

For downstream neighbours, the velocity is sampled over a rotor-disc area. The effective flow direction θ_{eff} is calculated from the velocity components in the horizontal plane. We take the root-mean-square of the wind direction $\theta_u \in \mathbb{R}^{n_u}$

sampled over n_u points to get one effective flow direction,

$$\theta_{\text{eff}} = \text{RMS}(\theta_u). \quad (4.15)$$

The proposed induced yaw angle γ_{ind} is then the difference between the effective inflow and the nominal wind direction,

$$\gamma_{\text{ind}} = \theta_{\text{eff}} - \theta. \quad (4.16)$$

The optimised yaw offset γ^* is the result of the optimisation with the FVW model. The new induced yaw angle reduces this optimised yaw offset to yield the commanded yaw angle γ_{ref} , which is sent to the wind turbine

$$\gamma_{\text{ref}} = \gamma^* - \gamma_i \quad (4.17)$$

with
$$\begin{cases} \gamma_i = \max(\min(\gamma^*, \gamma_{\text{ind}}), 0) & \text{if } \gamma^* > 0, \\ \gamma_i = \min(\max(\gamma^*, \gamma_{\text{ind}}), 0) & \text{otherwise.} \end{cases}$$

The conditional application of the induced yaw ensures the yaw reference does not compensate for induced yaw to achieve zero offset.

Figure 4.2 shows how this induced yaw angle contributes to approximating the secondary steering effects. The wake deflection is defined as the position where potential power from a virtual rotor placed in the stream would be minimal, as used in e.g. [96, 110]. The FVW results are based on individually simulated wakes which have been combined using root-sum-square superposition of the wake deficit. The induced yaw angles from the first two upstream turbines for each turbine are added to the actual yaw misalignment with respect to the free-stream inflow. The downstream turbines operate at a smaller yaw offset magnitude, but achieve similar levels of wake redirection. This captures the secondary steering effects for implementation in the control optimisation strategy. Note that the induced yaw effects are not applied on turbines that are operating without yaw offset, as it would lead to unwanted offsets.

The downstream deflection from the second turbine onwards is captured better for negative yaw misalignments. Wake redirection with positive yaw offsets appears to lead to more deflection on downstream turbines in the LES simulations due to rotating flow in the wake and ground effects, but modelling this asymmetry is out of scope of this chapter. In future work, an asymmetric thrust-yaw curve could be implemented or further refinements could be incorporated in a model adaptation stage in a closed-loop control implementation.

4.2.3 Directed graph network

The communication protocol between upstream and downstream neighbours is constructed based on a directed graph network, similar to, for example, the work




Figure 4.2: Deflection of the wake centre comparing the FVW modelling of induced yaw effects with time-averaged flow from LES. The individual FVW wakes have been combined for this comparison using a root-sum-square superposition of the wake deficit. The cumulative effect of wake steering is captured as a reduced yaw offset is required for similar levels of wake deflection when operating in the wake of yaw-misaligned turbines. The model is symmetric, whereas the LES data shows greater wake deflection from the second turbine onwards when implementing positive yaw misalignments.

by Starke *et al.* [123]. The structure of this network naturally changes with the wind direction as wakes propagate with the flow through the farm. The relevant neighbouring turbines are selected based on arc sectors around the wind turbine as illustrated in Figure 4.3. The arc sectors are defined by a radius of influence and a spreading angle around the predicted inflow. Separate directed graphs are constructed for the upstream and downstream connection, although they may be symmetric.

The upstream graph is used for propagating the induced yaw effects to account for the secondary effects of wake steering. The downstream graph is used to determine which turbines are relevant in the optimisation for wake redirection control. Simulated wake length and prediction horizon are both important in determining suitable arc radius settings; downstream turbines, for example, should only be included in the optimisation problem if adequately covered by the simulated wake length and the prediction horizon. The spreading angle limits the connection to only those wakes that may actually interact through the streamwise wake propagation. It should be wide enough to cover the width of the wake and possible deflection due to yaw misalignment. An example network of FVW models is shown in Figure 4.4 using a symmetric upstream and downstream graph, illustrating how the wake models are connected along the flow direction through the farm.




Figure 4.3: Selection of up- and downstream neighbours based on arc sectors around the inflow wind direction. The resulting directed graph connects the wind turbine in the farm in the direction of flow.




Figure 4.4: Representation of the wakes in a wind farm using a network of FVW models, with an indication of the graph connecting the wakes which have been individually simulated. The highlighted wakes show a wake and its immediate upstream and downstream neighbours. The upstream wake simulation provides induced yaw estimates for incorporating secondary steering in the control signal for the downstream turbine. The downstream neighbour is accounted for in the optimisation for wake steering to minimise negative effects from wake interaction.

4.3 Controller synthesis

In this section we develop an economic model-predictive wind farm controller around the network of FVW models. Section 4.3.1 describes a reduction of the dimensionality of the optimisation using a B-spline basis. The optimisation problem for the open-loop receding horizon control strategy is then defined in Section 4.3.2.

Ch. 4

4.3.1 Basis functions for control signal

Previous work [76, 81] uses a control signal that may be freely chosen at every simulation time step. However, the current implementation of the model uses forward-mode automatic differentiation for constructing the gradients for optimisation, as opposed to the manual derivation of the adjoint method developed by van den Broek *et al.* [81] (Chapter 2). The automatic differentiation framework yields additional flexibility in model development and facilitates improvements in computational performance by minimising code complexity. Furthermore, it drastically reduces the memory requirements for gradient calculation compared to the manual adjoint derivation, which required storing all partial derivatives at every time step. As a trade-off, it comes with a computational cost that scales linearly with the number of control degrees of freedom. For that reason, the current work aims to limit the possible search space to improve optimisation performance.

The dimensionality of the problem is reduced by constructing the control signal using B-splines. For the optimisation, the control signal needs to be defined over a prediction horizon of N_h steps from the current step $k = k_0$. The reference turbine yaw heading ψ is calculated from a spline $s(k, \mathbf{c})$ defined on the range $k \in [k_0; k_0 + N_h]$ as

$$\psi_k = s(k, \mathbf{c}), \quad (4.18)$$

at time step k with n_b the number of B-spline basis functions with the corresponding coefficients $\mathbf{c} \in \mathbb{R}^{n_b}$. Figure 4.5 illustrates the construction of a control signal from an example B-spline basis with $n_b = 7$ splines, starting at $k_0 = 0$ and with a prediction horizon $N_h = 80$ steps.

To further reduce the dimensionality, not all coefficients are left to be free variables in the optimisation problem. The first coefficient is chosen equal to the current yaw angle to ensure a continuous yaw signal, $c_1 = \psi_{k_0}$. The turnpike effect [88], also illustrated by [81] (Chapter 2), leads turbines to always return to greedy control towards the finite optimisation horizon. Therefore, in the example illustrated in Figure 4.5, the final three coefficients, c_5, c_6, c_7 , are chosen equal to the wind direction which leaves the remaining coefficients, c_2, c_3, c_4 , free as the control parameters for the optimisation problem.

The smoothness of the B-spline basis improved the behaviour of the gradient for optimisation with the FVW in trial optimisations. The basis functions average




Figure 4.5: B-spline basis with $n_b = 7$ splines for constructing a yaw control signal over the prediction horizon. The first coefficient is fixed to the preceding yaw reference for continuity, $c_1 = \psi_{k_0}$. The final three coefficients, c_5, c_6, c_7 , are set equal to the wind direction θ at the associated time steps because the optimisation returns to greedy control towards the finite horizon. The remaining three coefficients, c_2, c_3, c_4 , are free in the optimisation.

out noisy contributions to the gradient and smoothen the optimisation landscape. This allows the optimisation problem in the current work to be defined with a lower input regularisation cost while still yielding smooth control signals. The dimensionality reduction from the use of basis functions does limit some of the flexibility in the control solutions that can be found compared to fully free optimisation.

4.3.2 Distributed optimisation

In order to scale the model-based control approach with the FVW to the wind-farm scale, a distributed approach is implemented as illustrated in Figure 4.6 and described in Algorithm 1. In this approach, each individual turbine has its own wake model. The optimisations for all turbines are then performed in parallel, where each of the turbines attempts to optimise its control signal considering wake effects on its immediate downstream neighbours given an expected inflow over the prediction horizon that is kept fixed during the iterations of the non-linear solver. This is an economic model-predictive control problem because the extremum for power maximisation is not known a priori, whereas conventional model-predictive

control is concerned with driving an objective function to zero, such as for tracking a power reference [78].

The full control optimisation problem is solved in a receding horizon control scheme, in which $N_c \geq 1$ is defined to be the number of samples executed before re-optimisation. Larger values reduce the computational requirements, but reduce flexibility under changing predictions as the control signal is re-optimised less frequently. At every re-optimisation step, information is shared between turbines in the farm.

The yaw reference for each individual turbine is defined by the coefficients of the spline basis, of which several are fixed and the n_m free coefficients gathered in the control vector $\mathbf{m} \in \mathbb{R}^{n_m}$. For every turbine, we construct the scalar objective function $J : \mathbb{R}^{n_m} \rightarrow \mathbb{R}$ to optimise the mean power production over the prediction horizon for the current turbine and its immediate downstream neighbours

$$J(\mathbf{m}) = \sum_{k=k_0}^{k_0+N_h} \left(R(\psi_k - \psi_{k-1})^2 + \sum_{i=1}^{n_{t,\text{sub}}} Q p_{k,i} \right). \quad (4.19)$$

The objective function uses an initial condition $\mathbf{q}_{k_0}$ for the wake model at the current time step $k = k_0$ with the state update according to (4.1) using a set of free-stream velocity predictions $\mathbf{u}_\infty$ over the horizon. The power p of turbine i at time step k is calculated following (4.11) and the yaw heading reference ψ following (4.18). The output weight $Q < 0$ such that minimisation of the objective maximises mean power production over the horizon and the input weight $R \geq 0$ balances the output and actuation cost. The number of turbines $n_{t,\text{sub}}$ is size of the subset of the wind farm consisting of the current turbine and its immediate downstream neighbours in the directed graph.

The objective function is then implemented in the optimisation problem

$$\min_{\mathbf{m}} J(\mathbf{m}) \quad \text{subject to } |\gamma_k| \leq \gamma_{\max}, \quad (4.20)$$

where the maximum yaw offsets $\gamma_{\max}$ are enforced as hard limits relative to the predicted inflow. The optimised yaw signal does not include the induced yaw effects, these are taken into account before sending the control signals to the wind turbine yaw controller for implementation in the wind farm. The problem is solved with the BFGS optimisation algorithm [92] although this approach was ineffective in previous work [81] (Chapter 2) because of the noisy optimisation landscape. The smoothing effect of the B-spline basis enabled better convergence trial optimisations.

The controller framework presented here is operated in open loop as data assimilation for state estimation and parameter updates are beyond the scope of the




Figure 4.6: The finite-horizon optimisation problem for economic model-predictive control is solved in parallel with a wake model for each turbine. A central controller communicates with the wind farm to update control set-points and incorporates the predicted inflow. It updates the graphs connecting upstream and downstream neighbours and distributes information. The current control framework is open loop and does not utilise wind farm measurements for state estimation or parameter updates.

current work. Additionally, management of fatigue loading is important for turbine operation, but left out of the control objective. Minimisation of fatigue loads could be achieved by integrating a surrogate model for turbine loads [124, 125], adding an associated cost to the objective function, and subsequently appropriately balancing objective weights.

Algorithm 1 Free-vortex wake controller

```

initialise wind farm from configuration
for turbine in wind farm do
    construct free-vortex wake model
     $q_0 \leftarrow$  run transient with initial inflow
end for
 $k \leftarrow 0$ 
 $k_{\text{final}} \leftarrow t_{\text{final}}/\Delta t$ 
while  $k < k_{\text{final}}$  do
     $u_{\infty} \leftarrow$  inflow over prediction horizon
    graphs  $\leftarrow$  update graphs with  $u_{\infty}$ 
    for turbine in wind farm do
        position, controls  $\leftarrow$  downstream neighbours from graphs
         $m \leftarrow$  minimise  $J(m)$  with  $u_{\infty}$ , position, controls
         $c \leftarrow$  combine fixed and optimised coefficients
         $\psi^* \leftarrow$  spline with coefficients  $c$ 
        for  $i$  in 1 to  $N_c$  do
             $q_{k+i} \leftarrow$  update model State  $q_k$  with  $\psi^*$ ,  $u_{\infty}$ 
             $\gamma_{\text{ind}} \leftarrow$  calculate induced yaw at position
        end for
         $\gamma_{\text{ind}} \leftarrow$  upstream neighbours ( $\leq 2$ ) from graphs
         $\gamma_{\text{ref}} \leftarrow$  reduce  $\gamma^*$  with  $\gamma_{\text{ind}}$ 
    end for
     $k \leftarrow k + N_c$ 
end while
  
```

4.4 Methods for controller validation

Given the novel control strategy constructed around the FVW model, it is imperative to validate its control performance with a suitable reference controller and realistic operating conditions. Section 4.4.1 describes the turbine yaw controller used to implement the reference signals from the wind farm controllers. The reference wind farm controllers are introduced in Section 4.4.2, followed by the settings for the FVW controller in Section 4.4.3. The wind farms for the test cases

are defined in Section 4.4.4 and a realistic time-varying wind signal for driving the simulation study is provided in Section 4.4.5. Finally, Section 4.4.6 describes the simulation environment that is used to measure controller performance.

4.4.1 Turbine yaw controller

The first aspect of testing the control strategy in a realistic wind farm setting is the implementation of a local turbine yaw controller. This yaw controller is used for all control strategies to follow the specified reference signal. The basic yaw controller is implemented based on a dead-band control strategy [51] with an 8° dead band. When the magnitude of the yaw error exceeds the dead band, the turbine will yaw with a constant 0.3° s^{-1} yaw rate until the error reaches zero. Additionally, to avoid persistent unintentional yaw misalignment, error integration is implemented similar to Kragh and Fleming [126]. The turbine will yaw until the error reaches zero if the cumulative error exceeds the equivalent of five degrees of misalignment for five minutes. This is set more strict than in the original work to facilitate a fair comparison of the control strategies.

4.4.2 Reference wind farm controllers

The standard baseline control strategy for wind farm control is greedy control, where each turbine operates individually to track the inflow wind direction without considering collective wind farm performance. This baseline is used in the current study to provide normalised output measures and quantify relative gains. However, a reference wake steering controller is necessary to assess the potential for dynamic model-predictive control.

The current industry state-of-the-art for implementing wake steering uses look-up tables with yaw angles optimised using steady-state engineering models. Therefore, we use FLORIS [47] with the cumulative curl model [49] and the serial-refine optimisation strategy [127] to generate a look-up table with yaw angle offsets optimised for power production in steady-state. A 2° hysteresis is applied on the wind direction signal to avoid excessive yaw actuation around wind directions where the yaw offset in the look-up table changes sign [51].

The model-predictive controller assumes a preview of the inflow over the optimisation horizon. For fair comparison, the greedy controller and the LUT controller use the same inflow information. However, these controllers lack preview and therefore utilise only the instantaneous flow conditions.

Recent work by Simley *et al.* [117] and Sengers *et al.* [118] explores LUT control with preview of the wind direction, selecting yaw offsets from the look-up table based on the inflow direction at a time in the future. With these studies in mind, we implement a preview-enabled look-up table (PLUT) controller to study whether results similar to the economic model-predictive controller might be realised by

utilising a simple control strategy. To do so, we use the same pre-optimised yaw offsets and hysteresis strategy that the LUT controller is based on. However, the yaw reference is selected based on the inflow direction θ at a time $t_{\text{preview}} = t + \Delta t$.

We relate the preview time to the time it takes for the effects of control actions to propagate to downstream turbines. A simple formulation relates the preview window Δt to turbine spacing Δx and the free-stream wind speed u_∞ as

$$\Delta t = \frac{\Delta x}{f_w \cdot u_\infty}, \quad (4.21)$$

where $f_w \leq 1$ is an approximate fraction of the free-stream wind speed at which the wake propagates. Simley *et al.* [117] and Sengers *et al.* [118] find an optimal preview window which, for their configuration, is equivalent to $f_w = 0.9$ and $f_w = 1.0$, respectively. For now, we implement the control strategy with $f_w = 1$ and an inter-turbine spacing of $\Delta x/D = 5$ which corresponds with the spacing along the main rows of wind turbines where wake steering will be applied for the layouts presented in Section 4.4.4. Note that this is a rough preview implementation; further exploration and refinement is outside the scope of the current work.

4.4.3 FVW controller settings

In the current study, the optimisation problem at the core of the FVW controller is solved over a prediction horizon of $N_h = 80$ steps. In order to save some computational expense, the first $N_c = 5$ samples of the optimised control signal are executed before re-optimisation, which is the first 6 % of the prediction horizon. The output weight is set to $Q = -1$ and the input weight $R = 0.001$ balances the output and actuation cost. The optimisation parameters were chosen based on results of exploratory parameter variations. The yaw offset results from the optimisation are limited to maximum yaw offsets $\gamma_{\max} = 30^\circ$.

A B-spline basis with seven coefficients is chosen to provide enough degrees of freedom for control on the given prediction horizon, which corresponds to the example illustrated in Figure 4.5. The first coefficient is chosen equal to the current yaw angle at time step k_0 to ensure a continuous yaw signal, $c_1 = \psi_{k_0}$, and the final three coefficients, c_5, c_6, c_7 , are chosen equal to the predicted wind direction at the associated time steps. The middle three coefficients remain free as the control parameters for the optimisation problem and are used to define the control vector $\mathbf{m} = [c_2 \ c_3 \ c_4]^T$.

The directed graph network is constructed using a spreading angle of 30° and a range of $8D$ for both the upstream and downstream connections. The $8D$ range is the limit for consistent power predictions with the current settings of the FVW model because a finite-length wake is simulated.

4.4.4 Wind farm definitions

The test wind farms use the DTU 10 MW reference turbine [128] with a rotor diameter $D = 178.3$ m and a hub height of 119 m.

The first test case is a three-turbine wind farm (TTWF), illustrated in Figure 4.7, is a relatively simple proof-of-concept to test the novel control strategy under a synthetic time-varying wind direction. The turbines are aligned with a 240° wind direction and spaced $5D$ apart. The case provides room for transitions between greedy control and wake steering. It also requires the controller to account for secondary steering effects to avoid excessive yaw misalignment.

The second wind farm test case is a subset of the Hollandse Kust Noord (HKN) wind farm, scaled by rotor diameter from the actual turbine to the DTU 10 MW reference turbine. The ten turbines in the South-West corner are selected as illustrated in Figure 4.7. For the first HKN test case, labelled HKNA, a synthetic wind direction signal is constructed to test controller performance for several transients and steady-state wind directions. The wind direction signals for the TTWF and HKNA cases are designed specifically to test the controller performance in the respective wind farm layouts.

4.4.5 Real-world wind signal

In order to set up realistic wind variations for the wind farm, we make use of publicly available wind measurements. The raw data from a ZephIR 300 m wind lidar at the HKN location is adapted from the KNMI Data Platform [129].

Two seven-hour time series of wind speed and wind direction are selected from the available measurements and illustrated in Figure 4.8. These time series drive the LES for test cases HKNB and HKNC. The selected data have wind directions $180^\circ \leq \theta \leq 270^\circ$ such that the South-West inflow boundaries can be used for driving the LES domain. Furthermore, the wind speeds are such that the wind turbines operate in region II, below-rated conditions. The measurements record wind conditions at 133 m above sea-level, which is close to the 119 m hub height of the DTU 10 MW reference turbine.

The raw data is cleaned up and interpolated from the original samples at approximately 17 s intervals to 1 s samples with cubic splines. A low-pass filter with a 1/600 Hz cut-off frequency is applied to generate a suitable signal for driving the LES. Higher frequency variations are naturally reintroduced in the turbulent variations of the simulation.

4.4.6 Simulation environment

The controllers are tested in large-eddy simulations (LES) with turbulent precursors using SOWFA [130]. Turbines are modelled with a rotating actuator-disc model of the DTU 10 MW reference turbine [128].




Figure 4.7: Layout of the wind farm test cases and simulation domains, and angle definitions for wind direction θ , turbine heading ψ , and yaw misalignment $\gamma = \theta - \psi$. The three-turbine wind farm has a $5D$ spacing and is aligned along $\theta = 240^\circ$. The ten-turbine subset of Hollandse Kust Noord (HKN) is the South-West corner of the wind farm, scaled to the 10 MW reference turbine.

The three-turbine wind farm is simulated in a $4 \text{ km} \times 2 \text{ km} \times 1 \text{ km}$ domain. The HKN cases are run in a $5 \text{ km} \times 5 \text{ km} \times 1 \text{ km}$ domain. The positioning of the turbines in the domains is illustrated in Figure 4.7. The base cell size is set to 20 m in all directions. A single refinement is applied to the bottom layer ($z < 300 \text{ m}$) to 10 m cells. This yields a total of approximately 9.7×10^6 grid cells. The simulations are run with a 0.5 s time step.

Turbulent precursors are prepared before the controller simulations by simulating for 20 000 s to develop turbulence and then forcing the specified wind direction and wind speed variations. The wind direction and speed appear to change almost uniformly throughout the flow field. The use of the same precursor data for all control strategies allows a comparison of the differences in output measures originating from control.




Figure 4.8: Time series of lidar measurements of wind direction and wind speed at the HKN location [129]. The raw data is post-processed and low-pass filtered with a 1/600 Hz cut-off frequency. Two seven-hour time series with wind direction $180^\circ \leq \theta \leq 270^\circ$ and below-rated wind speeds are selected for driving the realistic wind variations in the LES.

4.5 Results and discussion

The performance of the novel FVW controller is first evaluated on the three-turbine wind farm in Section 4.5.1. Subsequently, it is tested on the ten-turbine subset of HKN with synthetic wind direction variation in Section 4.5.2 and with realistic wind variations in Section 4.5.3. Section 4.5.4 comments on the limitations of optimisation with finite-length wakes on a finite horizon and Section 4.5.5 provides a perspective towards closed-loop control. A benchmark of computational performance is presented in Section 4.5.6 to discuss the steps towards real-time optimisation. Finally, Section 4.5.7 discusses the potential for preview-enabled look-up table control.

4.5.1 Three-turbine wind farm

The three-turbine test case is a relatively simple proof-of-concept to test the novel control strategy. The yaw offsets implemented by the three controllers are illustrated in Figure 4.9. Intentional yaw misalignment is applied to turbines 1 and 2 in all control strategies. The maximum offsets utilise the $\pm 30^\circ$ bounds applied to the optimisation problem. No offsets are applied to turbine 3, which is the most downstream turbine. It is always controlled towards alignment with the local free-stream wind direction for the range of wind directions studied here. The magnitude of yaw misalignment on turbine 2 is lower than on turbine 1 for both FVW and LUT controllers. This is the result of accommodating for secondary steering effects in the yaw control strategy. The induced yaw effect from operating in the wake of the yaw-misaligned turbine 1 lowers the required angle of misalignment for a similar level of wake redirection.




Figure 4.9: Yaw offsets realised for the three-turbine test case. Turbines 1 and 2 implement intentional yaw misalignment for wake steering around turbine 3. The FVW controller anticipates wind direction changes and accounts for secondary steering effects.

An important feature of the yaw reference generated by the novel FVW controller is the anticipation of changes in wind direction – the turbines yaw before the wind has actually rotated. The LUT controller, on the other hand, reacts to changes as they happen. The basic PLUT implementation realises an effect on the yaw reference for turbines 1 and 2 that is similar to the FVW controller behaviour by anticipating the transients. However, turbine 3, which is most downstream, tracks the instantaneous wind direction in the FVW controller, but yaws in advance of the transients with the PLUT approach. This leads to a longer time spent in misaligned operation, where yaw-aligned operation would be optimal.

The gains in power production of the FVW controller over the LUT appear mainly during the transients in wind direction as illustrated in Figure 4.10, with the PLUT controller achieving similar results. The power lost due to misaligned operation is initially sacrificed as the controller anticipates changes, which results in a gain in production following the transient. The FVW controller makes use of the dynamics of propagation of the wakes for long-term gains in power production, which can be seen in the normalised energy E produced since the start of the simulation. The power gains with FVW and PLUT controllers highlight the importance of considering the wake propagation dynamics when dealing with time-varying inflow conditions. The optimisation over future inflow conditions with the FVW, as well as the inclusion of preview in the LUT can both produce control signals that provide better performance than the LUT based on steady-state assumptions.




Figure 4.10: Wind farm performance for the three-turbine test case, where the top two plots show the driving wind direction and wind speed for the simulation. The third row shows energy produced relative to greedy control and the bottom row shows relative power production. The FVW controller improved power generation during and following transients by anticipating changes and performs approximately equivalent to the LUT in steady state.

The performance in steady-state is approximately equivalent between the three wake steering controllers.

The cumulative results for the TTWF are shown in Figure 4.11 and listed in Table 4.2. In terms of power production with respect to greedy control, the implementations of wake redirection with the FVW and the PLUT controllers yield a 3.8 % gain which exceeds the 2.7 % achieved with the LUT approach. The demand on the yaw actuators is measured using the yaw travel $\Delta\psi$, which is the total angular distance covered during the length of the simulation. The power improvements with the FVW are achieved with only a slightly increased demand on the yaw actuators as the total yaw travel increase compared to the greedy baseline is 58.1 % for the LUT and 69.3 % for the FVW controller. The yaw travel for the LUT and PLUT controllers is identical as they are based on the same wind direction signal and yaw offsets.

4.5.2 Ten-turbine subset of HKN

We expand the results from the three-turbine case by considering the ten-turbine subset of the South-West corner of the HKN wind farm with a synthetic wind




Figure 4.11: Cumulative results for the three-turbine test case in terms of total relative power production and yaw travel. The FVW controller improved power production at a slight increase in yaw travel compared to the LUT controller. In this case, the PLUT achieves the same improvement in power production as the FVW.

direction variation defined in Figure 4.15. A series of flow snapshots from the LES are provided in Figure 4.12 to illustrate the discussion of controller performance.

The cumulative performance of the FVW with respect to the LUT and PLUT is illustrated in Figure 4.13 and listed in Table 4.2. The FVW controller produces a 3.2 % gain in mean power production which exceeds the gain of 2.5 % from the LUT controller. This gain is consistent with the improvement over the LUT controller found in the TTWF case. The PLUT controller only realises a 2.6 % power gain, which is a slight improvement over the LUT, but much less than is achieved with the FVW. The FVW notably reduces the yaw travel demand, increasing 41.5 % over greedy control, whereas the LUT and PLUT controllers lead to a 55.4 % increase. This contrasts the results from the TTWF case where a slight increase in yaw travel was observed.

Turbine 1, which is upstream in all simulated wind directions, loses a bit more power comparing the FVW to the LUT as it operates under yaw-misaligned conditions for longer. However, this is offset by the power gain coming mostly from turbines 2 to 7, which are relatively close together along the wind directions considered. Unlike the TTWF case, the power gains from the PLUT controller are not equivalent to that of the FVW controller.

Turbines 8 to 10 are further downstream and are therefore not always accounted for in the optimisation with the FVW as, for the wind directions considered, they are often beyond the finite length of the simulated wakes given the current controller settings. The implementation of preview on these downstream turbines leads to a slight loss in performance comparing the PLUT to the LUT controller.




Figure 4.12: Series of hub height flow snapshots from LES of the HKNA test case with the FVW controller. In the initial transient, (a) all wind turbines are aligned with the mean inflow direction. Wake steering solutions are illustrated in (d) and (e) for the southern row of wind turbines 1 – 3 – 5 and (k) and (l) for the western row of wind turbines 1 – 2 – 4 – 6 – 8. Waked turbines have a reduced yaw offset because of the modelling of secondary steering effects. For certain wind directions, long wakes impact farm performance which are not accounted for in the FVW due to the limited prediction horizon. For example, (d) and (e) show the wake from turbine 4 impinging on turbine 9 and (h) and (i) show turbine 8 operating in the wake from turbine 3.




Figure 4.13: The FVW controller improves power production with respect to the LUT and PLUT approach and reduces the increase in total yaw actuation required. The gain comes mostly from turbines 2 to 7, whereas turbines 1 and 8 to 10 lose some power with respect to the LUT and PLUT controllers.

Table 4.2: Mean power production and cumulative yaw travel for the four test cases, where HKNA, HKNB, and HKNC feature the same ten-turbine wind farm and TTWF features a three-turbine wind farm. Increases are noted relative to the greedy control baseline.

		power (MW)		yaw travel (°)	
HKNA	greedy	61.23		1247	
	LUT	+1.51	+2.5%	+691	+55.4%
	PLUT	+1.59	+2.6%	+691	+55.4%
	FVW	+1.96	+3.2%	+518	+41.5%
HKNB	greedy	22.79		1749	
	LUT	+0.49	+2.2%	+2928	+167.4%
	PLUT	+0.59	+2.6%	+2788	+159.4%
	FVW	+0.53	+2.3%	+615	+35.1%
HKNC	greedy	27.86		2052	
	LUT	+1.46	+5.2%	+4232	+206.3%
	PLUT	+1.85	+6.6%	+4237	+206.5%
	FVW	+1.58	+5.7%	+2780	+135.5%
TTWF	greedy	16.36		364	
	LUT	+0.44	+2.7%	+211	+58.1%
	PLUT	+0.63	+3.8%	+211	+58.1%
	FVW	+0.63	+3.8%	+252	+69.3%

This lack of wake redirection away from far downstream turbines is also apparent in the yaw offsets applied as illustrated in Figure 4.14. Turbines 1, 3, and 4 have steady-state segments where no yaw misalignment is applied in the FVW controller, even though the LUT prescribes offsets for these wind directions. Their downstream neighbours are beyond the length of the simulated wakes with the FVW and can therefore not be accounted for in the model-predictive control optimisation with the current controller configuration.

The power generation over time for this test case is illustrated in Figure 4.15. The underperformance of the FVW controller in the initial segment is due to the lack of yaw misalignment on turbines 1 and 4, which leads their wake to impinge on turbines 9 and 10, whereas the yaw misalignment specified by the LUT controller minimises this negative aerodynamic interaction. The final segment of the simulation shows particular benefit from wake steering as the wind direction is aligned with the western row of turbines 1–2–4–6–8. The gains in power for the FVW controller over the LUT controller emerge during the transients in wind direction. Accounting for the propagation dynamics of the wakes leads to fewer instances of loss compared to greedy control. In steady state, the FVW controller with the current controller settings performs approximately equivalent to, or slightly worse than the LUT controller.

The PLUT controller notably underperforms even with respect to the LUT controller for a large part of this simulation. The simple preview implementation produces some gains following the transients, but sacrifices more power to achieve this. These losses may be due to the large wind direction variations and the wind farm layout, in addition to the implementation of preview on turbines that should be in yaw-aligned operation.

4.5.3 Realistic wind variations

The previous two cases highlighted the potential for the gains in terms of power generation and yaw travel reductions that may be achieved with the FVW controller. The wind direction variations were, however, specifically designed to test the added value of the dynamic model-predictive control framework and therefore lack realism. The two cases HKNB and HKNC are simulated using measured wind data to demonstrate controller performance under real variations in wind speed and direction.

Figure 4.16 summarises the total improvement in power production with respect to greedy control, which is also listed in Table 4.2. In case HKNB, the increase in power generation by wake redirection is improved from 2.2 % with the LUT to 2.3 % with the FVW controller and 2.6 % with the PLUT. The increased yaw travel is limited to only 35.1 % with the FVW compared to the 167.4 % with the LUT approach. The minor differences in yaw travel between LUT and PLUT




Figure 4.14: Realised yaw offsets for LUT and FVW controllers for the HKNA test case. Notably, there are some steady-state segments where turbines 1, 3, and 4 are not misaligned by the FVW controller where the LUT does prescribe a yaw offset. This is due to the limitations of the simulated wake length and prediction horizon in the current settings of the FVW controller.




Figure 4.15: Relative energy produced and power production for the HKNA test case. During transients in wind direction, the LUT approach loses power with respect to greedy control. The FVW controller loses a bit as it anticipates changes, but then gains power over the LUT controller. The initial steady-state segment also shows underperformance with respect to the LUT approach. The PLUT controller appears less effective during transients and only slightly improves over the LUT approach.

controllers are due to the treatment of the end of the time-series simulation. Case HKNC shows an increase of power production of 5.7 % with the FVW compared to 5.2 % with the LUT and 6.6 % with the PLUT, as well as a reduction of additional yaw travel from 206.3 % to 135.5 %. The losses of the FVW with respect to the LUT controller appear on turbines 9 and 10, which are far downstream from their upstream neighbours, beyond $12D$ downstream for most of the simulated wind directions.

These results show that some of the improvement in wind farm performance from a dynamic economic model-predictive control approach is maintained under realistic, time-varying wind conditions, where both wind direction and speed change over time. However, under certain conditions, unnecessary losses are incurred with respect to the LUT controller due to the limitations of the FVW controller with the current settings. A simple preview implementation appears more effective in accounting for the effects of wake propagation, while not being limited by finite wake length simulation and receding horizon predictions.

The performance over time is shown in Figure 4.17. The relative energy production over time shows that the power gains from the FVW controller over the




Figure 4.16: Controller performance relative to greedy control in terms of power production. In both cases, the PLUT controller yields the biggest increase in power production compared to greedy control, although the FVW controller also outperforms the LUT approach. Additionally, the FVW achieves these gains with a lower total cost in terms of yaw actuator duty. The FVW controller shows a tendency to underperform compared to the LUT for turbines 9, and 10 which are more than $12D$ away from their upstream neighbours for most of the simulated wind directions.

LUT controller are consistent throughout most of the simulated time series. The performance of the FVW and PLUT controller is equivalent for large parts of the simulation. The energy production with the PLUT sometimes exceeds the FVW, but some additional losses are incurred that bring it back to the same level.

The final segment of the HKNC test case exhibits a pattern with some large performance differences between the LUT and FVW controllers. This is where the PLUT controller achieves a large gain with respect to the FVW controller, whereas they realised similar production until that point in time. This segment is illustrated in more detail in Figure 4.18 with relative power production and the yaw heading of turbine 1. The wind direction oscillates slightly around $\theta = 201.5^\circ$, which is aligned with the western row of turbines 1 – 2 – 4 – 6 – 8. The yaw action of turbine 1 is representative of the control signal applied to turbines 2, 4, and 6 further downstream.

Due to the limits of the prediction horizon in the FVW controller, the FVW controller produces a control signal that switches the direction of wake steering with the oscillations in the wind direction. On the contrary, the implementation of hysteresis in the LUT controllers produces a consistent yaw offset reference to one side when combined with the local turbine yaw controller. Without hysteresis, the LUT controllers would present the same switching behaviour currently observed in the FVW controller.

This difference in control signal leads to significant variations in relative power production. For this wind direction variation from approximately $t = 22\,000$ s to $23\,700$ s, the predictive action of the FVW controller anticipates gains that are not fully realised. The losses from the yaw movements exceed the gains from the wake steering in the optimal direction. The final segment from $t = 23\,700$ s onwards shows how the predictive controller anticipates the wind direction variation to yield a net gain in power production compared to the LUT. The PLUT controller is able to realise these gains without the losses incurred with the FVW control signal and ends up with the largest average power production.

4.5.4 On wake length and the prediction horizon

The results from the control test cases show some of the limitations of the proposed model-predictive control approach. The finite-horizon optimisation can not account for turbines that are outside the simulated wake length or beyond the prediction horizon.

If wind turbines are placed along a straight line, the simulated wake and optimisation horizon only needs to be long enough to cover optimisation from one turbine to the next downstream neighbour to trigger wake steering. However, for longer rows of turbines, the segment from HKNC shown in Figure 4.18 demon-




Figure 4.17: Relative energy produced and power over time for the two data-driven test cases, HKNB and HKNC. The driving wind direction and wind speed are shown in the top two rows. The third row shows the cumulative energy time series normalised with respect to the greedy baseline controller. Both the LUT, PLUT, and the FVW controller exhibit significant improvements over greedy control. The fourth, bottom row shows the power production of the controllers normalised by the greedy baseline.




Figure 4.18: Segment of the HKNC test case shown in Figure 4.17. The inflow wind direction oscillates around the western row of turbines 1 – 2 – 4 – 6 – 8, which is aligned at $\theta = 201.5^\circ$. The yaw heading of turbine 1, ψ_1 , is representative of that applied to turbines 2, 4, and 6. The FVW controller switches wake steering directions from $t = 22\,000$ s to $23\,700$ s, whereas the hysteresis in the LUT controller produces a constant yaw offset. The excessive yaw action in the FVW results in underperformance for this segment. Beyond $t = 23\,700$ s, the FVW correctly anticipates the wind direction variation producing a net gain in power production.

strates that longer horizons will probably be beneficial to avoid excessive switching of the wake steering direction.

For large inter-turbine spacing without intermediate downstream turbines, long wakes will need to be simulated with long prediction horizons to be able to properly account for the downstream effects and reach wake steering yaw control solutions. This limitation is apparent in the lack of performance improvement for turbines 9 and 10 in all the HKN cases. For most of the wind directions under consideration, they are too far downstream to be accounted for in the finite-horizon optimisation with the FVW. Very long prediction horizons would be necessary to account for downstream effects, but long prediction horizons come at considerable computational cost as both simulating longer wakes and longer prediction horizons increase computational expense. Doubling both the length of the wake and the prediction horizon would lead to roughly an eightfold increase in computation time. Maintaining a similar degree of freedom in the control signal by also doubling the number of free spline coefficients then yields an optimisation problem that is approximately $16\times$ more expensive. Additionally, longer wakes stretch

the limits of what can be predicted with the physical model due to inherent instabilities in the free-vortex methods.

The steady-state optimisation with FLORIS does include these long wakes because it essentially solves a mean-flow, infinite horizon version of the control problem. For steady wind directions, the optimal yaw angles for wake steering from the steady optimisation can yield higher power production than those found through receding horizon control with finite-horizon optimisation.

Furthermore, due to the bimodal nature of wake steering, the receding horizon controller may end up implementing yaw offsets in the suboptimal direction, where the cost to switch directions may not outweigh the gain in power over the finite horizon, even though that may be optimal in an infinite-horizon sense. The steady-state optimisation does not suffer from this limitation, but will lose power when atmospheric conditions violate the mean steady-state assumptions too much. The LUT approach might then apply yaw misalignment to redirect wakes around turbines which will not propagate there due to variations in wind direction. This sacrifices power generated for an expected return that is never achieved. This is the result of a lack of inclusion of dynamic effects such as continuously varying wind conditions and propagation of wakes.

The optimal control approach might combine aspects from both receding horizon control and infinite horizon optimisation. This could enable synthesis of a controller that consistently converges to optimal solutions in steady state, while incorporating the dynamics of wake propagation for power gains during inflow transients.

4.5.5 Closing the loop

The current performance achievements are realised with an open-loop controller architecture by assuming a reasonably accurate model and wind speed and velocity predictions. However, Figure 4.2 already highlights differences between the simulation framework and modelled wake deflection. It shows that the incorporation of secondary effects of wake steering is only an approximation. Furthermore, the modelled deflection is symmetric for positive and negative yaw misalignment, whereas a clear asymmetry appears in the LES data.

For adaptation of model errors and incorporation of measurements into the model state, a closed-loop control framework is required. This would allow the control strategy to adapt to varying atmospheric conditions such as veer, shear, and turbulence intensity, as well as tune model parameters such as the turbulent growth parameter δ or the yaw exponents β_p or β_t .

One strategy that is promising for closing the loop is the Ensemble Kalman filter (EnKF), which has previously been developed for state estimation adaptation of steady-state models [50, 79]. Becker *et al.* [80] developed the EnKF for wind




Figure 4.19: The benchmark for the computational cost of control optimisation shows how the optimisation time scales approximately linearly with the number of steps in the prediction horizon. The computational expense scales quadratically with the wake length, determined by the number of rings n_r in the wake simulation. The horizontal lines provide an indication of the level below which real-time control optimisation is achieved for $u_\infty = 4 \text{ m s}^{-1}$ and $u_\infty = 9 \text{ m s}^{-1}$ with an update every five discrete time steps.

field estimation in a model-based setting with a dynamic engineering wake model. Additionally, Shapiro *et al.* [131] showed that closing the loop allows inclusion of unmodelled dynamics.

4.5.6 Towards real-time control

In order to verify the potential for real-time control, a small benchmark is run on a regular laptop running Windows 10 on an i7-8650 CPU at 1.90 GHz with 8 GB RAM. The benchmark is run in Julia 1.8.0 using the BenchmarkTools module.

The results of the benchmark are shown in Figure 4.19 to illustrate the scaling of the computational cost for solving the optimisation problem. The optimisation time scales linearly with the length of the prediction horizon and quadratically with the wake length, which is determined by the number of rings n_r in the wake simulation. The cost of the optimisation also scales linearly with the number of control degrees of freedom, which is set to three as is done throughout the current work.

The non-dimensionalisation of the FVW by rotor diameter and wind speed leads to a dependency on wind speed in measuring the performance relative to real-time. Therefore, we report values for rotor diameter $D = 178.3 \text{ m}$, and relative to the inflow wind speeds $u_\infty = 4 \text{ m s}^{-1}$ and $u_\infty = 9 \text{ m s}^{-1}$. With the configuration

as used in the current work, a simple forward run of the wake model with power predictions for two downstream neighbours over the full prediction horizon requires approximately 0.7 s. This means predictions can be made $1528\times$ faster than real-time at 4 m s^{-1} and $679\times$ faster than real-time at 9 m s^{-1} . The current update rate in the model-predictive controller is fixed at every five discrete time steps; this is equivalent to an update every 67 s at 4 m s^{-1} or 30 s at 9 m s^{-1} . With the current optimiser settings, every re-optimisation step takes about 21 s per wake, which is, respectively, $3.2\times$ or $1.4\times$ faster than real-time for optimising control updates.

This means that the current optimisation set-up realises real-time optimisation for model-predictive wind farm flow control in below-rated conditions. For that, a single processor per wake needs to be available to distribute the optimisation problems. Faster wind speeds require faster optimisation to achieve real-time model-predictive control. This might be within reach with improvements in the numerical algorithm or using a more performant processor.

4.5.7 Preview-enabled look-up table control

Under the realistic wind variations that drive the HKNB and HKNC cases, our simple preview implementation combines the effectiveness of the steady-state optimal yaw offsets with a simple strategy for accounting for wake propagation. The PLUT controller achieves a further increase in power production over the FVW controller, whereas, in the HKNA case, it underperforms significantly. The difference between these cases appears to originate from the magnitude of wind direction changes, where the FVW controller is more flexible to adapt to a broader range of circumstances.

Despite the lack of flexibility, the results demonstrate that a simple preview approach may realise power gains equal to, or greater than a more complex, economic model-predictive controller with limited simulated wake length and prediction horizon. Further refinement is required to maximise the gains that may be achieved by preview-enabled look-up table control and realise consistent performance, avoiding the losses on downstream turbines and for large magnitude wind direction variations.

Such further refinements should consider tuning the preview window to the wind farm layout, where a dependence on the wind direction would allow the preview controller to account for variable turbine spacing along different rows. Additionally, the LUT should be referenced without preview for turbines whose wake does not impinge on downstream rotors, which means that yaw-aligned operation is optimal. These adjustments are already naturally included in the economic model-predictive control optimisation, which may, therefore, provide a foundation for refining preview control.

4.6 Conclusions

A novel distributed, model-based approach to dynamic wind farm flow control is presented with a focus on yaw control for wake redirection. Previous optimisation results with the FVW are extended to economic model-predictive control at the wind-farm scale by parallelising optimisation, connecting individual models into a directed graph network, and incorporating secondary steering effects. The low computational cost enables real-time optimisation in below-rated conditions.

The novel controller is tested in a large-eddy simulation environment and compared against the industry state-of-the-art approach to wake steering, which is based on look-up tables, as well as an extension with wind direction preview. Given two wind farm configurations under synthetic wind direction variations, the FVW controller achieves improvements in power production during wind direction transients. In the simple three-turbine wind farm, equivalent gains are achieved by the PLUT, whereas it underperforms in a ten-turbine subset of the HKN wind farm. Under realistic inflow variations, the PLUT controller yields the largest improvement in power production over the LUT. The FVW yields a smaller power gain because some undesired effects still appear in the control signal. However, in most cases, the FVW controller reduces the increased demand on yaw actuation for wake steering which is advantageous for practical application in large wind farms.

The results with the FVW and PLUT both emphasise the value of including the dynamics of wake propagation for wake steering control. Further refinements in preview-enabled control are worth investigating and perhaps insights from the model-predictive control solutions can guide the development of preview strategies for look-up table controllers.

Improvements in the FVW control strategy could be achieved by considering longer prediction horizons to accommodate wake steering for longer wakes. However, this comes at a significant computational cost for the receding horizon optimisation. The FVW dynamics are a simplified representation of reality, in this case the LES, resulting in model errors that may be minimised. For example, the inclusion of asymmetry in wake steering is also important for maximising the potential gains in wind farm power production.

Lastly, closing the loop with state feedback is an essential next step to realising dynamic yaw control in a realistic setting as it enables adaptation of model parameters to changing environmental conditions. Furthermore, the results should be extended to use realistic forecasting of future inflow conditions.

Conclusions and recommendations

The results in this thesis demonstrate that dynamic wind farm flow control can improve wind farm power production under realistic, time-varying conditions, as well as reduce demands on yaw actuation. The economic model-predictive approach to achieving these gains in power relies on non-linear control optimisation with a simplified flow model for predicting the dynamics of the wind turbine wake. This flow model, being a free-vortex wake model with an actuator-disc representation of the wind turbine, is demonstrated to be suitable for gradient-based control optimisation. This physics-based model is validated under steady and dynamic conditions using data from wind tunnel experiments. A distributed approach to optimisation allows scaling of economic model-predictive control to the wind-farm scale. Experiments in large-eddy simulation show that the novel controller outperforms the state-of-the-art reference controller. Future work should consider extending the dynamic controller with state estimation to close the feedback loop. Control algorithms can be further refined towards real-time control and a hybrid approach combining dynamic and steady model results may treat finite-horizon effects in control optimisation. Finally, the development of realistic simulation environments is essential for the validation of wind farm flow control, combined with a reference wind farm controller that implements wake steering.

Chapter contents

- 5.1 Conclusions123
 - 5.1.1 Flow modelling for control 123
 - 5.1.2 Wake model validation 123
 - 5.1.3 Wind farm flow control 124
- 5.2 Recommendations125
 - 5.2.1 Closed-loop wind farm flow control 125
 - 5.2.2 Real-time control 125
 - 5.2.3 Horizon effects in predictive control 126
 - 5.2.4 Validation of wind farm flow control 126

5.1 Conclusions

The results in this thesis demonstrate that dynamic wind farm flow control can improve wind farm power production under transient inflow, as well as reduce demands on yaw actuation. The economic model-predictive approach to achieving these gains in power relies on non-linear control optimisation with a simplified flow model for predicting the dynamics of the wind turbine wake. The distributed, parallel solution of the optimisation problem is essential for real-time performance and scaling to large wind farms. This is a major step towards practical dynamic flow control implementation in wind farms, contributing to efficient use of the wind resource and existing infrastructure.

One drawback to the non-linear model-predictive control approach is the significant investment in terms of development time and the learning curve required to synthesise such a control strategy. Upon proof-of-concept demonstration, considerable effort will be necessary to make the model and optimisation sufficiently fast and reliable. Additionally, rigorous testing needs to be undertaken to ensure robust behaviour in all conditions which is non-trivial for a controller based on online non-linear optimisation. Therefore, although the results in this dissertation demonstrate the potential benefits of dynamic wind farm control, the development of the current research line to industrial applications is not straightforward.

5.1.1 Flow modelling for control

The free-vortex wake model presented in Chapter 2 is at the core of realising the performance gains in dynamic wake steering. The wake simulation based on an actuator-disc representation of the wind turbine rotor exhibits the characteristic curled wake dynamics for a rotor operating under yaw misalignment. The comparison of a two-dimensional and three-dimensional model further reinforces that the inclusion of curled-wake dynamics is essential for predicting power gains from wake steering.

Gradient calculation based on the adjoint demonstrates the potential for efficient optimisation. The initial results on simple test cases show that the economic model-predictive control framework achieves power gains with dynamic control signals and under time-varying conditions. However, the manual derivation of the adjoint system is cumbersome and prone to errors – it demands a lot of development time and is inflexible to model updates or adjustments. Next to that, the associated demands on memory usage are high for storing all the partial derivatives necessary for constructing the gradient.

5.1.2 Wake model validation

Validation of the free-vortex wake model demonstrates the effectiveness of the control-oriented model for predicting available power in the mid to far wake. In

comparison with three sets of data from wind tunnel experiments, we find a good correspondence in steady state and under dynamic variations in yaw misalignment. These results show the strength of the physics-based modelling approach, although there is a limit to the wake length that can be predicted because of the numerical instabilities that occur.

The actuator-disc model provided a good balance between accuracy and computational complexity. The further model extensions – a rotating actuator disc and a lifting-line model – showed limited improvement, whereas the additional complexity from two- to three-dimensional flow simulation was essential for modelling wake deflection under yaw misalignment. The addition of rotation to the actuator-disc model did not significantly improve accuracy of power predictions, and neither did the modelling of individual blades with a lifting-line model. Initial results indicated that the inclusion of ground effects may contribute to better modelling asymmetry for rotors close to the ground.

5.1.3 Wind farm flow control

Finally, optimisation with the free-vortex wake model is integrated in a distributed wind farm flow control strategy. The economic model-predictive controller is tested in a large-eddy simulation environment and compared with an industry state-of-the-art reference controller based on steady-state look-up tables, as well as an extension with wind direction preview. The novel controller anticipates the wind direction changes and yields improved power production during these transients. It does so with a reduction in the additional demand on yaw actuation from wake steering. The preview-enabled look-up table controller yields further gains in power production under realistic conditions. However, it does not reduce yaw actuator usage compared to the look-up table controller.

The physics-based model allows the generalisation of the control strategy to multiple scenarios with varying wind speeds and wind directions, where data-driven control methods would be limited to the scope of training data. The connection of individual models to a wind farm does require accounting for secondary steering effects, for which an approximation of the cumulative effects of wake steering is proposed. The presented results demonstrate the potential for distributed solution of the wind farm flow control problem, enabling the scaling to large wind farms.

The transition from the initial manual derivation of the adjoint to automatic derivation of the gradients has allowed more flexibility in model development and reduced code complexity, facilitating improvements in computational performance. These improvements enabled the demonstration of the potential for real-time optimisation of yaw misalignments for wake steering in wind farms.

5.2 Recommendations

This thesis has demonstrated the potential gains that arise from accounting for the dynamics of wake propagation in wake steering control for large wind farms. The physics-based model is an attractive approach to dynamic power predictions in the wake of wind turbines operating under yaw misalignment. However, as the development of a complete closed-loop control approach remains an open challenge, we make several recommendations for further research based on the work in this thesis.

Ch. 5

5.2.1 Closed-loop wind farm flow control

The final economic model-predictive controller presented in Chapter 4 is still an open-loop implementation. An essential next step is closed-loop control as state estimation and model adaptation are essential to realistic control performance evaluation. The Ensemble Kalman filter [50, 79, 80] is one possible candidate, but code suitable for automatic differentiation also enables optimisation-based approaches to data assimilation.

Closed-loop control strategies should be developed with realistic measurements and flow predictions. If using only the standard wind turbine measurements of local wind conditions and power, the wind farm flow is poorly observable and the bandwidth of feedback control on the scale of wind farms is limited due to the delays in wake propagation between turbines. Including lidar measurements increases observability and improves the potential performance of a feedback controller. The investment required for lidar installation should be offset against additional gains from predictive wake steering control.

Given the power of feedback in control, perhaps simpler models could be formulated that contain sufficient fundamental dynamics of wake deflection and wake propagation for dynamic wake steering control. The larger model mismatch with reality could be accommodated by means of data assimilation to provide on-line updates of parameters and state estimates. Furthermore, the model adaptation may account for variations in wake dynamics under changing environmental conditions.

5.2.2 Real-time control

The wind farm flow control results demonstrated real-time optimisation performance. However, further developments are required to provide more frequent control updates or potentially include state estimation in the same framework. In addition to the simple use of faster processors, further increases of computational efficiency in the model simulations and optimisation strategy could be achieved by specific algorithm development. For example, the specific design of optimisers for receding horizon control to reuse information from previous time steps.

5.2.3 Horizon effects in predictive control

The economic model-predictive control optimisation demonstrates potential gains in power production during transients. However, the limitations of prediction horizons and simulated wake length also lead to losses compared to the state-of-the-art reference controller for wake steering. A hybrid approach may combine the strengths of both approaches. Particularly, finite-horizon effects in the dynamic model-based approach may be addressed by utilising pre-optimised steady-state controls.

Alternatively, further refinements in preview-enabled look-up table control may account for sufficient wake propagation dynamics to achieve power gains during wind direction transients with yaw offsets that are steady-state optimal. These preview strategies can be guided by the yaw reference signals found through dynamic optimisation and economic model-predictive control.

5.2.4 Validation of wind farm flow control

A major remaining challenge is testing of effectiveness of wind farm flow control strategies in real wind farms through field experiments. Switching between a new control strategy and a baseline controller takes a long time to acquire sufficient data for generating statistically significant results as gains may be small and the wind farm environment is uncertain with a large variability in conditions. Additionally, it is infeasible to be toggling control throughout the entire year for entire farms. Thus, the question remains; how do we discern small improvements in wind farm performance in the real world, where there is no control experiment?

For this reason, simulation studies and wind tunnel experiments need to be combined as they allow repeatable experiments and comparison of control strategies under similar conditions. Specific attention should go towards the replication of realistic atmospheric conditions with variations in wind speed and wind direction for the validation of control strategy performance.

A reliable reference control strategy is necessary for comparing novel controllers against a suitable baseline. Such a reference controller already exists for turbine-level control in the form of the Reference Open Source Controller [132], developed from the Delft Research Controller [133]. On the farm scale, greedy control is no longer a suitable benchmark given the industry adoption of wake steering control. Therefore, a reference wind farm controller should provide a solution to wake steering that is easily adopted to different layouts and environmental conditions.

References

- [1] IPCC, *Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* (Geneva, Switzerland, 2023).
- [2] K. van der Wijst, E. Byers, K. Riahi, R. Schaeffer, and D. van Vuuren, *Data for Figure SPM.5 of Summary for Policymakers of the Synthesis report of Sixth Assessment Report*, Available at: <https://ipcc-browser.ipcc-data.org> (2023).
- [3] IPCC, *Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, 1 (Cambridge University Press, Cambridge, UK and New York, NY, USA, 2022).
- [4] European Commission, *Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions: An EU Strategy to harness the potential of offshore renewable energy for a climate neutral future*, COM(2020) 741 Final, 27 (2020).
- [5] RVO, *Dutch Offshore Wind Market Report*, Tech. Rep. April (2023).
- [6] IRENA, *Renewable power generation costs in 2022* (International Renewable Energy Agency, Abu Dhabi, 2022).
- [7] IRENA, *Renewable Capacity Statistics 2023* (The International Renewable Energy Agency, 2023).
- [8] RystadEnergy and WindEurope, *The State of the European Wind Energy Supply Chain - A «what-would-it-take» analysis of the European supply chain's ability to support ambitious capacity targets towards 2030*, Tech. Rep. April (2023).
- [9] C. Grabow, I. Manak, and D. Ikenson, *The Jones Act: A Burden America Can No Longer Bear*, Policy Anal. **845**, 1 (2018).
- [10] A. Scassola, I. Gabella, J. Lee, and F. Zhao, *Global Wind Workforce Outlook 2022-2026*, Tech. Rep. (Global Wind Energy Council, 2022).

- [11] GEBCO Compilation Group, *GEBCO 2023 Grid*, Available at: <https://download.gebco.net/> (2023).
- [12] NEWA Consortium, *New European Wind Atlas*, Available at: <https://www.neweuropeanwindatlas.eu> (2023).
- [13] P. Veers, K. Dykes, E. Lantz, S. Barth, C. L. Bottasso, O. Carlson, A. Clifton, J. Green, P. Green, H. Holttinen, D. Laird, V. Lehtomäki, J. K. Lundquist, J. Manwell, M. Marquis, C. Meneveau, P. Moriarty, X. Munduate, M. Muskulus, J. Naughton, L. Pao, J. Paquette, J. Peinke, A. Robertson, J. S. Rodrigo, A. M. Sempreviva, J. C. Smith, A. Tuohy, and R. Wiser, *Grand challenges in the science of wind energy*, *Science* **366**, eau2027 (2019).
- [14] A. Betz, *Introduction to the Theory of Flow Machines* (Elsevier, 1966).
- [15] J. W. van Wingerden, P. A. Fleming, T. Göcmen, I. Eguinoa, B. M. Doeke-meijer, K. Dykes, M. Lawson, E. Simley, J. King, D. Astrain, M. Iribas, C. L. Bottasso, J. Meyers, S. Raach, K. Kölle, and G. Giebel, *Expert Elicitation on Wind Farm Control*, *J. Phys. Conf. Ser.* **1618**, 022025 (2020).
- [16] A. C. Kheirabadi and R. Nagamune, *Real-time relocation of floating offshore wind turbine platforms for wind farm efficiency maximization: An assessment of feasibility and steady-state potential*, *Ocean Eng.* **208**, 107445 (2020).
- [17] J. Meyers, C. Bottasso, K. Dykes, P. Fleming, P. Gebraad, G. Giebel, T. Göcmen, and J. W. van Wingerden, *Wind farm flow control: prospects and challenges*, *Wind Energy Sci.* **7**, 2271 (2022).
- [18] D. van der Hoek, S. Kanev, J. Allin, D. Bieniek, and N. Mittelmeier, *Effects of axial induction control on wind farm energy production - A field test*, *Renew. Energy* **140**, 994 (2019).
- [19] J. Bartl and L. Sætran, *Experimental testing of axial induction based control strategies for wake control and wind farm optimization*, *J. Phys. Conf. Ser.* **753**, 032035 (2016).
- [20] J. Annoni, P. M. O. Gebraad, A. K. Scholbrock, P. A. Fleming, and J. W. van Wingerden, *Analysis of axial-induction-based wind plant control using an engineering and a high-order wind plant model*, *Wind Energy* **19**, 1135 (2017).
- [21] J. A. Frederik, R. Weber, S. Cacciola, F. Campagnolo, A. Croce, C. Bottasso, and J. W. van Wingerden, *Periodic dynamic induction control of wind farms:*

- Proving the potential in simulations and wind tunnel experiments*, Wind Energy Sci. **5**, 245 (2020).
- [22] W. Munters and J. Meyers, *Towards practical dynamic induction control of wind farms: Analysis of optimally controlled wind-farm boundary layers and sinusoidal induction control of first-row turbines*, Wind Energy Sci. **3**, 409 (2018).
- [23] J. A. Frederik, B. M. Doekemeijer, S. P. Mulders, and J. W. van Wingerden, *The helix approach: Using dynamic individual pitch control to enhance wake mixing in wind farms*, Wind Energy **23**, 1739 (2020).
- [24] D. van der Hoek, J. Frederik, M. Huang, F. Scarano, C. Simao Ferreira, and J. W. van Wingerden, *Experimental analysis of the effect of dynamic induction control on a wind turbine wake*, Wind Energy Sci. **7**, 1305 (2022).
- [25] D. van der Hoek, B. van den Abbeele, C. S. Ferreira, and J. W. van Wingerden, *Maximizing wind farm power output with the helix approach – experimental validation and wake analysis using tomographic PIV*, (2023), arXiv:2306.12849 .
- [26] M. F. Howland, J. Bossuyt, L. A. Martínez-Tossas, J. Meyers, and C. Meneveau, *Wake structure in actuator disk models of wind turbines in yaw under uniform inflow conditions*, J. Renew. Sustain. Energy **8**, 043301 (2016).
- [27] M. Bastankhah and F. Porté-Agel, *Experimental and theoretical study of wind turbine wakes in yawed conditions*, J. Fluid Mech. **806**, 506 (2016).
- [28] J. Bartl, F. Mühle, J. Schottler, L. Sætran, J. Peinke, M. Adaramola, and M. Hölling, *Wind tunnel experiments on wind turbine wakes in yaw: Effects of inflow turbulence and shear*, Wind Energy Sci. **3**, 329 (2018).
- [29] P. Fleming, J. Annoni, L. A. Martínez-Tossas, S. Raach, K. Gruchalla, A. Scholbrock, M. Churchfield, and J. Roadman, *Investigation into the shape of a wake of a yawed full-scale turbine*, J. Phys. Conf. Ser. **1037**, 032010 (2018).
- [30] P. Hulsman, M. Wosnik, V. Petrović, M. Hölling, and M. Kühn, *Development of a curled wake of a yawed wind turbine under turbulent and sheared inflow*, Wind Energy Sci. **7**, 237 (2022).
- [31] P. M. O. Gebraad, F. W. Teeuwisse, J. W. van Wingerden, P. A. Fleming, S. D. Ruben, J. R. Marden, and L. Y. Pao, *Wind plant power optimization through yaw control using a parametric model for wake effects—a CFD simulation study*, Wind Energy **19**, 95 (2016).

- [32] M. T. van Dijk, J. W. van Wingerden, T. Ashuri, and Y. Li, *Wind farm multi-objective wake redirection for optimizing power production and loads*, *Energy* **121**, 561 (2017).
- [33] F. Campagnolo, V. Petrović, J. Schreiber, E. M. Nanos, A. Croce, and C. L. Bottasso, *Wind tunnel testing of a closed-loop wake deflection controller for wind farm power maximization*, *J. Phys. Conf. Ser.* **753**, 032006 (2016).
- [34] M. Bastankhah and F. Porté-Agel, *Wind farm power optimization via yaw angle control: A wind tunnel study*, *J. Renew. Sustain. Energy* **11**, 023301 (2019).
- [35] F. Campagnolo, R. Weber, J. Schreiber, and C. L. Bottasso, *Wind tunnel testing of wake steering with dynamic wind direction changes*, *Wind Energy Sci.* **5**, 1273 (2020).
- [36] M. F. Howland, S. K. Lele, and J. O. Dabiri, *Wind farm power optimization through wake steering*, *Proc. Natl. Acad. Sci. U. S. A.* **116**, 14495 (2019).
- [37] P. Fleming, J. King, K. Dykes, E. Simley, J. Roadman, A. Scholbrock, P. Murphy, J. K. Lundquist, P. Moriarty, K. Fleming, J. van Dam, C. Bay, R. Mudafort, H. Lopez, J. Skopek, M. Scott, B. Ryan, C. Guernsey, and D. Brake, *Initial results from a field campaign of wake steering applied at a commercial wind farm – Part 1*, *Wind Energy Sci.* **4**, 273 (2019).
- [38] P. Fleming, J. King, E. Simley, J. Roadman, A. Scholbrock, P. Murphy, J. K. Lundquist, P. Moriarty, K. Fleming, J. van Dam, C. Bay, R. Mudafort, D. Jager, J. Skopek, M. Scott, B. Ryan, C. Guernsey, and D. Brake, *Continued results from a field campaign of wake steering applied at a commercial wind farm - Part 2*, *Wind Energy Sci.* **5**, 945 (2020).
- [39] P. Fleming, M. Sinner, T. Young, M. Lannic, J. King, E. Simley, and B. M. Doekemeijer, *Experimental results of wake steering using fixed angles*, *Wind Energy Sci.* **6**, 1521 (2021).
- [40] B. M. Doekemeijer, S. Kern, S. Maturu, S. Kanev, B. Salbert, J. Schreiber, F. Campagnolo, C. L. Bottasso, S. Schuler, F. Wilts, T. Neumann, G. Potenza, F. Calabretta, F. Fioretti, and J. W. van Wingerden, *Field experiment for open-loop yaw-based wake steering at a commercial onshore wind farm in Italy*, *Wind Energy Sci.* **6**, 159 (2021).
- [41] E. Simley, P. Fleming, N. Girard, L. Alloin, E. Godefroy, and T. Duc, *Results from a wake-steering experiment at a commercial wind plant: Investigating*

- the wind speed dependence of wake-steering performance*, Wind Energy Sci. **6**, 1427 (2021).
- [42] M. F. Howland, J. B. Quesada, J. J. P. Martínez, F. P. Larrañaga, N. Yadav, J. S. Chawla, V. Sivaram, and J. O. Dabiri, *Collective wind farm operation based on a predictive model increases utility-scale energy production*, Nat. Energy **7**, 818 (2022).
- [43] Siemens-Gamesa Renewable Energy, *Wake Adapt – Reducing wake losses*, (2023).
- [44] E. Simley, D. Millstein, S. Jeong, and P. Fleming, *The value of wake steering wind farm control in U.S. energy markets*, Wind Energy Sci. Discuss. **2023**, 1 (2023).
- [45] D. Kumar, M. A. Rotea, E. J. Aju, and Y. Jin, *Wind plant power maximization via extremum seeking yaw control: A wind tunnel experiment*, Wind Energy **26**, 283 (2023).
- [46] J. Liew, T. Göçmen, W. H. Lio, and G. C. Larsen, *Model-free closed-loop wind farm control using reinforcement learning with recursive least squares*, Wind Energy , 1 (2023).
- [47] NREL, *FLORIS. Version 3.0*, Available at: <https://github.com/NREL/floris> (2022).
- [48] M. Bastankhah and F. Porté-Agel, *A new analytical model for wind-turbine wakes*, Renew. Energy **70**, 116 (2014).
- [49] L. A. Martínez-Tossas, J. Annoni, P. A. Fleming, and M. J. Churchfield, *The aerodynamics of the curled wake: a simplified model in view of flow control*, Wind Energy Sci. **4**, 127 (2019).
- [50] B. M. Doekemeijer, D. van der Hoek, and J. W. van Wingerden, *Closed-loop model-based wind farm control using FLORIS under time-varying inflow conditions*, Renew. Energy **156**, 719 (2020).
- [51] S. Kanev, *Dynamic wake steering and its impact on wind farm power production and yaw actuator duty*, Renew. Energy **146**, 9 (2020).
- [52] C. R. Shapiro, J. Meyers, C. Meneveau, and D. F. Gayme, *Wind farms providing secondary frequency regulation: Evaluating the performance of model-based receding horizon control*, Wind Energy Sci. **3**, 11 (2018).

- [53] G. A. van Kuik, J. Peinke, R. Nijssen, D. Lekou, J. Mann, J. N. Sørensen, C. Ferreira, J. W. van Wingerden, D. Schlipf, P. Gebraad, H. Polinder, A. Abrahamson, G. J. van Bussel, J. D. Sørensen, P. Tavner, C. L. Bottasso, M. Muskulus, D. Matha, H. J. Lindeboom, S. Degraer, O. Kramer, S. Lehnhoff, M. Sonnenschein, P. E. Sørensen, R. W. Künneke, P. E. Morthorst, and K. Skytte, *Long-term research challenges in wind energy – A research agenda by the European Academy of Wind Energy*, Wind Energy Sci. **1**, 1 (2016).
- [54] P. M. Gebraad and J. W. van Wingerden, *A control-oriented dynamic model for wakes in wind plants*, J. Phys. Conf. Ser. **524**, 012186 (2014).
- [55] M. Lejeune, M. Moens, and P. Chatelain, *A Meandering-Capturing Wake Model Coupled to Rotor-Based Flow-Sensing for Operational Wind Farm Flow Prediction*, Front. Energy Res. **10**, 884068 (2022).
- [56] M. Becker, B. Ritter, D. van der Hoek, U. Konigorski, D. Allaerts, and J. W. van Wingerden, *The revised FLORIDyn model: implementation of heterogeneous flow and the Gaussian wake*, Wind Energy Sci. **7**, 2163 (2022).
- [57] J. P. Goit and J. Meyers, *Optimal control of energy extraction in wind-farm boundary layers*, J. Fluid Mech. **768**, 5 (2015).
- [58] J. P. Goit, W. Munters, and J. Meyers, *Optimal coordinated control of power extraction in LES of a wind farm with entrance effects*, Energies **9**, 29 (2016).
- [59] W. Munters and J. Meyers, *Dynamic Strategies for Yaw and Induction Control of Wind Farms Based on Large-Eddy Simulation and Optimization*, Energies **11**, 177 (2018).
- [60] N. Janssens and J. Meyers, *Towards real-time optimal control of wind farms using large-eddy simulations*, Wind Energy Sci. Discuss. **2023**, 1 (2023).
- [61] S. Boersma, B. M. Doekemeijer, M. Vali, J. Meyers, and J. W. van Wingerden, *A control-oriented dynamic wind farm model: WFSim*, Wind Energy Sci. **3**, 75 (2018).
- [62] M. J. van den Broek and J. W. van Wingerden, *Dynamic Flow Modelling for Model-Predictive Wind Farm Control*, J. Phys. Conf. Ser. **1618**, 022023 (2020).
- [63] M. J. van den Broek, B. Sanderse, and J. W. van Wingerden, *Flow Modelling for Wind Farm Control: 2D vs . 3D*, J. Phys. Conf. Ser. **2265**, 032086 (2022).
- [64] M. J. van den Broek, *FRED - Framework for wind farm flow Regulation and Estimation with Dynamics*, Available at: <https://github.com/TUdelft-DataDrivenControl/FRED> (2021).

- [65] J. W. van Wingerden, L. Pao, J. Aho, and P. Fleming, *Active Power Control of Waked Wind Farms*, IFAC Pap. **50**, 4484 (2017).
- [66] M. Vali, V. Petrović, S. Boersma, J. W. van Wingerden, L. Y. Pao, and M. Kühn, *Adjoint-based model predictive control for optimal energy extraction in waked wind farms*, Control Eng. Pract. **84**, 48 (2019).
- [67] J. G. Leishman, *Principles of Helicopter Aerodynamics* (Cambridge University Press, 2000).
- [68] J. Katz and A. Plotkin, *Low-Speed Aerodynamics*, 2nd ed. (Cambridge University Press, 2001).
- [69] M. J. Bhagwat and J. G. Leishman, *Transient rotor inflow using a time-accurate free-vortex wake model*, in *39th Aerosp. Sci. Meet. Exhib.* (2001) p. 993.
- [70] J. G. Leishman, *Challenges in modeling the unsteady aerodynamics of wind turbines*, Wind Energy **5**, 85 (2002).
- [71] F. J. Simoes and J. M. Graham, *Application of a free vortex wake model to a horizontal axis wind turbine*, J. Wind Eng. Ind. Aerodyn. **39**, 129 (1992).
- [72] J. B. de Vaal, M. O. Hansen, and T. Moan, *Validation of a vortex ring wake model suited for aeroelastic simulations of floating wind turbines*, J. Phys. Conf. Ser. **555**, 012025 (2014).
- [73] K. Brown, D. Houck, D. Maniaci, and C. Westergaard, *Rapidly recovering wind turbine wakes with dynamic pitch and rotor speed control*, in *AIAA Scitech 2021 Forum* (2021) p. 1182.
- [74] J. Dong, A. Viré, C. S. Ferreira, Z. Li, and G. van Bussel, *A modified free wake vortex ring method for horizontal-axis wind turbines*, Energies **12**, 3900 (2019).
- [75] H. Lee and D. J. Lee, *Effects of platform motions on aerodynamic performance and unsteady wake evolution of a floating offshore wind turbine*, Renew. Energy **143**, 9 (2019).
- [76] M. J. van den Broek, D. van den Berg, B. Sanderse, and J. W. van Wingerden, *Optimal Control for Wind Turbine Wake Mixing on Floating Platforms*, IFAC Pap. **56**, 7656 (2023).

- [77] T. Berdowski, C. S. Ferreira, A. van Zuijlen, and G. van Bussel, *Three-dimensional free-wake vortex simulations of an actuator disc in yaw*, in *2018 Wind Energy Symp.* (2018) p. 0513.
- [78] L. Grüne and J. Pannek, *Nonlinear Model Predict. Control*, 2nd ed. (Springer, 2017).
- [79] M. F. Howland, A. S. Ghate, S. K. Lele, and J. O. Dabiri, *Optimal closed-loop wake steering - Part 1: Conventionally neutral atmospheric boundary layer conditions*, *Wind Energy Sci.* **5**, 1315 (2020).
- [80] M. Becker, D. Allaerts, and J. W. van Wingerden, *Ensemble-Based Flow Field Estimation Using the Dynamic Wind Farm Model FLORIDyn*, *Energies* **15**, 8589 (2022).
- [81] M. J. van den Broek, D. De Tavernier, B. Sanderse, and J. W. van Wingerden, *Adjoint optimisation for wind farm flow control with a free-vortex wake model*, *Renew. Energy* **201**, 752 (2022).
- [82] M. J. van den Broek, *Free-vortex wake with discrete adjoint – model code*, Available at: <https://github.com/TUdelft-DataDrivenControl/vortexwake> or <https://doi.org/10.4121/2027862> (2022).
- [83] M. J. van den Broek, *Free-vortex wake with discrete adjoint - dataset*, Available at: <https://doi.org/10.4121/20278620> (2022).
- [84] G. C. Larsen, H. Madsen Aagaard, F. Bingöl, J. Mann, S. Ott, J. N. Sørensen, V. Okulov, N. Troldborg, N. M. Nielsen, K. Thomsen, T. J. Larsen, and R. Mikkelsen, *Dynamic wake meandering modeling* (2007).
- [85] D. R. Houck, N. DeVelder, and C. L. Kelley, *Comparison of a mid-fidelity free vortex wake method to a high-fidelity actuator line model large eddy simulation for wind turbine wake simulations*, *J. Phys. Conf. Ser.* **2265**, 042044 (2022).
- [86] T. Burton, D. Sharpe, N. Jenkins, and E. Bossanyi, *Wind Energy Handbook* (Wiley, 2001).
- [87] L. Masset, O. Bröls, and G. Kerschen, *Partition of the circle in cells of equal area and shape*, (2011).
- [88] R. Dorfman, P. A. Samuelson, and R. M. Solow, *Linear Programming and Economic Analysis* (McGraw-Hill, New York, 1958).

- [89] S. Gros, *An Economic NMPC Formulation for Wind Turbine Control*, in *Conf. Decis. Control* (IEEE, 2013).
- [90] L. Grüne, *Economic receding horizon control without terminal constraints*, *Automatica* **49**, 725 (2013).
- [91] T. Lauß, S. Oberpeilsteiner, W. Steiner, and K. Nachbagauer, *The Discrete Adjoint Gradient Computation for Optimization Problems in Multibody Dynamics*, *J. Comput. Nonlinear Dyn.* **12**, 031016 (2017).
- [92] R. Byrd, P. Lu, J. Nocedal, and C. Zhu, *A limited memory algorithm for bound constrained optimization*, *J. Sci. Comput.* **16**, 1190 (1995).
- [93] D. P. Kingma and J. L. Ba, *Adam: A method for stochastic optimization*, in *Int. Conf. Learn. Represent.* (2015) pp. 1–15, arXiv:1412.6980 .
- [94] U. Ciri, M. A. Rotea, and S. Leonardi, *Model-free control of wind farms: A comparative study between individual and coordinated extremum seeking*, *Renew. Energy* **113**, 1033 (2017).
- [95] C. Cossu, *Wake redirection at higher axial induction*, *Wind Energy Sci.* **6**, 377 (2021).
- [96] M. J. van den Broek, D. De Tavernier, P. Hulsman, D. van der Hoek, B. Sanderse, and J. W. van Wingerden, *Free-vortex models for wind turbine wakes under yaw misalignment – a validation study on far-wake effects*, *Wind Energy Sci.* **8**, 1909 (2023).
- [97] M. J. van den Broek, *Simulation data and code underlying the publication: Free-vortex models for wind turbine wakes under yaw misalignment – a validation study on far-wake effects.* (2023).
- [98] P. Hulsman, M. Wosnik, V. Petrović, M. Hölling, and M. Kühn, *Data Supplement for 'Curled Wake Development of a Yawed Wind Turbine at Turbulent and Sheared Inflow' - Wind Energy Science Journal*, (2021).
- [99] D. van der Hoek, M. J. van den Broek, and J. W. van Wingerden, *Data underlying the publication: Free-vortex models for wind turbine wakes under yaw misalignment - a validation study on far-wake effects*, (2023).
- [100] P. Hulsman, V. Petrović, M. Wosnik, and M. Kühn, *Evaluating open-loop wake steering controllers by replicating dynamic wind direction changes in the wind tunnel*, (2023), (in preparation).

- [101] E. Branlard, L. A. Martínez-Tossas, and J. Jonkman, *A time-varying formulation of the curled wake model within the FAST.Farm framework*, Wind Energy **26**, 44 (2023).
- [102] W. Munters and J. Meyers, *Optimal dynamic induction and yaw control of wind farms: Effects of turbine spacing and layout*, J. Phys. Conf. Ser. **1037**, 032015 (2018).
- [103] T. Sebastian and M. Lackner, *Analysis of the Induction and Wake Evolution of an Offshore Floating Wind Turbine*, Energies **5**, 968 (2012).
- [104] M. Jeon, S. Lee, T. Kim, and S. Lee, *Wake influence on dynamic load characteristics of offshore floating wind turbines*, AIAA J. **54**, 3535 (2016).
- [105] M. J. van den Broek, D. van den Berg, B. Sanderse, and J. W. van Wingerden, *Optimal Control for Wind Turbine Wake Mixing on Floating Platforms*, IFAC Pap. **56**, 7656 (2023).
- [106] G. A. M. van Kuik, *The Fluid Dynamic Basis for Actuator Disc and Rotor Theories* (2018).
- [107] H. B. Squire, *The Growth of a Vortex in Turbulent Flow*, Aeronaut. Q. **16**, 302 (1965).
- [108] P. Hulsman, C. Sucameli, V. Petrović, A. Rott, A. Gerds, and M. Kühn, *Turbine power loss during yaw-misaligned free field tests at different atmospheric conditions*, J. Phys. Conf. Ser. **2265**, 032074 (2022).
- [109] J. Schottler, A. Hölling, J. Peinke, and M. Hölling, *Wind tunnel tests on controllable model wind turbines in yaw*, in *34th Wind Energy Symp.* (AIAA, 2016).
- [110] J. Schottler, J. Bartl, F. Mühle, L. Sætran, J. Peinke, and M. Hölling, *Wind tunnel experiments on wind turbine wakes in yaw: Redefining the wake width*, Wind Energy Sci. **3**, 257 (2018).
- [111] V. Petrović, J. Schottler, I. Neunaber, M. Hölling, and M. Kühn, *Wind tunnel validation of a closed loop active power control for wind farms*, J. Phys. Conf. Ser. **1037**, 032020 (2018).
- [112] M. Lin and F. Porté-Agel, *Large-eddy simulation of a wind-turbine array subjected to active yaw control*, Wind Energy Sci. **7**, 2215 (2022).

- [113] M. J. van den Broek, M. Becker, B. Sanderse, and J. W. van Wingerden, *Dynamic wind farm flow control using free-vortex wake models*, Wind Energy Sci. Discuss. **2023**, 1 (2023).
- [114] M. J. van den Broek, *Simulation data and code accompanying the publication: Dynamic wind farm flow control using free-vortex wake models*, (2023).
- [115] B. A. M. Sengers, M. Zech, P. Jacobs, G. Steinfeld, and M. Kühn, *A physically interpretable data-driven surrogate model for wake steering*, Wind Energy Sci. **7**, 1455 (2022).
- [116] J. King, P. Fleming, R. King, L. A. Martínez-Tossas, C. J. Bay, R. Mudafort, and E. Simley, *Control-oriented model for secondary effects of wake steering*, Wind Energy Sci. **6**, 701 (2021).
- [117] E. Simley, P. Fleming, J. King, and M. Sinner, *Wake Steering Wind Farm Control with Preview Wind Direction Information*, Proc. Am. Control Conf. **2021-May**, 1783 (2021).
- [118] B. A. M. Sengers, A. Rott, E. Simley, M. Sinner, G. Steinfeld, and M. Kühn, *Increased power gains from wake steering control using preview wind direction information*, Wind Energy Sci. **8**, 1693 (2023).
- [119] Z. Li and X. Yang, *Large-eddy simulation on the similarity between wakes of wind turbines with different yaw angles*, J. Fluid Mech. **921**, 1 (2021).
- [120] K. S. Heck, H. M. Johlas, and M. F. Howland, *Modelling the induction, thrust and power of a yaw-misaligned actuator disk*, J. Fluid Mech. **959**, 1 (2023).
- [121] J. Liew, A. M. Urbán, and S. J. Andersen, *Analytical model for the power-yaw sensitivity of wind turbines operating in full wake*, Wind Energy Sci. **5**, 427 (2020).
- [122] P. Fleming, J. Annoni, M. Churchfield, L. A. Martinez-Tossas, K. Gruchalla, M. Lawson, and P. Moriarty, *A simulation study demonstrating the importance of large-scale trailing vortices in wake steering*, Wind Energy Sci. **3**, 243 (2018).
- [123] G. M. Starke, P. Stanfel, C. Meneveau, D. F. Gayme, and J. King, *Network based estimation of wind farm power and velocity data under changing wind direction*, in *2021 Am. Control Conf. (AACC, 2021)* pp. 1803–1810.
- [124] K. Shaler, J. Jasa, and G. E. Barter, *Efficient Loads Surrogates for Waked Turbines in an Array*, J. Phys. Conf. Ser. **2265**, 032095 (2022).

- [125] E. Bossanyi, *Surrogate model for fast simulation of turbine loads in wind farms*, J. Phys. Conf. Ser. **2265**, 042038 (2022).
- [126] K. A. Kragh and P. A. Fleming, *Rotor Speed Dependent Yaw Control of Wind Turbines Based on Empirical Data*, in *50th AIAA Aerosp. Sci. Meet.* (AIAA, 2012).
- [127] P. A. Fleming, A. P. Stanley, C. J. Bay, J. King, E. Simley, B. M. Doekemeijer, and R. Mudafort, *Serial-Refine Method for Fast Wake-Steering Yaw Optimization*, J. Phys. Conf. Ser. **2265**, 032109 (2022).
- [128] C. Bak, R. Bitsche, A. Yde, T. Kim, M. H. Hansen, F. Zahle, M. Gaunaa, J. Blasques, M. Døssing, J. J. W. Heinen, and T. Behrens, *Light rotor: The 10-MW reference wind turbine*, Eur. Wind Energy Conf. Exhib. 2012 **1**, 532 (2012).
- [129] KNMI, *Wind – lidar wind profiles measured at North Sea wind farm TenneT platforms 1 second raw data*, Available at: <https://dataplatfom.knmi.nl/dataset/windlidar-nz-wp-platform-1s-1>, last accessed: 2023-06-29 (2023).
- [130] M. J. Churchfield, S. Lee, J. Michalakes, and P. J. Moriarty, *A numerical study of the effects of atmospheric and wake turbulence on wind turbine dynamics*, J. Turbul. **13**, N14 (2012).
- [131] C. R. Shapiro, G. M. Starke, C. Meneveau, and D. F. Gayme, *A wake modeling paradigm for wind farm design and control*, Energies **12**, 2956 (2019).
- [132] N. J. Abbas, D. S. Zalkind, L. Pao, and A. Wright, *A reference open-source controller for fixed and floating offshore wind turbines*, Wind Energy Sci. **7**, 53 (2022).
- [133] S. P. Mulders and J. W. van Wingerden, *Delft Research Controller: An open-source and community-driven wind turbine baseline controller*, J. Phys. Conf. Ser. **1037**, 032009 (2018).

Acknowledgments

Four years (and a few months) ago, I embarked on the PhD adventure after finding during my master's thesis that I was enjoying research and maybe not entirely done with university life. It's been a winding road that's led through valleys and up peaks, where the right turn to take was not always clear. However, the path has been rewarding and it's good to be able to look back at the experiences of the past years, wrap up the work in this dissertation, and prepare for the next step. Now is also the time to give thanks to some people for their role in this journey.

First of all, none of this would have been possible without Jan-Willem. From the moment we first talked I was impressed by the energy you put into your work and our group. I am still amazed by the boundless pool of ideas floating around in your mind. Keep connecting people and letting up small balloons – who knows what they'll grow into. It's been a pleasure working with you these four years.

Benjamin, your rigour and critical questions have been of tremendous value in shaping my research. I admire the thoroughness with which you approach problems and the balance you create between work and life. Delphine, your aerodynamic expertise and introduction to free-vortex methods have been essential to make this dissertation possible. Paul, it's been good to collaborate with you; I wish you the best for your own PhD.

The wind energy group has gotten so big that we have almost taken over the second floor of the department. Marcus, you have been an invaluable sparring partner. I have enjoyed our discussions on wind farm flow control as well as your critical eye for graphic design. Daniel, thanks for your contributions on floating wind and your grumpy rants on various topics. Keep running strong! Daan, for paving the way towards the dissertation and sharing the struggles of experimental campaigns. Livia and Atin, thanks for sharing the troubles and joys of research as we worked towards our final dissertations. Jonas and Emanuel, be sure to keep bothering the control engineers about CFL numbers! And many thanks to the rest of the wind group – Bart, Joeri, Sebastiaan, Unai, David, Amr, Claudia, Mees, Rogier, Guido, Marion, Tim, Matteo, Bert, Jean, Zhixin, Jesse, Uwe – for the collaborations, discussions, and coffee breaks.

Of course, the rest of the department should not be forgotten. Coen, Roger, Jacques, Sander; I might have to come back around for a game of table football. Alex, Clara, Eva, Frida, Emilio, Paul, Thomas, Suad, Leila, Pascal, Leonore, Fred-

erick, and everybody else; thanks for the positive vibes, creative ideas, chats over lunch or coffee, and the Friday night drinks at Bebop.

Thanks to my friends outside of work for reminding me that life is more than just the PhD and your contribution to keeping me (somewhat) sane throughout all this. B70, EJD, EJJ, Oud Trek, Karen, Hennie, the RA-B runners, Scouting VWP; you've all had your part!

Finally, I would like to thank my parents, Jan and Elma, who have always been there for me; my little brother, Niels, for the food in particular; Ruud, Teuny, and Karen, for their interest and support; and Vera, for being by my side for the highs and lows of the past years.

*Maarten Jan van den Broek
Rotterdam, January 2024*

Curriculum vitæ

Maarten Jan van den Broek

- 2019–2023 Doctoral degree in Control Engineering
Delft University of Technology, the Netherlands
Thesis: *Dynamic wind farm flow control using free-vortex wake models*
Promotor: prof. dr. ir. Jan-Willem van Wingerden
Copromotor: dr. ir. Benjamin Sanderse
- 2014–2019 Master of Science in Mechanical Engineering
Delft University of Technology, the Netherlands
Thesis: *Fast self-stable planar bipedal running*
Biomechanical engineering track with specialisation biorobotics
- 2011–2014 Bachelor of Science in Aerospace Engineering (*cum laude*)
Delft University of Technology, the Netherlands
- 2009–2011 Pre-University College
Leiden University, the Netherlands
- 2005–2011 Highschool – VWO (*cum laude*)
Alfrink College, Zoetermeer, the Netherlands
Internation Baccalaureate English
- 1993 Born on 29 September 1993
Zoetermeer, the Netherlands

List of publications

Journal publications

-  10. Maarten J. van den Broek, Marcus Becker, Benjamin Sanderse, Jan-Willem van Wingerden: *Dynamic wind farm flow control using free-vortex wake models*, Wind Energy Science Discussions, in review, 2023
-  9. Maarten J. van den Broek, Delphine De Tavernier, Paul Hulsman, Daan van der Hoek, Benjamin Sanderse, Jan-Willem van Wingerden: *Free-vortex models for wind turbine wakes under yaw misalignment – a validation study on far-wake effects*, Wind Energy Science, 8, 1909–1925, 2023
-   8. Maarten J. van den Broek, Delphine De Tavernier, Benjamin Sanderse, Jan-Willem van Wingerden: *Adjoint optimisation for wind farm flow control with a free-vortex wake model*, Renewable Energy, 201(November), 752–765, 2022.
- 7. Tuhfe Göçmen, Filippo Campagnolo, Thomas Duc, Irene Eguinoa, Søren Juhl Andersen, Vlaho Petrović, Lejla Imširović, Robert Braunbehrens, Jaime Liew, Mads Baungaard, Maarten Paul van der Laan, Guowei Qian, Maria Aparicio-Sanchez, Rubén González-Lope, Vinit V. Dighe, Marcus Becker, Maarten J. van den Broek, Jan-Willem van Wingerden, Adam Stock, Matthew Cole, Renzo Ruisi, Ervin Bossanyi, Niklas Requate, Simon Strnad, Jonas Schmidt, Lukas Vollmer, Ishaan Sood, and Johan Meyers: *FarmConnors wind farm flow control benchmark – Part 1: Blind test results*, Wind Energy Science, 7, 1791–1825, 2022

 Included in this thesis.

 Awarded TU Delft *Best Energy Paper* and *Best Wind Energy Paper* of 2022.

Conferences

- 6. Maarten J. van den Broek, Marcus Becker, Benjamin Sanderse, Jan-Willem van Wingerden: *Model-predictive control for dynamic wake steering in large wind farms*. Oral presentation at NAWEA WindTech 2023.
- » 5. Maarten J. van den Broek, Daniel van den Berg, Benjamin Sanderse, Jan-Willem van Wingerden: *Optimal control for wind turbine wake mixing on floating platforms*. IFAC-PapersOnLine, 56(2):7656–7661, 2023 Oral presentation at IFAC World Congress 2023.
- 4. Maarten J. van den Broek, Benjamin Sanderse, Jan-Willem van Wingerden: *Dynamic wind farm flow control with a free-vortex wake model*. Oral presentation at WESC 2023.
- 3. Maarten J. van den Broek, Benjamin Sanderse, Jan-Willem van Wingerden: *Free-vortex wake model with discrete adjoint for wind farm flow control*. Oral presentation at NAWEA WindTech 2022.
- » 2. Maarten J. van den Broek, Benjamin Sanderse, Jan-Willem van Wingerden: *Flow modelling for wind farm control: 2D vs. 3D*. Journal of Physics Conference Series, 2265(3):032086, 2022. Oral presentation at TORQUE 2022.
- » 1. Maarten J. van den Broek, Jan-Willem van Wingerden: *Dynamic flow modelling for model-predictive wind farm control*. Journal of Physics Conference Series, 1618(2):022023, 2020. Oral presentation at TORQUE 2020.

» Paper published in conference proceedings.