

DATA-DRIVEN WIND PLANT CONTROL

DATA-DRIVEN WIND PLANT CONTROL

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K. C. A. M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op maandag 1 december 2014 om 15:00 uur

door

Pieter Marinus Otto GEBRAAD

werktuigbouwkundig ingenieur
geboren te Capelle aan den IJssel, Nederland.

Dit proefschrift is goedgekeurd door de promotor:

prof. dr. ir. M. Verhaegen

Copromotor: dr. ir. J.W. van Wingerden

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
prof. dr. ir. M. Verhaegen,	Technische Universiteit Delft, promotor
dr. ir. J.W. van Wingerden,	Technische Universiteit Delft, copromotor

Onafhankelijke leden:

prof. dr. ir. J. Meyers,	Katholieke Universiteit Leuven
prof. dr. ir. M.R. de Baar,	Technische Universiteit Eindhoven
prof. dr. ir. G.A.M. van Kuik,	Technische Universiteit Delft
dr. T. Knudsen,	Aalborg University
dr. P.A. Fleming,	National Renewable Energy Laboratory
prof.dr.ir. J. Hellendoorn,	Technische Universiteit Delft, reservelid



The work presented in this thesis was supported by the Far Large Offshore Wind (FLOW) project 'Offshore wind power plant control for minimal loading'.



Printed by: Sieca Repro

Front & Back: Front cover picture of results of the simulation of the flow field in a wind plant, courtesy of Matt Churchfield and Sang Lee.
Back cover photo of Windpark Distripark by Rien Gebraad.

Copyright © 2014 by P.M.O. Gebraad

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

*Ziezo, daar staan ze, mijn stellingen.
Ferm op de poten, voor ieder te kijk en niet omver te werpen.
Durf er eens aan te komen.*

uit 'De televisie, haar werking en haar nut',
een verhaal uit de bundel 'Alle dagen feest' van

Remco Campert

CONTENTS

1	Introduction	1
1.1	Motivation	1
1.2	Background and state-of-the-art	3
1.3	Thesis objectives	14
1.4	Methodologies	15
1.5	Thesis contributions and outline	16
2	Evaluation of control degrees-of-freedom for wind plant control	19
2.1	Introduction	20
2.2	Simulation scenarios	20
2.3	Evaluation of axial-induction-based methods.	24
2.4	Evaluation of wake redirection and turbine repositioning methods	33
2.5	Conclusions.	46
3	Axial-induction-based optimization control for wind plants	49
3.1	Introduction	49
3.2	Maximum power-point tracking control for wind plants	52
3.3	Benchmark wind plant control algorithm with a game theoretic approach	59
3.4	Wind plant simulation model	59
3.5	Simulation examples	64
3.6	Conclusions.	70
4	Yaw-based optimization control for wind plants	75
4.1	Introduction	75
4.2	Characterizing wake effects through simulations in SOWFA, a high-fidelity CFD wind plant simulator.	78
4.3	Data-driven parametric wind plant model: FLORIS.	80
4.4	Wind plant yaw optimization using a game-theoretic approach	90
4.5	Wind plant yaw optimization simulation examples	91
4.6	Conclusions.	99
5	A control-oriented dynamic model for wakes in wind plants	101
5.1	Introduction	101
5.2	FLORIDyn model description	103
5.3	Simulation case study.	115
5.4	Using a Kalman filter to correct the predicted velocity field	117
5.5	Application of the FLORIDYN model for optimized yaw control	124
5.6	Discussion, conclusions and future work	128
6	Conclusions and Recommendations	131
6.1	Conclusions.	131
6.2	Recommendations	133

Bibliography	137
Summary	147
Samenvatting	151
List of Publications	155
Curriculum Vitae	157

1

INTRODUCTION

1.1. MOTIVATION

Wind energy offers a potential to reduce carbon emissions and mitigate worldwide climate change, although a faster wind energy deployment is needed in order to avoid dangerous global warming (Barthelmie et al., 2014). Also, a further growth of wind energy offers political and economical advantages for countries, by having the potential to reduce their reliance on the import of energy.¹

While on a global level, there are no insurmountable technical constraints hindering a further growth of wind energy (Edenhofer et al., 2011), on a local level, policy makers are faced with growth-limiting factors that include limited financial means of private and public parties after the recent global economic crisis, the competitive cost-levels of energy from fossil fuels, and in many areas, limitations on the available appropriate locations for the placement of large wind turbines.

The Netherlands is a country where each of these growth-limiting factors play a role. After the recent financial crisis, it became more difficult to develop wind energy projects in The Netherlands, since external investors demand a higher investment participation from project owners, and the focus of government subsidy programs shifted towards reducing the cost of energy before making large investments in wind energy capacity. While the cost of energy generated with onshore wind turbines is becoming competitive with that from fossil fuel-fired plants, the available locations for onshore wind turbines are becoming scarce. The future development of Dutch wind energy will therefore for a large part take place in the sea: the government plans to realize an offshore wind energy capacity of 3,450MW, by 2023. Currently however, offshore wind energy is about

¹The European Union's (EU's) import of energy currently amounts to 53% of its consumption, and this import represents a value of more than 1 billion euros a day. This makes the EU economy vulnerable to fuel-price volatility and disruptions of the supply, which may have natural or political causes. In order to reduce this vulnerability, the European Commission recommends the EU member states to increase their own energy production, pointing out wind energy as a form of renewable energy that is increasingly competitive with other energy sources (European Commission, 2014).

60% more expensive than onshore wind energy, and in the heavily-used Dutch part of the North Sea, the space available for wind energy is limited. The general policy, both offshore and onshore, is that new wind turbines are placed in clusters (wind plants). (De Boer, 2013; Int. Energy Agency, 2014; Ministerie IenM, 2014; Littel, 2014)

Grouping the turbines in wind plants helps to reduce land- or sea-area use and landscape impact, and reduces the costs of installation and maintenance, and of connecting the turbines to the grid through cabling. A downside of placing wind turbines in larger plants, is that the aerodynamic interaction between the turbines may have a negative effect on the total electrical power production of the wind turbines, and may increase the loads experienced by turbines. The aerodynamic interaction effects are caused by the turbine wakes, which are the flow structures that form behind each turbine (cf. Figure 1.1 for an illustration). The negative effects of this aerodynamic interaction can be mitigated by placing the turbines further away from each other in the more prevailing wind direction, resulting in a lay-out optimization step in the design of the wind plant (Mosetti et al., 1994; Samorani, 2013), and/or by using wind plant control techniques during the operation of the wind plant (this thesis).



Figure 1.1: Clouds forming in the Horns Rev offshore wind plant in Denmark, showing the flow structures behind each turbine (the wind turbine wakes), that cause interaction with the turbines downstream. Source: Christian Steiness. See also Hasager et al. (2013) for more information on the picture.

The concept of wind plant control, that was first proposed in Steinbuch et al. (1988), aims at improving the performance of the wind plant as a whole, through coordinating the control operations across the wind turbines. Wind plant control has the possibility to reduce wind energy costs both by maximizing the power output of the plant as a whole (instead of each individual turbine, which can be suboptimal), as well as by helping to reduce the loads experienced by turbines. Also, through combined optimization of wind plant control and wind plant lay-out, a further increase of the power density

(generated electrical power per land area) of new wind plants can be achieved (Fleming et al., 2014c).

As part of the Dutch program for wind energy, researchers in industry and academia participating in the Far and Large Offshore Wind (FLOW) research program aim to realize a cost of energy reduction in 2015 of 20% for offshore wind energy compared to 2010, through various innovations (FLOW, 2010a). In the area of controls engineering, this research has aimed at innovations in the field of wind plant controls, and consist of the work in FLOW (2010b), and this thesis.

In this thesis, we specifically focus on improving the wind plant performance by taking into account the aerodynamic interaction between the turbines in the wind turbine control algorithm. When considering the optimization of the plant performance, we specifically aim at maximizing the power output and/or mitigating the loads on the wind turbines, rather than regulating their output towards certain set-point in order to maintain the balance between generation and demand on the electricity grid (cf. Aho et al. (2012)), although this may be a possible extension of some of the work.

1.2. BACKGROUND AND STATE-OF-THE-ART

In the previous section, we argued that wind plant control that takes into account wake interaction effects, can enhance the performance of the wind plant. In this section, we will provide background information on the characteristics of the wakes (Section 1.2.1), and on the state-of-the-art of individual wind turbine control (Section 1.2.2), before discussing the current work in wind plant control research (Section 1.2.3). Because most of the state-of-the-art wind plant control methods are model-based, also background information on wind plant models is provided in Section 1.2.4.

1.2.1. WIND TURBINE WAKES

The wind turbine wake is the flow structure downstream of a wind turbine, that is characterized by:

- a reduced flow velocity caused by the extraction of energy from the flow by the turbine,
- an expansion of the wake cross-sectional area: as the flow decelerates under the influence of the blade forces, the streamtube of the wake will expand as an effect of the principle of conservation of mass,
- an increased turbulence intensity caused by the obstruction of the flow by the turbine, and the resulting velocity gradients in the flow (shear),

The properties of the wake have been studied extensively, see Crespo et al. (1999); Vermeer et al. (2003); Sanderse et al. (2011) for literature overviews. The above characteristics are of interest for the control of wind turbines in wind plants, since the velocity deficits will cause a decrease of power production of turbines standing in the path of a wake of another turbine, and the increased turbulence and shear in the wake may increase the loads on those downstream turbines.

The amount of wake interaction between turbines is not only dependent on the overlap area of the wake of the upstream turbine with the rotor of the downstream turbine, but also on the distance between the turbines. This is because as the flow in the wake moves downstream, it will recover to the surrounding flow conditions (turbulence intensity, speed and main direction) through convection and diffusion. The turbulence in the flow promotes this process of wake recovery by mixing of the flow in the wake and the surrounding stream. Apart from the above-mentioned turbulence and shear caused by turbine, also the roughness of the Earth's surface, and thermal effects in the atmosphere cause turbulence and (vertical) shear in the inflow and downstream in the wake and the surrounding flow. Therefore, the ambient atmospheric conditions, consisting of the inflow turbulence intensity and the atmospheric thermal stability conditions, influence the amount of wake recovery (Abkar et al., 2014; Barthelmie et al., 2010), and thereby also the amount of interaction between turbines in a wind plant.

In descriptions in literature, wakes are often subdivided in a near wake region (close to the rotor) and a far wake region (further downstream). In the near wake, the turbine geometry directly affects the flow. In this region, vortices (rotating flow structures) caused by the tips and roots of the blades are present, that are characterized by a high turbulence intensity and large velocity gradients. In the far wake, the flow velocity profile has evened out more, and the effect of the rotor is only seen through more large-scale effects of velocity deficit, and increased turbulence intensity.

Another effect observed in wakes structures, called wake meandering, consists of large oscillating movements of the velocity deficit area in the lateral and vertical direction (see Aubrun et al. (2012), for example). Suggested causes for wake meandering are the large-scale turbulent structures in the atmosphere perturbing the flow direction of the wake (España et al., 2011; Larsen et al., 2007, 2008), and the regular formation and shedding (releasing) of vortices at the turbine rotor blades (Medici et al., 2006).

The concept of wind plant controls relies on the fact that the amount of wake interaction, i.e., the effect that one turbine has on another turbine through its wake, is dependent on the control settings of the turbines, since they influence the forces that the wind turbine exerts on the wind flow. In the next section, it is explained what the control degrees-of-freedom of a modern turbine are, and how they are used to control the power conversion of the turbine, as well as the loads on the system.

1.2.2. INDIVIDUAL WIND TURBINE CONTROL

In Figure 1.2 the main components of a horizontal-axis wind turbine² are shown, as well as a subsystems-level model scheme for the power conversion in the wind turbine, and the control of this process. The wind turbine consists of a rotor, most often with three rotor blades, that is attached to a generator through a drive-train. The generator and drivetrain are housed in a nacelle, which is supported by a tower. The nacelle can be rotated around the tower axis by a yaw mechanism. The rotor blades convert the momentum from the wind field passing the rotor-swept plane, into forces driving the rotor.

²Although there exist numerous architectures for wind turbines (see Hau (2013b) for an overview), this thesis focuses on horizontal-axis wind turbines, currently the most-used architecture in the wind industry.

The drive-train transfers the aerodynamic torque on the rotor, T_r , to the generator shaft, either directly (which is referred to as direct-drive), or through a transmission (a gear-box). The generator converts the rotational kinetic power into electrical power P , and generates a reactive torque on the shaft, T_g .

The efficiency of the power conversion from wind field kinetic power to wind turbine electrical power production in steady-state, is referred to as the power coefficient C_P . With this definition, the steady-state electrical power production of the wind turbine is (Molenaar, 2003):

$$P = \frac{1}{2} \rho A U^3 C_P \quad (1.1)$$

where ρ is the air density, and A the rotor-swept area. Scalar U is the effective speed of the wind field passing through the rotor, being the averaged wind speed that would be measured if the rotor was not present. The state-of-the-art power and rotor speed control mechanism of a horizontal wind turbine is the variable-speed, variable-pitch control concept. It makes use of the fact that the effectiveness of the rotor to generate torque on the generator, and thus the C_P -factor, is dependent on the inclination angle of the wind on the rotor blades, that can be adjusted by changing the pitch angle of the blades (rotating the blade around its longitudinal axis) and/or the tip-speed-ratio of the rotor, defined as:

$$\lambda = \frac{\omega R}{U} \quad (1.2)$$

where ω is the rotational speed and R the radius of the rotor. In Figure 1.2 an example C_P -curve is shown, showing the dependence of this factor on tip-speed ratio λ , and blade pitch β . Also, an expression for the aerodynamic torque T_r on the rotor is given, in which η_g is the efficiency of the conversion of the rotor rotation kinetic power to generated electrical power. Together, these relations are used to characterize the steady-state rotor aerodynamics. To regulate the amount of power conversion and the rotor speed, a controller sets a desired generator torque T_d , tracked by an internal controller of the generator, and a desired pitch angle of the blades β_d , a signal that is tracked by a pitch servo motor on each blade.

A simplified objective power curve for the controller is shown in Figure 1.3. For wind speeds below a certain threshold $U_{\text{cut-in}}$, the generation of electrical energy is not worthwhile, this range of wind speeds is referred to as Region 1. In the range between the cut-in wind speed $U_{\text{cut-in}}$ and the rated wind speed U_{rated} (Region 2), the controller aims to have the power conversion at maximum efficiency, by choosing the pitch at the setting β_{opt} , and regulating the tip-speed-ratio towards λ_{opt} , an operating point that corresponds to the maximum C_P value, $C_{P,\text{opt}}$ (see the C_P -surface in Figure 1.2). The optimal tip-speed-ratio λ_{opt} is reached in steady-state if the generator torque is set as:

$$T_{g,\text{opt}} = K \omega^2 \text{ with } K = \frac{\rho A R^3 C_{P,\text{opt}} r}{2 \lambda_{\text{opt}}^3} \quad (1.3)$$

where $r = \omega/\omega_g$ is the gearbox ratio. The above control law follows from substituting the optimal operating point in the relation for T_r in Figure 1.2, and removing the explicit dependency on the effective wind speed (that is not directly measurable) using the

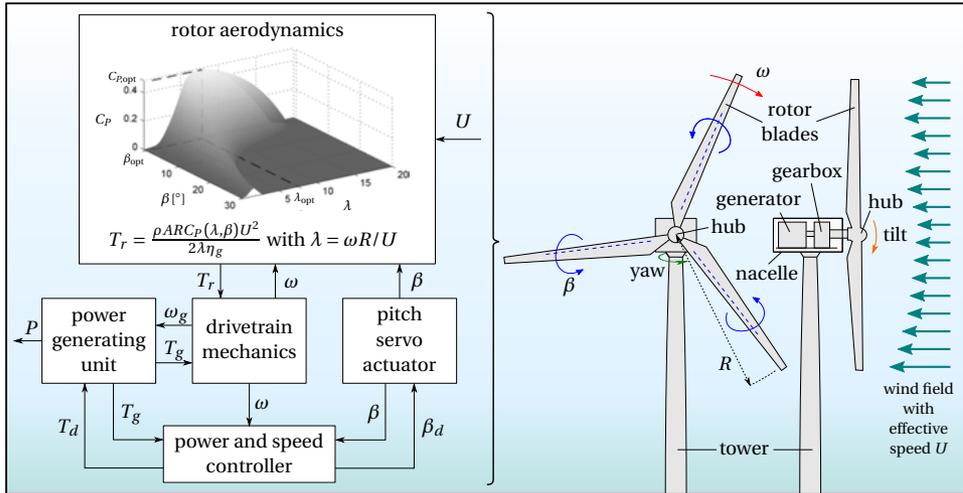


Figure 1.2: Subsystems-level model scheme for power conversion in a three-bladed horizontal-axis variable-speed, variable-pitch wind turbine. The forces exerted by the wind on the rotor blades generate a torque on the rotor T_r . This torque T_r is dependent on the effective wind speed U , the rotor speed ω and the pitch angle of the blades β . The rotor torque is transferred by the drive-train to a generator producing electrical energy. The internal controller of the pitch actuator tracks a desired pitch angle β_d . The power generating unit has an internal controller tracking a desired generator speed ω_d . These control inputs are used by a power and speed controller to regulate the electrical power produced by the wind turbine and the rotor speed.

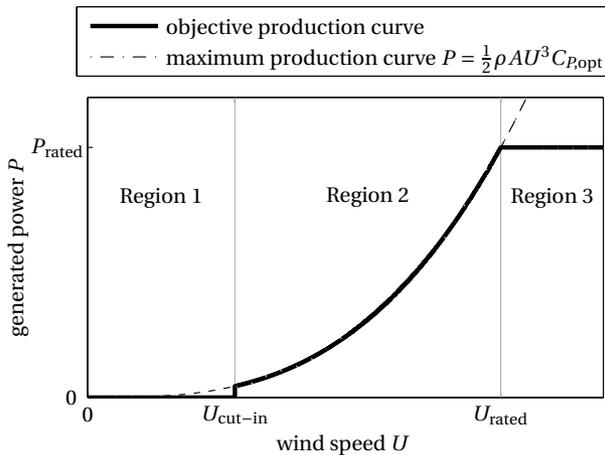


Figure 1.3: A possible objective power curve for the controller of an individual wind turbine. In Region 2, the turbine controller tracks the maximum power generation, while for higher wind speeds it tracks a rated power P_{rated} . Note that while the objective power generation varies with the wind speed, the control commands itself are mostly based on rotor speed measurements, rather than the (less reliable) wind speed measurements.

tip-speed ratio. Above the rated wind speed U_{rated} , in Region 3, the objective of the controller is to keep the electric power production, as well as the rotor speed and generator torque, below a certain maximum allowed (rated) value, in order to avoid overload of the mechanical structure and the power electronics. This is done by keeping the generator torque constant to the rated value, and adjusting the pitch angle of the blades to track the rated rotor speed. Thereby the pitch is increased with the wind speed to reduce the rotor efficiency. An often-used strategy to track the maximum rotor speed, is by using a Proportional-Integral (PI) control law for the pitch angle. In this strategy, each of the blades have the same pitch angle β , therefore this control loop is referred to as collective pitch control.

In individual wind turbine control, the yaw mechanism is mainly used to keep the rotor faced into the wind, which maximizes the energy extraction. In order to reduce actuator usage, the yaw reference is only updated when a certain threshold offset of the yaw with the wind direction is measured. The minimum time between yaw reference updates on a wind turbine may be in the order of tens of seconds to minutes (examples are discussed in Hau (2013a) and Kragh et al. (2013b)). The tilt angle is not actuated, a small constant offset of the rotor axis with the horizontal plane is often used to prevent the blades hitting the tower in high winds.

The power and speed control strategies described above are somewhat simplified. Between the different control regions shown in Figure 1.3, additional transition regions are defined in order to improve start-up behavior or to avoid high rotor speeds (see Jonkman et al. (2009) for an example). Further, the power and speed control scheme described above may be extended with additional control loops. Most often, load control is the goal of these extensions. Controls reducing fatigue and extreme loads will reduce damage to components and lengthen the turbine lifetime, and will possibly allow the use of lighter designs (e.g. more slender blades and lighter support structures), enabling a reduction of the manufacturing and operational costs of wind turbines (Lantz et al., 2012). The rotor blades experience a spatially- and time-varying wind field, causing vibrational loads on the blades that also propagate to the drive-train and the support system. Because the blades rotate through the field, the experienced loads have large periodic components. Control measures to reduce these vibrational loads include:

- Adjustments of the control laws to avoid certain rotor speeds at which periodic loads due to rotational sampling of the wind field would excite resonances of the system (Bossanyi, 2003b).
- Using the generator torque or the collective pitch for feedback control to dampen dynamic loads on the tower and drive-train (Bossanyi, 2003b).
- Using individual pitch control (IPC) techniques, where the forces on each rotor blade are individually controlled using different blade pitch offsets for each blade. By feeding back measured loads (e.g. strain of the blade roots), and transforming them to a non-rotating reference frame, the periodic loads on the blades and support structure can be reduced, at the cost of more high-frequent use of the pitch actuator (Bossanyi, 2003a).
- Using local actuation on the blades (e.g. using flaps) to locally affect the forces of

the wind field on the system (van Wingerden et al., 2008).

The above technologies are ordered, in an arguable way, by their maturity in development. While the first ones are more common in industrial practice, the latter are in research and development stages. Apart from the trend towards increasing the number of control loops in the turbine for load reduction, also there is a trend towards using *system identification* (using measured input-output data to derive models to be used controls development) and advanced *data-driven techniques* (in a more general sense, using measured data to perform control actions). These techniques can be used to adjust control laws to time-varying dynamics of the wind turbine, or to possible errors between the real wind turbine and theoretical models. Examples are found in Houtzager (2011); van der Veen (2013).

1

1.2.3. WIND PLANT CONTROL

In the previous section, we have described how individual wind turbines are controlled. When applying wind plant control for optimization of plant-wide electrical power production and wind turbine load reduction, some of the turbine controllers may need to deviate from the locally optimal control settings for the wind turbine. For example, an upstream turbine may need to reduce its power production in order to reduce the wake effects on downstream turbines, and increase the total power production of the wind plant. In this wind plant control approach, the control settings of the individual turbines are optimized with a global objective (e.g., total power production of the wind plant). Hence, in the context of this thesis, wind plant control consist of cooperative control of the turbines: the turbines exchange information which each other, or with a supervisory controller, and act on this information in order to reach a global objective (Bai et al., 2011).

Wind plant control strategies have different ways of using the controls degrees-of-freedom (DOFs) of the turbine. In Section 1.2.2 the conventional control DOFs were identified as the pitch angles of the blades, the generator torque, and the yaw angle. Since each of these control DOFs affect the power extraction of the wind turbine, they also affect the velocity deficit in the wake. When we only use the control DOFs to affect the wake velocity deficit, we refer to this as *axial-induction-based control*. Most wind plant control approaches in literature use this strategy, using generator torque and blade pitch as control DOF. Further, the direction of the wake can be changed using control, such that the overlap of the wake with downstream rotors can be avoided or reduced. This type of control is referred to as *wake redirection control*. Wake redirection can be achieved using the yaw DOF (Dahlberg et al., 2003) and by using individual pitch control (this thesis, Section 2.4.2).

For both wake-redirection and axial-induction based wind plant control, most current research studies on wind plant control methods take a model-based approach. In these studies the optimal control settings in wind plants are found, using models of wake effects in wind plants that range from simplified engineering models (e.g. Horvat et al. (2012), Heer et al. (2014) for axial-induction-based control, and Park et al. (2013) for wake redirection control) to mid- and high-fidelity simulations tools (Soleimanzadeh et al. (2013), Schepers et al. (2007), Annoni et al. (2014b), Goit et al. (2014) for axial-induction-

based control).

In the model-based control approach, a computationally-efficient engineering model of the wake effects in a wind plant is useful for quickly finding optimized control settings using iterative algorithms, while a more high-fidelity wind plant model, that tends to be more computationally complex, may be used for validation (before final tests on a real wind plant). In the next Section 1.2.4, we will provide an overview of wind plant models. A challenge for model-based strategies is that the optimal settings are dependent on the wake recovery properties, and that those properties are affected by the atmospheric conditions that are varying with time (cf. Section 1.2.1). Therefore, if model-based optimization is used to find the optimal settings for a wind plant at a particular time instant, the model parameters should match the specific atmospheric conditions at that time.

With this in mind, a model-free data-driven optimization method has been proposed in Marden et al. (2013) that tests the control settings, evaluates the effect on the total wind plant performance, and reiterates. Similar data-driven control methods, based on extremum-seeking control techniques, are presented in Johnson et al. (2012). A challenge with data-driven approaches, is that within the wind plant, there are significant delays between a control settings change on one turbine, and the effect on downstream turbines, since the wake effects have to propagate through the flow field. Therefore, the time-efficiency of data-driven methods depends strongly on the efficiency of the optimization.

1.2.4. WIND PLANT MODELING

In the context of this thesis, a wind plant model describes the interaction of a wind turbine with the atmospheric boundary layer. This interaction with the atmosphere also includes the interaction that the turbines may have with each other, through the wakes that form in the atmosphere behind each turbine. In recent years, a large number of wake models have been presented in literature. These models have different levels of complexity and fidelity, and may have different purposes, e.g., wind plant performance evaluation, lay-out optimization or controls development. A short recent overview of wind plant models can be found in Moriarty et al. (2014). In Table 1.1 some examples are given to illustrate the variety of models.

A wind plant model for use in wind plant control includes the following elements (as also illustrated in Figure 1.4):

- **A model describing the flow characteristics in the wind field around the turbines.** These flow field models range from simplified heuristic or analytic models consisting of explicit expressions describing the properties of the flow field (e.g., the velocity, turbulence intensity, direction) as a function of a limited number of input parameters (we refer to them as parametric models), to computational fluid dynamics (CFD) models that rely on solving the Navier-Stokes equations of motion in the flow field to predict its physics, cf. Table 1.1. Within these categories, several levels of fidelity for the prediction of different properties of the flow field can be distinguished, e.g., static (steady) versus dynamic (unsteady) models, two-dimensional (2D) versus three-dimensional (3D) models, etc.

		model name	flow field submodel	rotor submodel	turbine submodel
level of fidelity	parametric models	Jensen model (Jensen, 1984)	Ⓟ static parametric description of steady-state wake velocity profile	actuator disk	• C_p, C_T relations for power and thrust (eq. 1.1, 1.4)
		Frandsen model (Frandsen et al., 2006)	Ⓟ static parametric description of steady-state wake velocity profile	actuator disk	• C_p, C_T relations for power and thrust
		quasi-steady wind farm flow model (Brand et al., 2010)	Ⓟ static parametric description of steady-state wake velocity profile Ⓣ parametric description of standard deviation accounting for turbulence	actuator disk	• C_p, C_T relations for power and thrust • relations for static bending loads
	CFD models	wind farm state-space model (Soleimanzadeh et al., 2014, 2012)	Ⓟ linearized 2D unsteady NS Ⓣ no additional turbulence model	actuator disk	• C_p, C_T relations for power and thrust • relations for static bending loads • relations for effect of yaw on wake velocity (deflection effects not included)
		dynamic wake meandering model (Larsen et al., 2008; Hao et al., 2014)	Ⓟ 2D (radial) simplified steady NS (Ainslie, 1988) Ⓣ mixing-length turbulence model Ⓢ corrections for wake meandering and vertical shear	actuator disk	• FAST aero-elastics dynamics model (Jonkman et al., 2005)
		FarmFLOW (Schepers et al., 2007)	Ⓟ 3D simplified (parabolized) steady NS Ⓣ $k - \epsilon$ turbulence model Ⓢ corrections in near wake	actuator disk	• C_p, C_T relations for power and thrust
		ActiveWindFarms LES model (Goit et al., 2014)	Ⓟ 3D unsteady NS Ⓣ LES	actuator disk	• C_p, C_T relations for power and thrust
		SOWFA (Churchfield et al., 2012b)	Ⓟ 3D unsteady NS with thermal buoyancy effects Ⓣ LES	actuator line	• FAST aero-elastics dynamics model (Jonkman et al., 2005)

- Ⓟ main model for wake velocity profile
 Ⓣ additional model for small-scale turbulence effects
 Ⓢ additional heuristic corrections

Table 1.1: Overview of a selection of wind plant models, ordered, in an arguable way, by their level of fidelity (level of realism). Flow field submodels that describe the velocity profile of the wake (Ⓟ), vary from simplified static parametric descriptions, to CFD simulations based on Navier-Stokes (NS), with several levels of fidelity. The parametric models are all static relations, and in the CFD models there are some models that describe the steady-state flow field, based on (quasi-)steady NS, while there are other models that describe the dynamics of the wind flow based on unsteady NS. The methods used for including effects of (small-scale, unresolved) turbulence (Ⓣ) are varying from the more simplified models ($k - \epsilon$, mixing length) to more advanced Large Eddy Simulation (LES) methods. In some of the CFD models, additional heuristic corrections (Ⓢ) are applied to compensate for simplifications of the Navier-Stokes flow model. Also in the turbine model, there are several levels of fidelity of the models; from static relations based on the C_p, C_T relations, with several augmentations to predict load effects, to an aero-elastics dynamics model of the turbine.

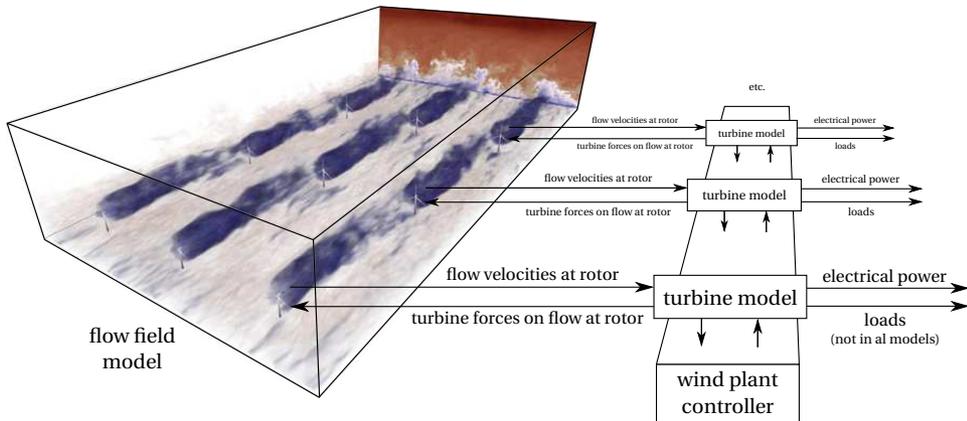


Figure 1.4: A generalized model scheme for wind plant models. A wind plant model includes a flow field model that predicts the wind velocity field in the wind plant, and a wind turbine model (one for each turbine in the flow). Based on the predicted velocity field at the rotor, a turbine model predicts the forces of the rotor on the flow, that are fed back to the flow field model, as well as the electrical power production and the structural loads on the wind turbine. Each turbine communicates with a wind plant controller. In this case, the flow field model is visualized by the results of the 3D large-eddy simulations in SOWFA (picture from Lee et al. (2012)), but in many other wake interaction models more simplified flow field models are used.

- A model of the forces of the wind turbine on the flow and the reactive forces of the flow on the wind turbine.** If the wake model has the purpose of being used for control development, the effect of the control degrees-of-freedom (pitch, rotor speed and possibly yaw or tilt) on the rotor forces are to be included in this part of the model. The exact representation of the blade bodies in the flow is generally too computationally costly for the simulation of wind plants, therefore modeling methods with simplified rotor representations are used. In the *actuator disk method*, the rotor is modeled as a disk of distributed forces exerted on the flow. As an example, in its most simple form, the thrust force F_T of the rotor on the flow is assumed to be uniformly distributed over the rotor disk, and modeled using a static relation:

$$F_T = \frac{1}{2} \rho A C_T U^2 \quad (1.4)$$

where, as before, A is the rotor area and ρ the air density, and C_T is the thrust coefficient that, similar to the C_p factor, is dependent on the control settings, and U is the effective wind speed estimated with the flow characteristics model. A more detailed representation is the *actuator line method*, introduced by Sørensen et al. (2002), where each blade is represented as a distribution of forces along a moving line, and at a number of points of the blade, these forces on the flow are calculated from the local flow velocities using the local lift and drag characteristics of the airfoil. An even more accurate, but also more computationally costly representation is one in which the blades are represented as planes of distributed forces (the *actuator surface method*). Refer to Sande et al. (2011) for a more comprehensive overview of rotor modeling methods.

- A model describing the effect of the flow forces on the performance of the wind**

turbine system, in terms of power production and, in some but not all models, loads on the system. This part of the model may range from a simplified static model describing the turbine performance as a function of the effective wind speed, to more comprehensive models describing the full structural dynamics. An example of a simplified steady-state power production model was given previously in Section 1.2.2 (eq. 1.1), and an example of a simplified structural loads model is the one presented in Soleimanzadeh et al. (2012), which relates the steady-state thrust (predicted by eq. 1.4) to the tower and blade bending loads. A comprehensive wind turbine structural dynamics model that has been included in wake interaction models is the FAST model (see below in SOWFA description).

- If the model is to be used for evaluating wind plant control models, a **wind plant controls structure** needs to be emulated. The wind plant controller communicates with each of the wind turbines, collecting local measurements at the turbine sensors and sending information or commands to each of the turbines.

In the context of this thesis, it is important to note that not all wind plant models are able to predict the effect of all the control DOFs on the flow field. More specifically (cf. Table 1.1), while the high-fidelity LES wind plant model SOWFA contains all the physics models to able to predict the effects of the yaw DOF and individual pitch control on the flow field in terms of the wake redirection, the simplifications in the wind plant state-space model (Soleimanzadeh et al., 2014) and the dynamic wake meandering model (Hao et al., 2014) make that these wake redirection effects are not captured in the model. Also the parametric models in Table 1.1 do not contain relations to predict the wake redirection effects of control³.

Below, we will discuss SOWFA in more detail, as it will be used in different parts of this thesis to evaluate wind plant controls concepts.

1.2.4.1. THE SOWFA WIND PLANT MODEL

One of the more comprehensive wind plant models available is the Simulator for On-shore/Offshore Wind Farm Applications (SOWFA). It is explained in more detail in Churchfield et al. (2012b). Here, a summary of the main features relevant for this thesis is given.

SOWFA's wind field model consists of a CFD simulation of the 3-dimensional wind flow around the turbine rotors in the atmospheric boundary layer, using a large-eddy simulation (LES) method (example results are shown in Figure 1.4 and 1.5). In the LES method, in order to limit the computational cost of the simulation, the larger scales of the flow field are resolved by solving the temporally and spatially discretized unsteady Navier-Stokes equations with additional stress terms that empirically model the effect of the turbulence in the smaller unresolved (subgrid) scales. In the Navier-Stokes model, Coriolis forces that account for the Earth's rotation are included in the momentum equations, as well as a buoyancy term that models the combination of gravitational and thermal effects on the flow. To generate the buoyancy term, also the advection-diffusion equations for temperature need to be resolved. The rotors are represented by rotating

³The parametric model used in (Park et al., 2013) for yaw-based wake redirection was left out of the overview, since it is not validated with high-fidelity simulations or measured wind plant data.

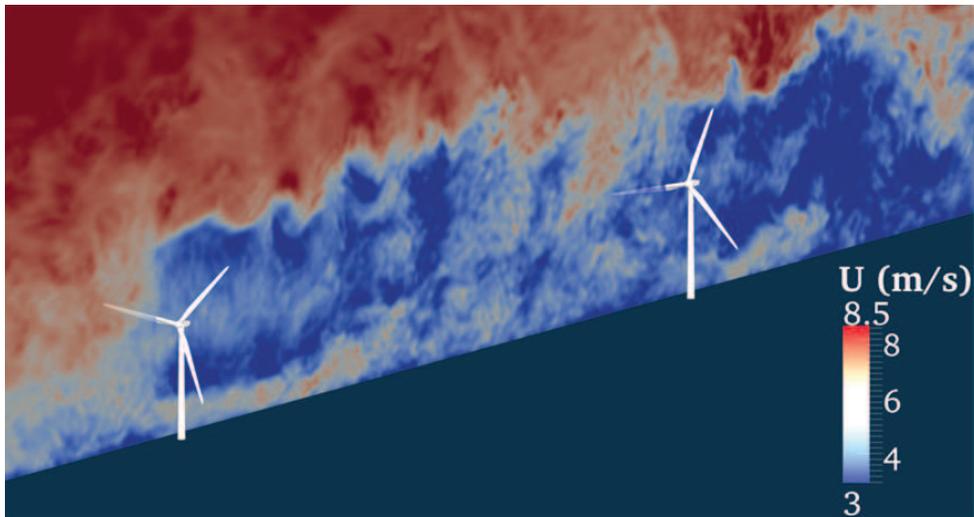


Figure 1.5: A vertical cut-through of the wind velocity field around two turbines as simulated by SOWFA. Note that while wind turbines are visualized in the flow, in the SOWFA simulation only the effect of the rotor forces are included in the flow simulation, through the actuator line method. Source: Churchfield et al. (2012b).

actuator lines in the flow field simulation (the nacelle and tower of the turbine are not represented). The local velocity vectors at points along the actuator line are sampled from the flow field. From these, the FAST wind turbine aero-elastic multibody dynamics simulation (Jonkman et al., 2005) then calculates at each simulation time-step:

- the blade forces on the flow distributed along the actuator lines, using the local lift and drag properties of the blade airfoil,
- the structural loading responses induced by the aerodynamic forces (e.g., blade, tower, and drive-train loads),
- the resulting rotor acceleration, speed, and position,
- the displacements of the actuator line caused by blade and tower bending, rotor rotation, and yaw,
- the electrical power production of the turbine.

Apart from the velocity field samples from the CFD simulation, inputs to the FAST aero-elastics simulation are control settings (reference signals for blade pitch, generator torque, and yaw angle) and the turbine design properties. A controller algorithm can be included in FAST to provide the control settings. Further, in Fleming et al. (2013a,b) a wind plant controls structure was introduced in the SOWFA simulation framework, that can send and receive measurements and commands from and to each wind turbine in simulation. This controller can be used to implement and test plant-wide control strategies in the SOWFA simulation environment.

Prerequisites for a wind plant simulation experiment in SOWFA are the initial conditions for the flow field, as well as the inflow to the simulated domain. These should rep-

1

resent the ambient atmospheric conditions that a researcher is interested in testing in simulation. They are generated by running a so-called precursor simulation. In this precursor simulation, there are no turbines present. The outflow properties are measured, and the boundary conditions on the inflow side of the domain are set the same, such that the flow is cycling through the domain. A pressure gradient over the domain is set to control the mean velocity and direction at hub-height to a certain desired value for the following simulation experiment. At the bottom of the domain, a ground surface roughness is set, that will cause turbulence to be generated, i.e., with this surface roughness, a desired turbulence intensity can be set. Further, a surface temperature flux is set, which also affects the turbulence intensity by imposing a certain atmospheric stability condition. The simulation is started with an initial velocity, pressure and temperature profile, and once the turbulence structures reach a quasi-steady state, they are sampled for the duration of the following simulation experiment, and stored. This stored flow field is then prescribed as the inflow for the simulation experiment with the turbine present in the flow.

If SOWFA is used for high-fidelity predictions of the wake interaction and the atmospheric effects on the wind turbine performance, a detailed representation of the flow field dynamics needs to be calculated. This includes resolving the blade-induced vortices in the near wake, as well as the turbulence properties in the far wake and the upstream induction zone (the part of the flow field upstream of the turbine that is affected by the rotor, cf. Simley et al. (2014)). To do this, a spatial discretization with relatively small dimensions (typically around 3 meters) is used around the turbines and in the wake. Also, the effects of larger turbulent structures in the atmosphere on the wake need to be included (in order to predict wake meandering, for example), resulting in a large simulated domain. Therefore, for these high-fidelity simulations, the computational complexity of SOWFA is rather large. A typical calculation will take tens of hours of distributed computation on clusters with a few hundred processors.

Validation of the SOWFA tool is an ongoing process. In Churchfield et al. (2012a), SOWFA was used to simulate the 48-turbine Lillgrund wind plant, and the results were then compared with field data, with good agreement throughout the first five rows.

1.3. THESIS OBJECTIVES

With the objective to contribute to the reduction of the cost of energy of offshore and onshore wind energy (cf. Section 1.1), this thesis aims to further develop control techniques that improve the performance of the wind plant, by taking into account wake interaction effects.

In Section 1.2.1 it was explained that the wake interaction is influenced by the atmospheric conditions. A wind plant control strategy that is to be applied on a real wind plant in changing atmospheric conditions, should therefore be able to adapt the control settings of the turbines to these time-varying conditions in real-time. In order to optimize the wind plant performance, it can make use of all of the control degrees-of-freedom of the turbine.

In existing work in wind plant controls research (cf. Section 1.2.3), we see develop-

ments of wind plant controls that mainly focus on axial-induction-based control (using generator torque or pitch). In most of this work, use is made of a model-based approach to perform the offline optimization of control settings. Generally, in the model-based approaches, and in the proposed model-free approaches, the time-efficiency of the optimization on the wind plant in real-time is not fully taken into account.

In this thesis, we contribute to the existing body of work, by:

Thesis Objective 1 Evaluating the potential of each of the currently-used control degrees-of-freedom of the wind turbine (generator torque, collective and individual blade pitch and rotor yaw) to affect the wake interaction effects between the turbines.

Thesis Objective 2 Developing methods to optimize the different control settings of the wind turbines in order to improve overall wind plant performance (in terms of power production and/or loads on the turbines), taking into account the time-efficiency of the optimization in real-time implementation on the wind plant.

In the next Section 1.4, we will go into the general methodologies used in this thesis to fulfill the above objectives. Then, in Section 1.5 we will further specify the thesis contributions in the thesis outline.

1.4. METHODOLOGIES

In the proposed wind plant control methods in this thesis, a data-driven approach is taken. By basing the optimization of the control settings on measured data, the wind plant controller is able to adapt to the time-varying atmospheric conditions (e.g., wind velocity, turbulence intensity and atmospheric stability), which influence the wake properties and therefore the optimal point of operation of the turbines.

In this section, we introduce the basic approaches for data-driven wind plant control taken in this thesis: a *direct data-driven optimization approach*, and a *data-driven model-based approach*, both illustrated in Figure 1.6.

In the *direct data-driven optimization approach* (Figure 1.6a), measured data is directly used in the optimization algorithm to find the optimal operation point, based on a memory of previously tested settings. In this approach, the efficiency of the optimization algorithm (in terms of the number of iterations) is of importance, since the optimization algorithm aims to quickly adapt the control settings to time-varying atmospheric conditions. Notice that while there is no model included in the scheme in Figure 1.6a, here we do not use the term model-free to describe the approach, because to perform the optimization in a time-efficient manner, some inherent assumptions on the behavior of the wind plant are taken based on physical reasoning, which can be viewed as a qualitative model.

In the *data-driven model-based approach* (Figure 1.6b), from the measured data, the parameters of a wind plant model are estimated in an identification procedure, and the state of the model (representing, for example the properties of the wind field) is updated using an observer. Then the optimization algorithm can iteratively test control settings on this model to find the model-predicted optimal operation point, before applying the

optimized control settings on the real wind plant. For the wind plant controls development using this model-based control strategy, not only the accuracy of the predictions, but also the computational complexity is a relevant property of the wind plant model. In case the optimization is taking place on-line in time-varying atmospheric conditions, it is required that the control updates take place in a relatively short time, therefore the computational complexity of evaluating the predictions of model should be low enough. Therefore, a high-complexity model like SOWFA is not suited for online optimization, and in this thesis more computationally efficient *control-oriented* parametric models are developed for which the parameters can be found from measured data. Control-oriented, in this context, means that the effects of the wind turbine control DOFs on the wake interaction that are relevant for the wind plant performance, can be predicted by the model.

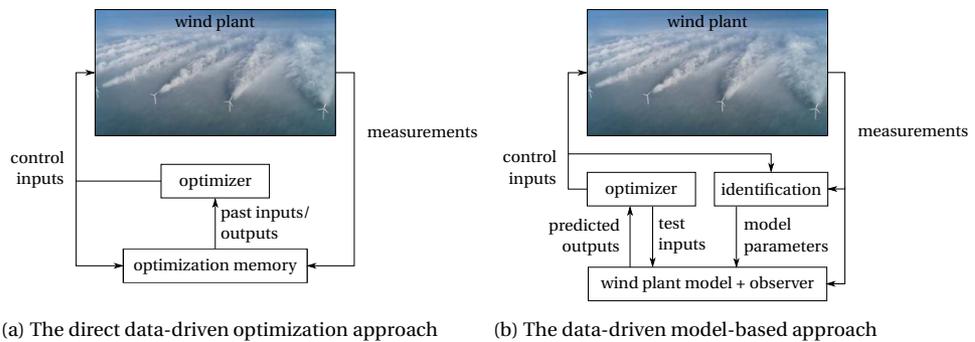


Figure 1.6: Block schemes illustrating the two basic approaches for data-driven wind plant control in this thesis

In this thesis, we provide proof-of-concepts of control approaches, both in terms of the effect of control DOFs on interaction effects (Thesis Objective 1), as well as the effectiveness of the associated algorithms for optimization of the control settings (Thesis Objective 2). At several points we use SOWFA high-fidelity simulation to provide these proof of concepts, since so far, we have not had the ability to test control concepts on a real wind plant. The SOWFA simulator includes relevant dynamics and is one of the more comprehensive CFD models available, but like any model it contains simplifications (e.g. in the discretization of the spatial domain). By not basing the control algorithm on the specific model formulation in SOWFA, but taking a more generic approach where data is extracted from SOWFA for use in data-driven control concepts, using the measurements that are also available on current wind plants (e.g., local wind direction at the turbines, and turbine power productions), we aim to develop procedures and control algorithms that can also be applied on real wind plants using readily available measurement data.

1.5. THESIS CONTRIBUTIONS AND OUTLINE

In this thesis, different aspects of the proposed data-driven control schemes as explained in the previous Section 1.4, are developed, in order to meet the objectives formulated in Section 1.3. We have presented the work in different chapters, that each stand on their own in the sense that they can be read independently, so some of the introductions and

definitions are repeated in the different chapters.

- In Chapter 2, the potential of the different control degrees-of-freedom of the wind turbine to affect the wake interaction effects in the wind plant in SOWFA high-fidelity simulations are evaluated, and the results are discussed.

Parts of this chapter have been published in Fleming, Gebraad, Lee, van Wingerden, Johnson, Churchfield, Michalakes, Spalart, and Moriarty (2014b,d); Annoni, Gebraad, Scholbrock, Fleming, and van Wingerden (2014a).

- In Chapter 3, a time-efficient *direct data-driven optimization approach* for *axial-induction-based control* is developed.

This chapter has been published in Gebraad and van Wingerden (2014b).

- In Chapter 4, a time-efficient *data-driven model-based approach* for *yaw-based wind plant control* is developed, in which the wake is redirected using yaw offsets. A control-oriented, data-driven, parametric model is developed for this approach, predicting the wake-redirection effects of yaw.

This chapter will appear as a journal publication in Gebraad, Teeuwisse, van Wingerden, Fleming, Ruben, Marden, and Pao (2014d).

- In Chapter 5, a *data-driven control-oriented model* is developed, that takes into account specific dynamics of the wake interaction. A Kalman filter for updating the flow predictions of the model using measured data is developed, and the model is used in a *model-based wind plant control* example.

This chapter has been published in Gebraad and van Wingerden (2014a).

Finally, in Chapter 6, the conclusions of this thesis are presented.

2

EVALUATION OF CONTROL DEGREES-OF-FREEDOM FOR WIND PLANT CONTROL

In this chapter the different wind turbine control degrees-of-freedom (DOFs) that can be used to affect the wake interaction effects in a wind plant, are evaluated. For each of the concepts, we explain the mechanism by which the wakes and the wake interaction effects are affected by a certain wind turbine DOF, present case studies using high-fidelity SOWFA simulations and discuss their results in terms of the effect of a DOF on turbine-to-turbine wake interaction. First we study the axial-induction-based concept, that consist of using pitch and/or generator torque control to adjust the power production on upstream turbines to increase the total electrical energy yield of the wind plant. From the simulation results combined with earlier results from literature, we can conclude that the effectiveness of the axial-induction-based control is dependent on the specific inflow and atmospheric conditions, and the turbine characteristics. The results even suggest that there are circumstances in which the concept of total wind plant power increase through axial-induction-based control is infeasible. Another concept in which the yaw DOF is used to induce wake redirection, is shown to be effective at reducing the wake interaction effects and increasing the power production of a simulated two-turbine setup. Other less conventional strategies that affect the wake interaction in the wind plant, being individual pitch control based wake redirection, rotor tilt wake redirection, and repositioning of floating turbines, are also tested in SOWFA simulations.

Parts of this chapter have been published in Fleming, Gebraad, Lee, van Wingerden, Johnson, Churchfield, Michalakes, Spalart, and Moriarty (2014b,d); Annoni, Gebraad, Scholbrock, Fleming, and van Wingerden (2014a).

2.1. INTRODUCTION

In this chapter we evaluate the different control degrees-of-freedom (DOFs) that can be used to affect the wake interaction effects in a wind plant. We study:

- conventional control DOFs on modern wind turbines (collective blade pitch and generator torque) that affect the flow mainly by influencing the axial induction of the rotor and thereby the velocity deficit in the wake,
- turbine DOFs that change the direction of the wake, being rotor yaw, individual blade pitch control, and the tilt of the rotor (an unconventional DOF),
- repositioning techniques to move floating turbines out of the wake.

For each of the concepts, we explain the mechanism by which the wakes and the wake interaction effects are affected by a certain control DOF, show simulation case studies in which the settings for the DOF are varied, and discuss the results of these simulations. Each of the concepts is studied in SOWFA simulations of setups with either one or two turbines. These simulation scenarios are explained in more detail in Section 2.2. Then, axial-induction-based techniques are discussed in Section 2.3, and wake redirection and turbine repositioning techniques are discussed in Section 2.4. Summarizing conclusions are provided in Section 2.5.

2

2.2. SIMULATION SCENARIOS

In the next two Sections 2.3 and 2.4, SOWFA simulations are used to study how the aforementioned control DOFs affect the wakes and the wake interaction effects. Two scenarios are simulated: a single turbine scenario and a scenario with two turbines aligned in the flow. The scenarios are explained in more detail in this section. An introductory explanation on the SOWFA simulator was given in Section 1.2.4.1.

2.2.1. SINGLE-TURBINE SIMULATION SCENARIO

In this scenario, we simulate an NREL 5-MW baseline turbine (described in Jonkman et al. (2009)) in turbulent inflow, in a domain that is 3 km by 3 km in the horizontal and 1 km in height. Details on the positioning of the turbine and meshing of the domain are given in Figure 2.1. It is shown that the mesh is refined in two steps in a rectangular region. The smallest mesh cells for the CFD calculation contain the turbine rotors, the axial induction zones of the rotor and a large part of the wake. Further away from the turbines the mesh is coarsened in order to reduce the computational cost of the simulation.

By studying the wake properties with different settings of the control DOF, we can investigate how these DOFs affect the wake. In the baseline case against which we compare the different DOF settings, the turbine is operating at a nominal below-rated (Region 2) tip-speed ratio of 7.55, with a zero yaw offset of the rotor with the mean wind direction.

The conditions simulated, which are based on the study in Churchfield et al. (2012b), are that of a neutral atmospheric boundary layer (ABL), with a low aerodynamic surface roughness value of 0.001 m that is typical for offshore conditions. The inflow is generated

in a precursor simulation of the same domain, but without the turbines present, without the mesh refinements, and with periodic boundary conditions. In this precursor simulation the turbulent structures develop in the flow, and the horizontally averaged wind speed is driven to 8 m/s at the turbine hub height, controlled through a time-varying mean driving pressure gradient. Finally, after a total of 17,000s of simulated time, the turbulence intensity of the inflow develops to 6%, and the vertical change in mean wind velocity across the rotor disk to 1.46 m/s. The wind comes from the southwest (300°). In the final 1,000s of the precursor simulation, the full flow field is sampled and stored. This stored flow field is then prescribed as inflow for the simulations with the turbine present in the flow.

SOWFA requires significant computational power in order to run high-fidelity simulations: using a sample time of 0.02 s, the time steps take an average 2.5 s to calculate on the Sandia National Laboratories/NREL Red Mesa supercomputer (National Renewable Energy Laboratory, 2012), using distributed computation with 256 processors. In most of the simulations, we use a simulated time length of 1,000s in order to let the wakes develop through the domain and collect data in a fully developed flow, which thus yields an execution time of 34.4 h for each simulation.

Figure 2.3 shows the time-averaged flow field in contour planes taken from the simulation of the single turbine in the baseline case, with the rotor tilted up 5° to prevent blade strikes on the tower (the default tilt setting for the NREL 5-MW, Jonkman et al. (2009)). The vertical slices through the wake at various downstream locations, shown in the bottom two rows of Figure 2.3, show the mean wake as viewed from upstream looking downstream. Note that the wake moves to the right with increasing downstream distance, even though there is no yaw misalignment. This deflection can be explained by vertical shear in the boundary layer and wake rotation: in reaction to a rotor rotating clockwise, low speed flow in the lower part of the boundary layer will be rotated up and to the right, and high speed flow in the upper part of the boundary layer will be rotated down and to the left, and as a result the wake deflects to the right.

2.2.2. TWO TURBINE SIMULATION SCENARIO

To study the wake interaction effects between turbines, a second scenario was developed to simulate two NREL 5-MW baseline turbines in turbulent inflow, with a downwind spacing of 7 rotor diameters (7D), a typical turbine spacing. Details on positioning of the turbines are given in Figure 2.2. The same domain and meshing properties are used as in the single-turbine simulation scenario described in Section 2.2.1 and shown in Figure 2.1. Also the same ABL properties, boundary conditions and inflow properties are used as in the single-turbine simulation. The computational cost of each case is approximately the same as for the single-turbine cases described in Section 2.2.1. Conservatively assuming the wake convection speed to be one-half the mean hub-height wind speed, and setting a length scale equal to the turbine-to-turbine spacing of 7 rotor diameters, then the 1000 s simulation time can be expressed as at least 4.5 wake flow-through periods.

2

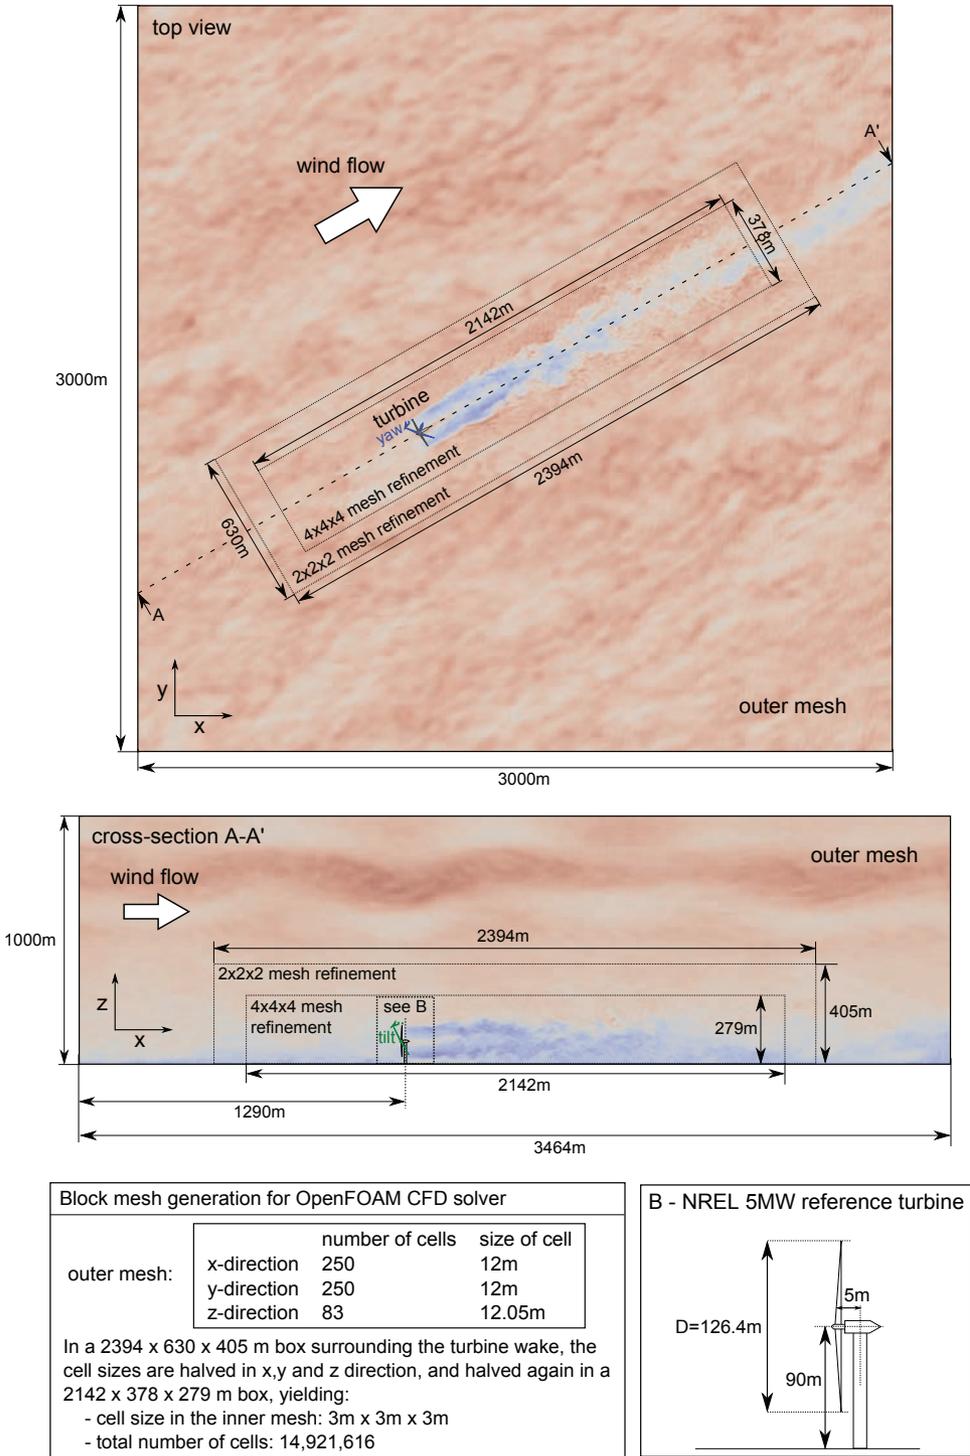


Figure 2.1: Overview of the simulation setup in the single-turbine baseline case.

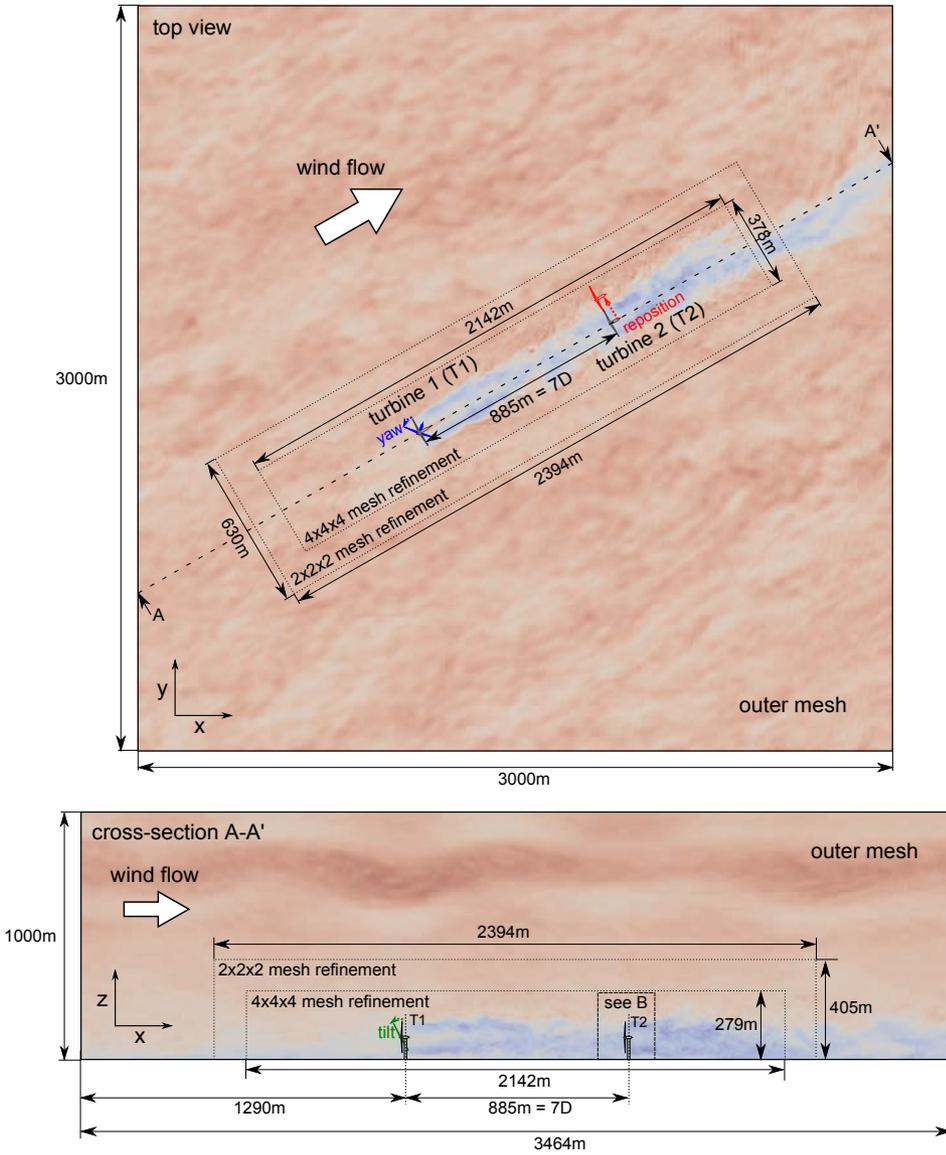


Figure 2.2: Overview of the simulation setup in the two-turbine baseline case.

2.3. EVALUATION OF AXIAL-INDUCTION-BASED METHODS

In this section we study the collective pitch and generator torque control DOFs of the turbine, that affect the flow mainly by influencing the axial induction of the rotor and thereby the velocity deficit in the wake. First the mechanism is explained in more detail in Section 2.3.1. Then in Section 2.3.2, a short literature overview is given of different studies investigating the potential benefit of the axial-induction in terms of potential power production increase. In Section 2.3.3, a simulation example is provided, showing some of the difficulties arising in the axial-induction control concept. A discussion of the results follows in Section 2.3.4.

2.3.1. THE MECHANISM OF AXIAL-INDUCTION-BASED WIND PLANT CONTROL METHODS

In this thesis, the wind plant control concepts are aimed at improving the overall wind plant performance by controlling the wake interaction effects in the wind plant. Most of literature on this type of wind plant control has focused on adjusting the power extraction of some of the turbines in the wind plant, in order to influence the velocity deficit in the wakes (Knudsen et al. (2014)). The power extraction can be adjusted by using conventional control DOFs on a turbine, such as blade pitch angle and generator torque. In this thesis, we refer to this type of control as axial-induction-based control methods, because the generator torque and blade pitch are adjusted to influence the axial induction factor of the rotor. The axial induction factor a is the fractional decrease in wind velocity between the free stream and the turbine rotor (see also Figure 2.4). The generator torque and blade pitch influence the axial induction of the rotor, and therefore also affect the velocity deficit in the wake the rotor generates. In the wake behind the rotor, the flow expands, and as the flow moves downstream, it recovers to the free-stream conditions because it extracts momentum from the surrounding flow through diffusion and convection. The rotor distorts the inflow, which, together with surface roughness thermal effects and velocity gradients in the wake, creates turbulence in the wake that acts as a ‘mixer’ of the free-stream and the wake, advancing the wake recovery (Sanderson et al., 2011). If another turbine downstream is standing in the path of the wake that is not yet fully recovered to the free-stream conditions, it experiences the reduced wind speed in the wake, which results in a lower electrical power production of the downstream turbine. Therefore in wind plants, in which turbines are placed relatively close to each other, the wake effect causes a coupling between the control settings of upstream turbines and the power productions of downstream turbines.

The amount of total power production gain that can be achieved from optimizing the control settings of the turbines, is dependent on the aerodynamic characteristics of the turbine, as well as on the atmospheric conditions. Two important characteristics of the turbine are the power coefficient C_P and the thrust coefficient C_T , both of which are a function of the tip-speed ratio (TSR) and the pitch of the blades (Bianchi et al., 2007). The TSR is given by:

$$\lambda = \frac{\omega R}{U} \quad (2.1)$$

where R is the rotor radius, ω the rotor speed, and U the free-stream wind speed (i.e.

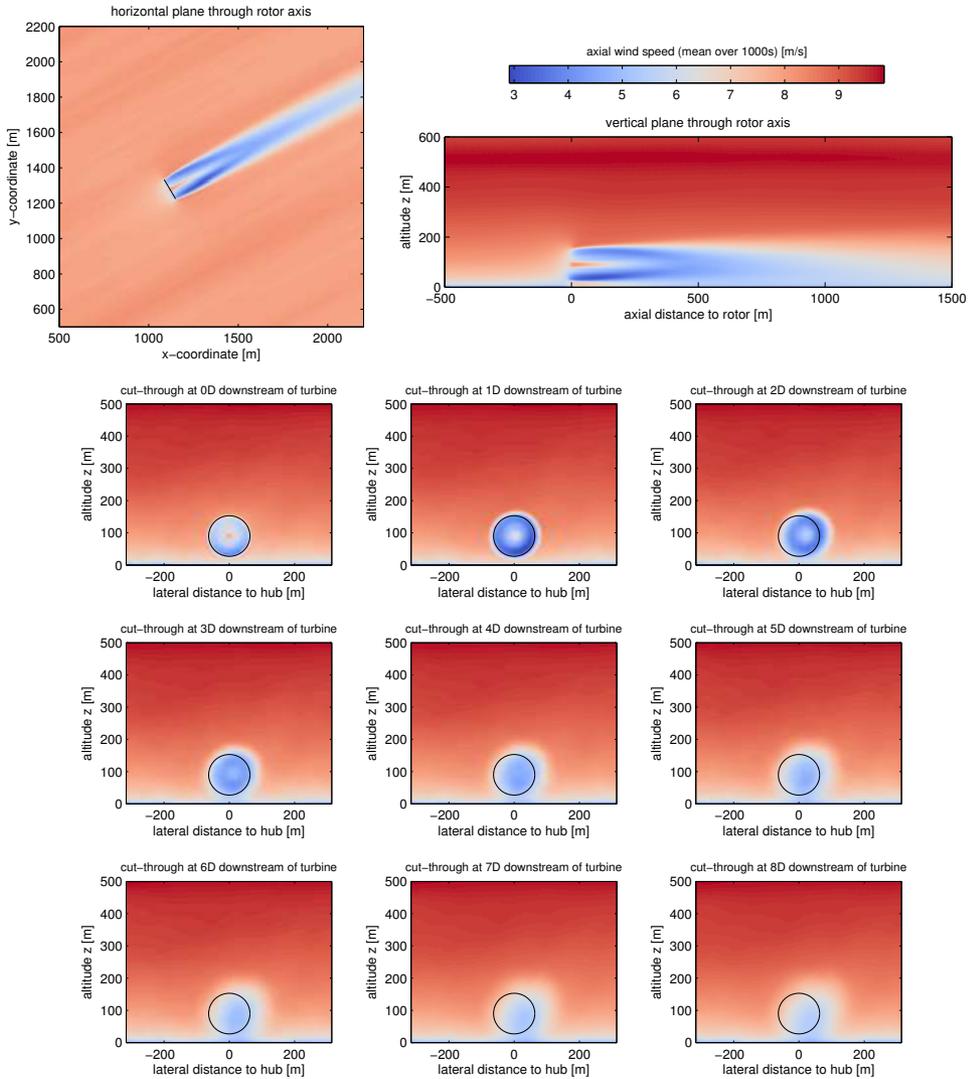


Figure 2.3: Averaged velocity profiles of the wake forming behind the NREL 5-MW turbine with no pitch or yaw control as calculated in the SOWFA simulation. D is rotor diameter. The x - y plane is a view from above and the downstream planes are as viewed from upwind looking downwind.

the wind speed not disturbed by the rotor, see Figure 2.4). The rotor speed, and thus the TSR, can be influenced by adjusting the generator torque or changing the lift forces on the rotor blades by adjusting the blade pitch. The power coefficient determines the efficiency of the rotor in power extraction; the steady-state power extraction of the rotor, P , is given by:

$$P = \frac{1}{2} \rho A C_P (\beta, \lambda) U^3 \quad (2.2)$$

with A the total area swept by the rotor, ρ the air density and C_P the power coefficient expressed as a function of the TSR λ and the collective blade pitch β . The thrust coefficient determines the thrust force of the rotor on the flow. The total thrust force of the rotor directed opposite of the flow, F_T , is given by:

$$F_T = \frac{1}{2} \rho A C_T (\beta, \lambda) U^2 \quad (2.3)$$

The thrust of the rotor determines the reduction of velocity over the rotor plane, i.e. the axial induction a (see also Figure 2.4). From actuator disc momentum theory it follows that if we assume that there is no recovery of the wake, the extraction of energy over the rotor makes that velocity in the wake behind the rotor drops to:

$$U_{\text{wake,min}} = U (1 - 2a) \quad (2.4)$$

with axial induction a being related to the thrust factor as follows:

$$a = \frac{1}{2} \left(1 + \sqrt{1 - C_T} \right) \quad (2.5)$$

In reality, there is wake recovery through convection and diffusion of momentum, therefore $U_{\text{wake,min}}$ can be considered to be a lower bound on the wind velocity in the wake.

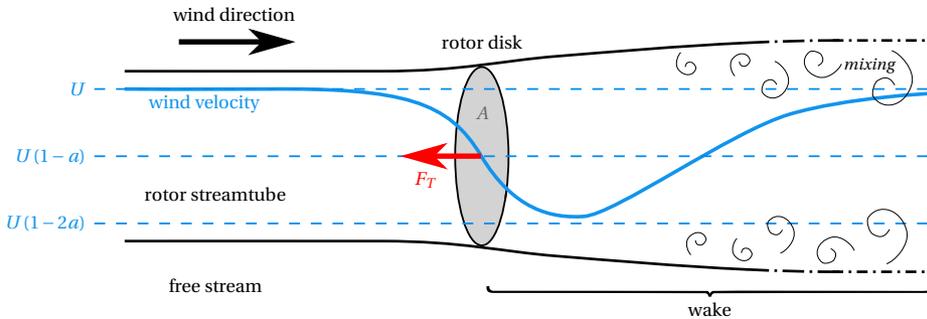


Figure 2.4: Simplified representation of the wake and upstream induction zone of the wind turbine rotor. The blue solid line is a possible time-averaged profile of the wind velocity, averaged over the cross-section of the streamtube of the flow that passes through the rotor. In the induction zone in front of the rotor, the flow velocity is already reduced, then there is a minimum of the velocity in the wake behind the rotor, after which the flow recovers to the free-stream conditions. U is the free-stream velocity, a the axial induction factor, F_T the rotor thrust force, and A the area of the rotor disk.

In below-rated wind conditions, the axial-induction wind plant control concept relies on the fact that at the operation point of maximum extraction of a single turbine,

the sensitivity of the power production to the pitch and generator torque is small, since the surface describing the dependence of C_P on TSR and pitch is flat around its optimum, while the thrust force is more sensitive to the pitch and TSR, since the C_T surface is not flat at the same pitch and TSR operation point (see the C_P and C_T surface in Figure 2.5 as an example). This means that by deviating a small amount from the point of maximum C_P by adjusting pitch and/or generator torque on an upstream turbine, the power production of that turbine will reduce only a small amount, while the axial induction will reduce enough to significantly increase the velocity in the wake (Schaak, 2006). Under the right circumstances, this increase in velocity downstream of the rotor will increase the power of a downstream turbines more than the loss in power production on the upstream turbine. Factors affecting the amount of production increase that can be achieved through axial-induction-based wind plant control techniques include:

- The specific rotor characteristics: the ratio of the gradients of the C_T -surface and the C_P -surface around optimum operation determine how much the velocity in the wake can be increased by a certain amount of reduction of the power production on the turbine creating the wake.
- The overlap of the wakes with the downstream turbine. If the overlap of the wake is small, the relative gain in power production that can be achieved is small. The wake overlap is affected by the wind direction, the expansion of the wake and the relative positioning of the turbines. Also the overlap changes over time under influence of vertical and horizontal movements of the wake (wake meandering).
- The amount of wake velocity recovery of the wake having taken place at the location of the downstream turbine through mixing. More mixing takes place if the turbulence intensity in the inflow is high. Also, the amount of turbulent mixing is dependent on the atmospheric thermal stability conditions. The effects of inflow turbulence intensity and atmospheric stability on wake recovery are discussed in Abkar et al. (2014) (a CFD simulation study) and Barthelmie et al. (2010) (a study of wind plant field data). If there is more turbulent mixing in the wake, the velocity in the wake recovers to the velocity of surrounding flow over a shorter distance, and the relative gain in power production that can be achieved through the axial-induction-based wind plant control is smaller.

2.3.2. OVERVIEW OF AXIAL-INDUCTION-BASED WIND PLANT CONTROL STUDIES

Several studies investigate the beneficial effect of axial-induction-based control for the optimization of wind plant total power production:

- Several simulation studies have shown a beneficial effect of axial-induction-based control on power production, using wake models that range from engineering models (e.g. Horvat et al. (2012), Johnson et al. (2012), Marden et al. (2013) and the study in Chapter 3) to more high-fidelity CFD simulations tools (e.g. Schepers et al. (2007), Annoni et al. (2014b), Goit et al. (2014)), with a wide range of reported potential increases (Knudsen et al. (2014) gives an overview). However, not all the simulation studies report a beneficial effect: such a counterexample is the LES

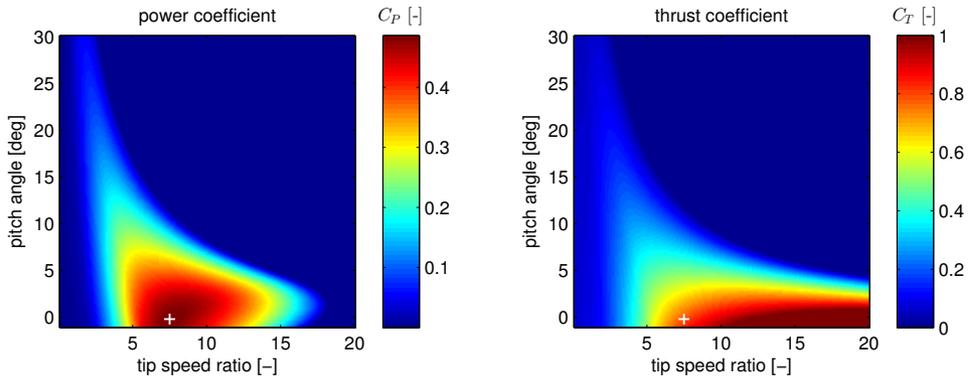


Figure 2.5: Power and thrust coefficient of the NREL 5MW reference turbine Jonkman et al. (2009) as a function of blade pitch and TSR. In both surfaces, the cross (+) indicates the operation point of maximum C_p . The C_p surface is smooth so around the maximum, there is a small sensitivity of the power with respect to the TSR and pitch. At the same operation point, the gradient of the C_T is larger, and thereby also the sensitivity of the thrust is larger. Data from Soltani et al. (2010).

study of the Lillgrund wind plant in Nilsson et al. (2014), where axial-induction control using a range of pitch offsets on the front turbines is tested, and a production increase on the downstream turbines is reported, but an increase of the total power production is not achieved since the gain in production on the downstream turbines does not exceed the loss on the front turbines.

- Also, wind tunnel experiments have been performed with scaled turbines. In the two-turbine tests in Adaramola et al. (2011) and the scaled wind plant test in Machiels et al. (2007), the results from such tests are an increase in total power production of turbine rows when reducing the power extraction of the front turbines using pitch control.
- Experimental data from wind plants with full-scale industrial wind turbines is scarce, as the overview in Knudsen et al. (2014) shows. A relevant analysis is found in Machiels et al. (2007), where a systematic increase of power production is reported on the first two turbines of a row of 2.5 MW turbines with 3.8 rotor diameter spacing when applying the axial-induction-based control by applying a pitch offset on the front turbine.

In general, one notices a wide range of reported production increases through axial-induction-based control, including a case where no increase can be found.

2.3.3. EVALUATION OF AXIAL-INDUCTION-BASED CONTROL STRATEGIES USING SOWFA

In this section, we present a simulation counterexample to the examples of successful implementation of axial-induction-based wind plant control for wind plant electrical power production increase listed in the previous section, in order to demonstrate some of the difficulties arising in this type of control. In this example it is shown that even

when two turbines are aligned in the wind direction at a relatively short distance, under certain realistic inflow conditions there may not be a significant increase in power achievable by deviating from the turbine-level optimal control settings.

In Figure 2.6a the results are shown of applying different pitch offsets on the front turbine of the wind-aligned two-turbine setup with 7 rotor diameter spacing described in Section 2.2.2. It shows that the turbine-level power optimal setting (zero pitch offset) also yields maximum power production for the total wind plant. Although the effect of reduction of the axial induction on the front rotor causes an increase on power of the second turbine in the row, the power lost on the first turbine by offsetting the pitch is not regained on the second turbine. Note that a pitch offset also influences the tip-speed ratio, but the sensitivity of the C_P and C_T on the tip-speed-ratio is small relative to the sensitivity to pitch (see Figure 2.5).

In Figure 2.6b, the results are shown of simulation cases in which there is a deviation from the turbine-level optimal generator torque settings by applying a scaling factor α on the regular below-rated rotor speed control law (Jonkman et al., 2009) of the front turbine, so that the applied generator torque is $T = \alpha * K * \omega^2$ with $K = 0.0179 \text{ Nm/RPM}^2$ the turbine-level optimal gain for maximum power production¹. A decrease in tip-speed-ratio is needed to decrease C_T and thereby the axial induction factor of the front rotor, which can be seen in the C_T -surface in Figure 2.5. This decrease in tip-speed-ratio is achieved when the generator torque of the upstream turbine is increased ($\alpha > 1$). As mentioned before, the C_T -coefficient is not very sensitive with respect to the tip-speed-ratio, and the possibility to increase the generator torque is limited by the fact that the rotor may stall when a temporary reduction of wind speed occurs, therefore the amount by which axial induction can be affected by adjusting torque is limited. When increasing torque on the upstream turbine, an increase of power at the downstream turbine can be observed. However, as in the pitch case, there is not enough power increase on the downstream turbine to compensate for the loss on power production on the front turbine, and a decrease in electrical power production on the total two-turbine setup results.

In Figure 2.7 the results are shown of a study to investigate the reason why in this particular two-turbine case the axial induction control concept does not yield a wind plant power production increase. For a single turbine setup, we compare a two degree pitch offset on the turbine to the baseline case with no offset. From the results of the flow simulation, the kinetic energy added to the wake by pitching the turbine is calculated. In Figure 2.7a the differences are shown between the pitch offset case and the baseline case, in the kinetic power density P_{density} of the wind flowing through cut-through planes downstream of the rotor, calculated as:

$$P_{\text{density}} = u_{\text{axial}} \left(\frac{1}{2} \rho \vec{U} \vec{U}^T \right) \quad (2.6)$$

where u_{axial} is the axial component of the velocity (along the rotor axis), and \vec{U} is the flow velocity vector, and ρ is the air density. When visualizing the rotor plane of a 'vir-

¹We found that the generator torque gain K to reach optimal tip-speed ratio for maximum power production in SOWFA simulations deviates from the optimal value $K = 0.0256 \text{ Nm/RPM}^2$ in Aerodyn simulations reported in Jonkman et al. (2009)

tual' downstream rotor of equal size placed downstream aligned in the wind direction in Figure 2.7a, it can be seen that the kinetic energy conserved in the flow by using a sub-optimal pitch angle on the turbine, is mostly going outside of that rotor plane. Therefore, the power reduction on the front turbine would mostly be 'lost' to speeding up the flow surrounding a downstream rotor, rather than increasing the power production of downstream turbine. The associated energy balance is made in Figure 2.7b, where the added power flow through a potential downstream turbine is summed over the 'virtual' rotor plane and compared to the power lost at the upstream turbine by pitching. When considering that the downstream turbine can operate at a maximum $C_{P,\max} = 0.48$, this balance predicts that it would not be possible to recover the energy lost by placing a turbine of equal size downstream at a realistic spacing (at least more than one rotor diameter). A second cause for a limited ability to improve production at the downstream turbine through control offsets on the upstream turbine, is suggested in Annoni et al. (2014a): a reduction in thrust force on the front turbine may reduce turbulence in the wake, and thereby the wake recovery, which has a negative effect on the velocity at the downstream turbine.

2

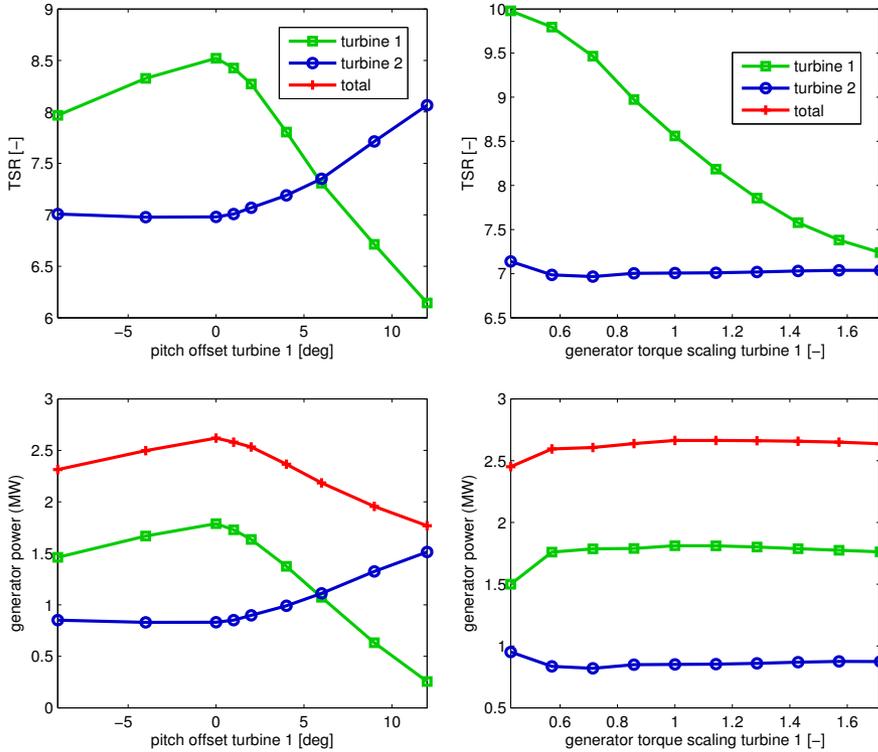
In this particular simulation case, the results suggest that the combination of the C_P and C_T turbine characteristics and the expansion and recovery properties of the wake in the simulated flow make the concept of power increase through axial-induction-based control optimization through adjusting pitch or generator torque infeasible.

2.3.4. DISCUSSION

The examples of successful implementation of the axial-induction-based wind plant control concept in simulation and scaled, and full-scale experiments, listed in the literature overview in Section 2.3.2, combined with the high-fidelity simulation counterexample in Section 2.3.3, show that the plant-level optimal settings and the potential gain from plant-wide instead of turbine-level optimized control are dependent on the particular inflow conditions, the wind plant configuration and the turbine characteristics. Plant-wide axial induction control for power optimization will only have potential if there is enough wake overlap and the wake recovery and expansion is small enough so that the energy 'sacrificed' by the upstream turbines to the downstream flow will not be lost to the flow passing the downstream turbines. Also, the limited efficiency of the downstream turbine in converting the increase in kinetic energy of the inflow to electrical energy (the power coefficient), has to be taken into account when evaluating the potential of axial-induction control methods.

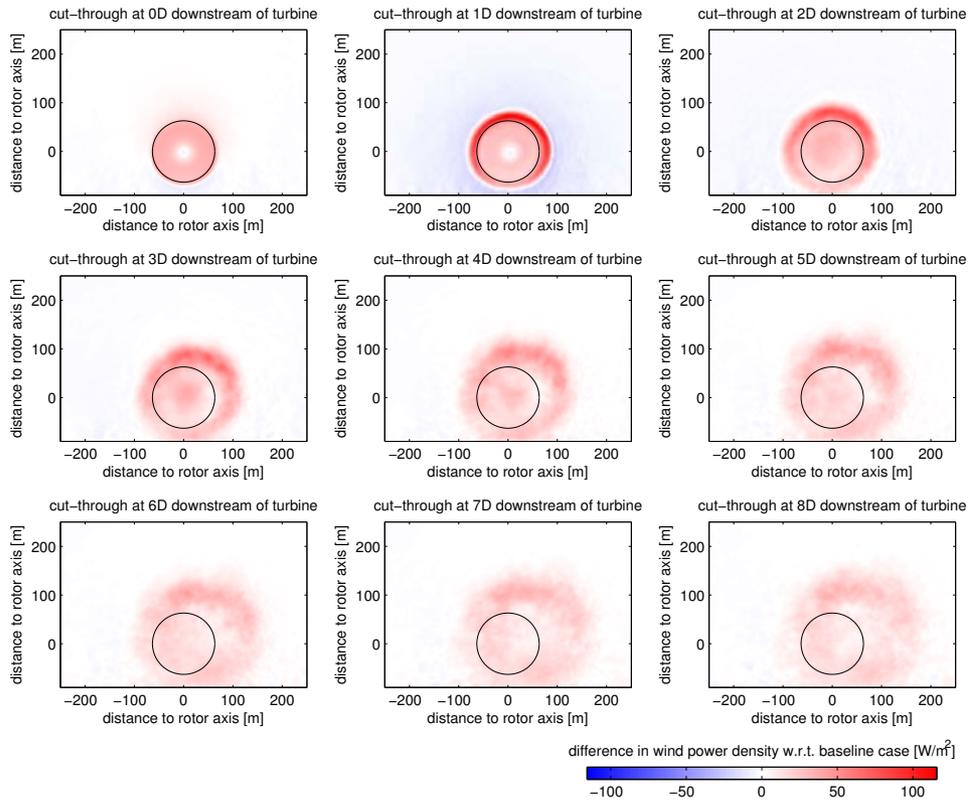
We conclude that for the axial induction control concept to be applied, an optimization strategy for the control settings on the turbines is needed that is adaptive to the particular atmospheric circumstances and the turbine characteristics at a certain operating point. In Chapter 3 we present such an adaptive approach (a direct data-driven approach) that is based on the available measured data, and knowledge of the lay-out of the wind plant.

An alternative control method is to use data-driven model-based wind plant control methods, in which control-oriented engineering models are used to perform the optimization (as proposed in Section 1.4). Although engineering wake models, such



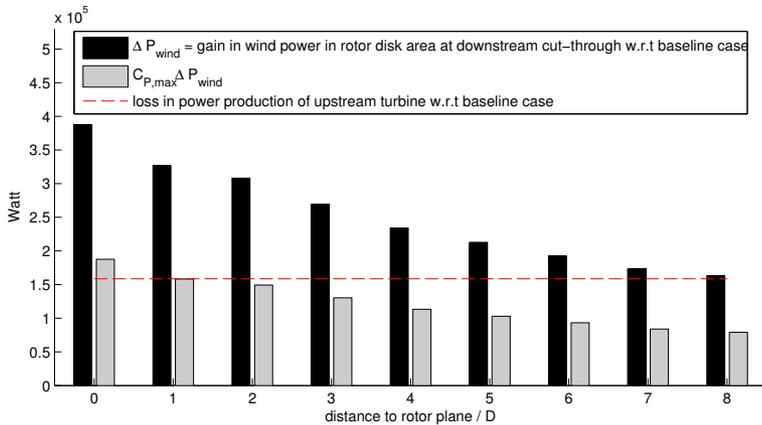
(a) Power results of changing pitch on upstream turbine 1 (b) Power results of changing generator torque on upstream turbine 1

Figure 2.6: SOWFA Simulation results of using the axial-induction control concepts on a two-turbine setup, using pitch (left) or generator torque (right) offsets on the front turbine (1) to affect the power on the downstream turbine (2). The total power is not improved by introducing offsets.



2

(a) Kinetic power added to the flow behind the NREL 5-MW turbine by introducing a 2° pitch offset. The black circle visualizes the location of a second turbine rotor that could potentially be placed downstream, aligned in the wind direction.



(b) Power lost on the turbine by introducing the pitch offset (red line) versus power added to the flow through the plane of a potential rotor downstream, relative to the baseline case with no pitch offset (black bar), and possible portion of that power to be recovered by the downstream turbine (gray bar).

Figure 2.7: Power balance for the wake of a single turbine with a 2° offset. The power lost on the front turbine by offsetting the pitch, cannot be recovered by a second turbine operating at maximum C_P that is placed downstream at a realistic spacing.

as the Jensen (Jensen, 1984; Katić et al., 1986), or the Frandsen model (Frandsen et al., 2006), take into account wake expansion and the turbine efficiency, the particular parameters used in these models will affect the predicted results in terms of power production increase potential. These parameters need to be adapted to the atmospheric conditions and the turbine characteristics. Comparing the simulation results presented in Section 2.3.3 and 2.3.2, however, suggest that there is some discrepancy between high-fidelity models and the engineering-type models in the prediction of the effect of axial-induction based control. Based, in part, on the simulation results in this section, and the engineering model presented in Chapter 4, Annoni et al. (2014a) discusses further these discrepancies, and proposes extensions to the engineering model to resolve them.

2.4. EVALUATION OF WAKE REDIRECTION AND TURBINE REPOSITIONING METHODS

An alternative approach to controlling wakes in wind plants is to redirect the wake rather than to only optimize induction. A SOWFA simulation-based evaluation of how much wake redirection can be achieved on a single turbine through different techniques can be found in Section 2.4.1. Then, in Section 2.4.2 we investigate the effects of the wake redirection techniques on a second turbine standing downstream of another turbine to which we apply the wake redirection techniques.

2.4.1. EVALUATION OF WAKE REDIRECTION CONTROL STRATEGIES FOR A SINGLE-TURBINE CASE USING SOWFA

In this section, first we evaluate the effect of different control strategies on the properties of the wake.

2.4.1.1. DESCRIPTION OF THE WAKE REDIRECTION STRATEGIES

In this section we consider three methods of achieving wake redirection, which are illustrated in Figure 2.8:

- By using the yaw drive to rotate the rotor and nacelle around the tower, the rotor is misaligned with the wind direction, which makes that the flow direction of the wake is changed in the horizontal direction, in the direction opposite to the yaw angle γ .
- By using actuation to rotate the rotor and nacelle in the tilt direction, where the rotor ‘leans forward or backwards’, the rotor is also misaligned with the wind direction. In this case the flow direction of the wake is changed in the vertical direction, in the direction opposite to the tilt angle τ .
- Finally, the pitch actuation of the rotor blades, that rotate the blades around their longitudinal axis, is used to induce wake deflection. We follow a novel Individual Pitch Control (IPC) strategy to induce an imbalance of the forces of the blades on the flow, causing redirection.

The proposed methods are discussed in more detail below:

Wake redirection through rotor yaw and tilt In conventional wind turbine control, the yaw drive is used to keep the rotor facing into the wind as the wind direction changes, the yaw misalignment strategy is thus an alternative way to use an existing control DOF of modern wind turbines. The tilt angle is not a controllable feature in present-day wind turbines. Still, knowledge of the capability of this possibility for wind plant control might be useful for novel wind turbine designs.

A simplified physical explanation for each wake redirection strategy is illustrated in Figure 2.8. For the yawing and tilting case, the thrust force f exerted by the rotor on the flow in the rotor axis direction is shown. When the wind inflow is at an angle to this direction, the thrust can be divided into components. In the yawing case these components are f_x and f_y . Force component f_x is parallel to the flow and slows down the wind, while force component f_y is in the cross-wind horizontal direction and causes the flow in the wake to deflect. Likewise, the thrust force of a tilted rotor has a vertical component f_z that causes wake deflection in the vertical direction. Since the thrust force component f_x is smaller in the yaw case, also the axial induction of the rotor decreases with increasing the yaw or tilt offset, which then increases the wind velocity in the wake. The above explanation is simplified, since in addition to the rotor force f , there is also a component of the rotor forces in the direction normal to the rotor axis. A more detailed analysis based on vortex cylinder theory is found in Burton et al. (2002b).

The yaw wake redirection method has been studied experimentally in wind tunnel tests with scaled turbines in Adaramola et al. (2011) and Medici (2005). Also, field tests with kW-scale turbines were performed in Wagenaar et al. (2012) with encouraging results, although the data was found to be too scattered to make clear conclusions. A correlation between yaw offset and a higher wind velocity downstream was demonstrated on a MW-scale turbine in Soleimanzadeh et al. (2014). In addition, the yaw wake redirection method was tested in CFD simulation in Jiménez et al. (2010). In a similar manner, vertical wake redirection through tilt has been investigated in Guntur et al. (2012) using a CFD model. Differences with the simulation study presented hereafter are that both Jiménez et al. (2010) and Guntur et al. (2012) use an actuator disk model of the turbine, and that the analysis in Guntur et al. (2012) assumes laminar flow.

Wake redirection through Individual Pitch Control Pitch control of the individual blades is a feature available in state-of-the-art modern wind turbines, that is generally used to mitigate loads on the rotor. In the IPC method evaluated in this section, we use IPC in an unconventional manner based on the notion that we can intentionally induce rotor force imbalances on the flow by pitching the blades in a coordinated manner. This IPC method is based on the IPC scheme using the Coleman transformation as specified in Houtzager (2011) which in turn is based on the initial idea in Bossanyi (2009). The difference with the original IPC scheme of Houtzager (2011), that is aimed at blade load reduction by removing yaw and tilt moments on the rotor, is that we intentionally create these moments on the rotor. This is done by introducing reference signals for the yaw and tilt moments, denoted respectively by $M_{r,yaw}$ and $M_{r,tilt}$ in the scheme in such a way

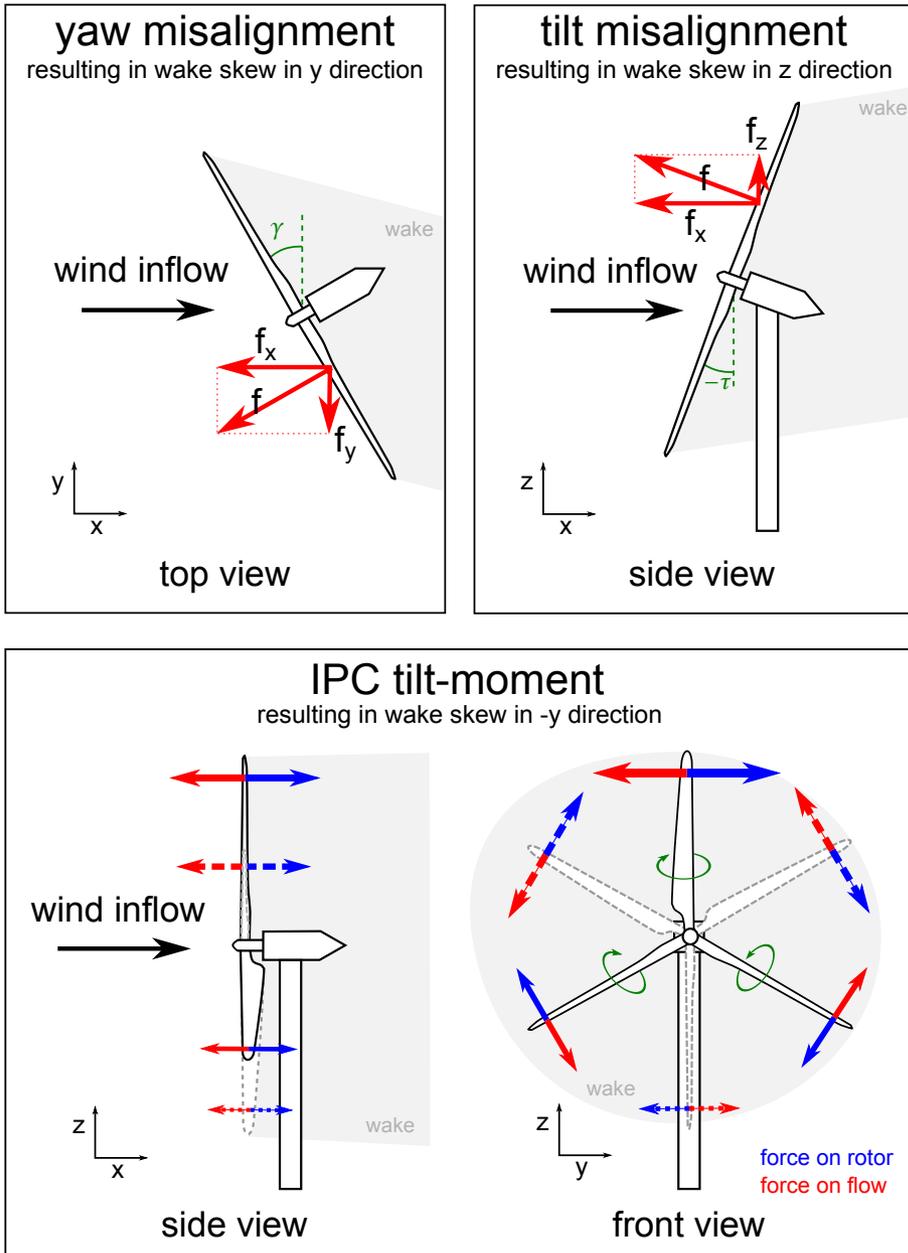


Figure 2.8: Demonstration of the mechanisms for yaw-, tilt- and IPC-based methods for achieving wake redirection. The figures on the top, based on the explanation given in Jiménez et al. (2010), shows the thrust force decomposed into components. In each case, one of these components causes wake redirection in a crosswind direction. IPC, shown on the bottom, generates an asymmetric rotor torque which can yield a force on the flow in the crosswind direction.

that they are tracked by the integral control in the Coleman-transformed domain. Further, we adjust the implementation such that it is active in below-rated operation with a varying rotor speed. In this IPC implementation the once-per-rotor-revolution (1P) and twice-per-revolution (2P) additive adjustments to the pitch, $\{\delta\theta_{i,jP}\}_{i=1}^3$, are given by:

$$\begin{aligned} \begin{bmatrix} \delta\theta_{1,jP} \\ \delta\theta_{2,jP} \\ \delta\theta_{3,jP} \end{bmatrix} &= L(s) P_{jP}(\varphi + \delta_{jP}) \begin{bmatrix} \frac{K_{i,jP,yaw}}{s} & 0 \\ 0 & \frac{K_{i,jP,tilt}}{s} \end{bmatrix} \begin{bmatrix} \delta\theta_{1,jP} \\ \delta\theta_{2,jP} \\ \delta\theta_{3,jP} \end{bmatrix} \\ &\times \left(\frac{2}{3} P_{jP}^T(\varphi) N_{jP}(s) \begin{bmatrix} M_{y,1} \\ M_{y,2} \\ M_{y,3} \end{bmatrix} - \begin{bmatrix} M_{r,yaw} \\ M_{r,tilt} \end{bmatrix} \right) \end{aligned} \quad (2.7)$$

for $j = 1, 2$, with Coleman transformation matrices:

$$P_{jP}(\varphi) = \begin{bmatrix} \cos(j\varphi) & \sin(j\varphi) \\ \cos(j(\varphi + 2\pi/3)) & \sin(j(\varphi + 2\pi/3)) \\ \cos(j(\varphi + 4\pi/3)) & \sin(j(\varphi + 4\pi/3)) \end{bmatrix}, \quad (2.8)$$

and with inverse notch filters \mathcal{N}_{jP} that amplify the 1P and 2P frequencies, and low-pass filter \mathcal{L} :

$$\mathcal{N}_{jP}(s) = K_{jP} \frac{2\zeta_{jP}\omega_{jP}s}{s^2 + 2\zeta_{jP}\omega_{jP}s + \omega_{jP}^2} I_{3 \times 3}, \quad \mathcal{L}(s) = \frac{\omega_L^2}{s^2 + 2\zeta_L\omega_L s + \omega_L^2} I_{3 \times 3}, \quad (2.9)$$

with $\omega_{jP} = j\varphi$, where the gain factors K_* , damping factors ζ_* , low pass-band frequency ω_L and phase offsets δ_{jP} are parameters of the controller.

When setting nonzero references $M_{r,yaw}$ or $M_{r,tilt}$, this IPC implementation creates an uneven distribution of thrust forces on the rotor blades over the course of a rotation (see Figure 2.8) and a tilt or yaw moment on the turbine rotor. Still, the thrust reaction forces on the flow are directed parallel to the inflow direction. Therefore, the uneven distribution of the thrust forces can cause differences in velocities in the wake, but it does not cause significant redirection of the flow. However, as shown in the turbine front view in Figure 2.8, IPC also causes the blade torques to be uneven over the course of a rotor rotation (in the sense that rightward torque is not matched by leftward torque). Therefore the in-plane reaction forces of the rotor on the flow are also unbalanced resulting in the fact that the turbine applies a net force on the flow perpendicular to the thrust direction, which does cause the flow to be redirected and the wake structure to be skewed. Notice that the IPC configuration drawn in Figure 2.8, which depicts the case when we apply a nonzero reference $M_{r,tilt}$, will yield a tilt moment on the turbine rotor (because the blade thrust is most different between the top and bottom azimuth positions) but a wake skew in the horizontal direction (because the reaction forces of the rotor on the flow are most different in the horizontal direction). Similarly, applying a nonzero reference $M_{r,yaw}$ will yield wake redirection in the vertical direction.

2.4.1.2. SIMULATION SETUP

In this study, the different proposed methods explained in Section 2.4.1.1 are applied on the single turbine scenario described in Section 2.2.1 in individual numerical simulations with SOWFA, each with different settings of yaw misalignment, tilt angle, or IPC

moment set-point. The wind inflow is the same for all simulations. From the simulations, we extract the turbine's average power over the simulation, as well as the metrics of loading for several components. From the flow, we use a correlation method to identify the wake-center at all locations downstream from the turbine. The results allow for a trade-off analysis of wake redirection potential and the effects on the turbine.

In each case, the turbine uses the baseline controller defined in Jonkman et al. (2009) for collective pitch and torque control. In the IPC-induced moment test cases, the IPC implementation of equations 2.7-2.9 is used. The IPC parameters $K_p, \zeta_p, \omega_L, \delta_{jP}$ as specified for the NREL 5-MW in Houtzager (2011). In these test cases, either $M_{r,yaw}$ or $M_{r,tilt}$ are chosen positive or negative, with the magnitude large enough such that the pitch angles vary with maximum amplitude, in order to find the maximum effect of IPC action on the wake. The filters are used in a Tustin discretized form with a sample time of 0.02 s. The pitch angles are saturated to a 5 degree amplitude, and the pitch rates are limited to 8 deg/s.

2.4.1.3. ANALYSIS AND DISCUSSION

A method developed at NREL is used to determine the mean wake center in the rotor hub-height horizontal plane at each downstream location of the turbine rotor: first the velocity profile in this plane are averaged over the total simulated time, then a Gaussian shape is fitted to the time-averaged velocity profile in the crosswind direction, and finally the maximum of that Gaussian is taken as the wake center position. Using the velocity profile in the rotor-axis-centered vertical plane, the same process is followed to find the vertical wake deflection; however, first the vertical profile of mean velocity of the flow not affected by a turbine is subtracted to remove the effect of vertical shear that is present in the atmospheric boundary layer. Figure 2.9 shows the result of the wake center-line identification algorithm in the horizontal plane for the yaw cases, Figure 2.10 shows the results in the vertical plane for the tilt cases, and Figure 2.11 shows the results for the IPC cases.

It should be remarked that because the velocity profile in the wake does not follow a perfect Gaussian shape, the method to track the wake center is somewhat arbitrary, although it gives insight into the mechanism of the wake redirection methods.

Wake center tracking results in the yaw and tilt simulations demonstrate significant displacements of wake center in the opposite direction to the yaw or tilt angle. In Figures 2.9 and 2.10, the wake displacements resulting from yaw at seven rotor diameters (7D) from the turbine, a typical location for a downstream turbine in a wind plant, are 'highlighted' and shown to go up to approximately $0.5D$ for a positive yaw offset of 40 degrees. Note that in the baseline wake centerline contours, there is also a small wake deflection as a consequence of the rotor rotation effects mentioned before in Section 2.2.1. For each of the methods, at larger distances, the direction of the wake will recover to the free-stream flow direction.

The advantage of using tilt angle in the positive direction, as we will later demonstrate in Section 2.2.2, is that deflecting the wake downwards in this way, promotes the attraction of higher-speed flow from the upper atmosphere, which would further increase the downstream velocity which could potentially increase further the power production of

downstream machines. It should be pointed out that to our knowledge, in industrial MW-scale wind turbines, currently there is no means of modifying tilt angle. Also, note that since positive tilt angles would cause the blades to come closer to the tower on up-wind turbines, this technique would be better suited to downwind machines.

In addition to measurements of the wake, data were collected from the FAST turbine output, in order to study the effects on power production and turbine loads. Here, for the yaw and tilt strategies we refer to the next Section 2.4.2 for the main results in terms of load and power characteristics, since these results do not significantly change when adding a second downstream turbine.

With the IPC-based methods significant wake skew is achieved for some cases, although not as large as can be achieved with rotor yaw or tilt misalignment. The results show that the largest vertical skew is obtained when inducing a high rotor yaw moment $M_{r,yaw}$ and the largest horizontal skew is created when using a high tilt moment $M_{r,tilt}$. This can be explained using the reasoning in Section 2.4.1.1.

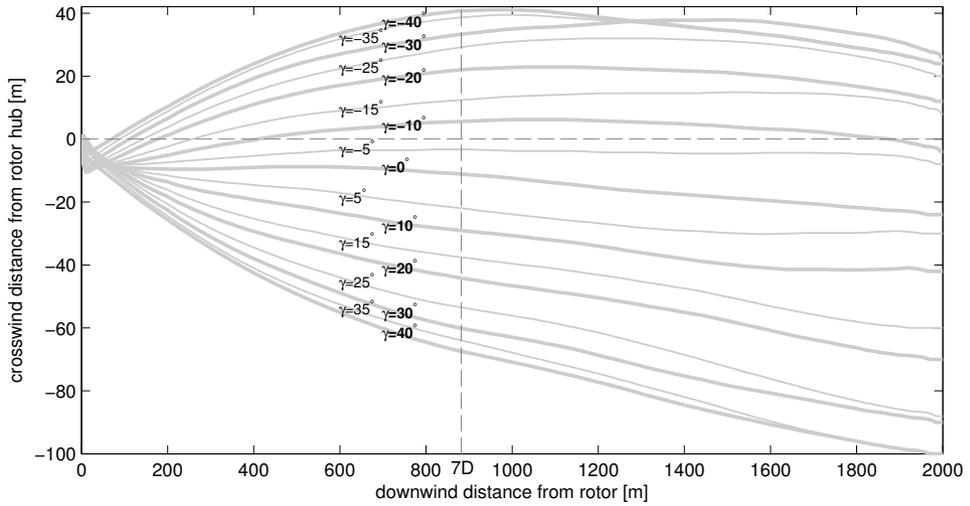
However, if we take into account the effects on the turbine itself, the result of the IPC-based method are mixed. Because the method is maximizing an asymmetric rotor moment, the blade loads are substantially increased (with up to 130%, cf. Fleming et al. (2014d)). This result implies that while it may be possible to achieve wake redirection with IPC, this particular IPC algorithm is not a good method. Finding an IPC controller that achieves wake skew with reduced blade loads would be very useful because IPC is already possible to implement on many existing turbines (unlike changes to tilt), and can be adjusted more quickly than yaw angle. However, finding such an IPC implementation is expected to be problematic, because a moment applied by the rotor blades on the flow through pitch will inherently result in a reactive moment on the rotor itself, which then results in an increase in loads. Because of the limited potential of the IPC wake-redirection method, we will not further investigate the strategy in this thesis.

2.4.2. EVALUATION OF WAKE REDIRECTION AND TURBINE REPOSITIONING CONTROL STRATEGIES FOR A TWO-TURBINE CASE USING SOWFA

The objective of this section is to compare the two most successful wake redirection strategies (rotor yaw and tilt misalignment) in a SOWFA simulation of a two turbine setup. Further, we test the potential of a turbine repositioning strategy for wake mitigation. This latter method is explained in more detail below.

2.4.2.1. TURBINE REPOSITIONING STRATEGY

In the turbine repositioning method for wake mitigation, the turbines are assumed to be floating and therefore repositionable. The turbines can then be moved out of each others wake, as proposed in van Wingerden (2011), to reduce the effects of wakes on loading and power production. The concept requires adjustments of the wind turbine positions to wind direction and possibly also to other time-varying conditions that affect the wake interaction in an offshore wind plant (e.g. turbulence and wind speed). In a wind plant, cable length and other constraints on the ability to move the turbines will limit the possible turbine displacement, therefore in a first investigation of the concept



2

Figure 2.9: Wake centerlines in the horizontal hub-height plane as estimated with the Gaussian function-fitting method, for different yaw angles γ .

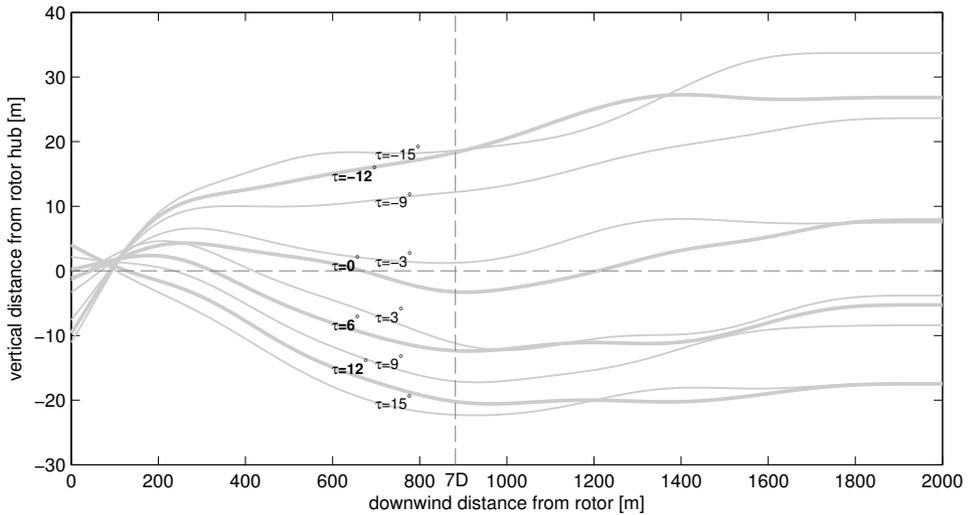
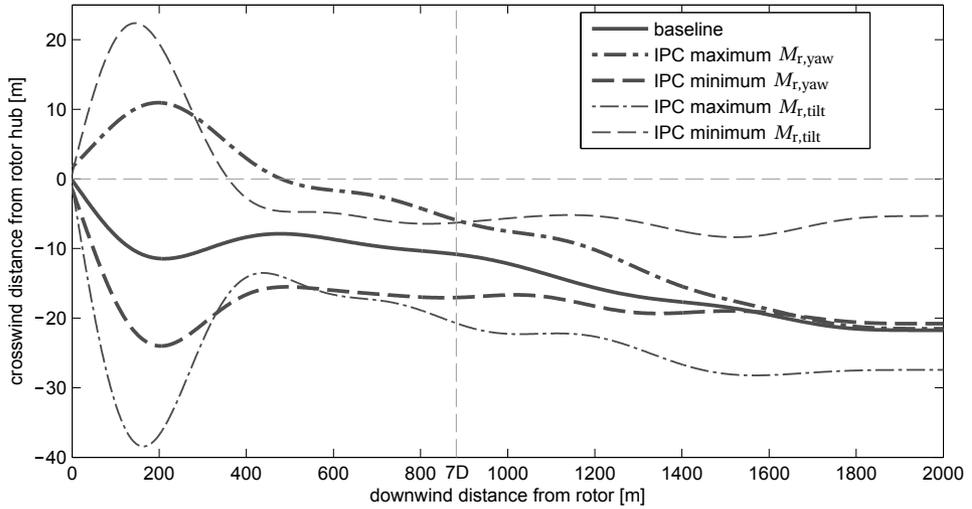
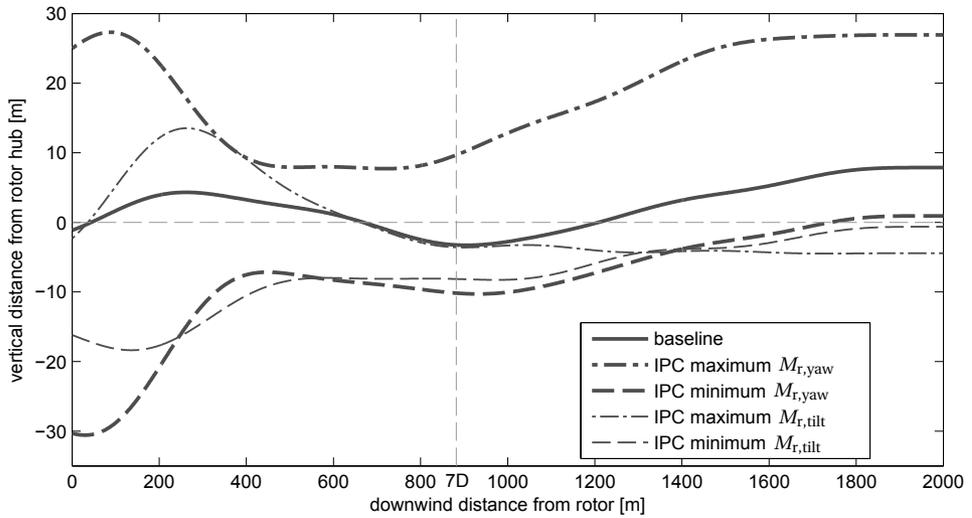


Figure 2.10: Wake centerlines in the rotor-axis-centered vertical plane as estimated with the Gaussian function-fitting method, for different tilt angles τ .



(a) Wake centerlines in the horizontal hub-height plane.



(b) Wake centerlines in the rotor-axis-centered vertical plane.

Figure 2.11: Wake centerlines as estimated with the Gaussian function-fitting method for different IPC-based strategies

it is important to know the size of the displacements that are needed to reduce the wake effects, and this is what will be investigated in this two-turbine case study.

2.4.2.2. SIMULATION SETUP

In simulations with the two-turbine scenario described in Section 2.2.2, the yaw, tilt, and position set-points of the upstream turbine are swept and held fixed for separate 1000-second simulations. In each case, the turbines use individual control for power and speed regulation, using the baseline pitch and torque control laws defined in Jonkman et al. (2009). Further, we investigate if the loads induced by moving from full to partial overlap can be mitigated through the use of load-reducing IPC with the standard collective pitch control. The IPC on the downstream turbine is switching on for the last 400 s of the simulation. As in the single-turbine case, the IPC implementation is based on the equations 2.7-2.9, but in these cases the set-points $M_{r,tilt}$ and $M_{r,yaw}$ are set to zero such that rotor loads are mitigated by the IPC. Again the parameters as specified in Houtzager (2011) are used, and we use the load-reducing IPC also in below-rated conditions. The SOWFA supervisory wind plant controller (described in Section 1.2.4.1) collects the data from the individual turbines.

Note that in floating turbines, additional hydrodynamics play a role, which are not currently modeled in SOWFA. Also, some control adjustments would be needed to account for the additional dynamics of floating turbines, as discussed in Van der Veen et al. (2012). Therefore, the scenario in which we test turbine repositioning control is somewhat artificial.

Figure 2.12 shows the flow fields for example cases for the different wake mitigation control strategies. It is shown that by misaligning the yaw angle of front turbine (T1) with the wind direction, the flow direction of the wake is changed in the direction opposite to the yaw angle, as was also seen in the single-turbine case study of Section 2.4.1. By controlling the deflection of the wake through yawing, the wake can be directed away from the downstream turbine (T2). In a similar manner, changing the rotor tilt on T1 redirects the wake in the vertical direction, so that its overlap with the rotor of T2 is reduced. Further, it is shown in Figure 2.12 that T2 can be moved out of the wake of T1. With each of the wake mitigation techniques, we aim at increasing the power production of T2 by increasing the velocity of the inflow into T2. The techniques reduce the overlap area of the wake of T1 with the rotor of T2, making that a larger part of the rotor encounters the higher free-stream velocity. Yawing or tilting also reduces the axial induction of the rotor of T1, and this further increases the wake velocity. In the case of yawing and tilting, a reduction of power production of T1 can be expected because the axial induction and the effective rotor area are reduced (Medici, 2005; Burton et al., 2002b). Changing from full overlap to partial wake overlap affects the turbulence and shear in the inflow to T2, which affects the loads on T2 (Yang et al., 2011). In the next part of this section, we focus on the effect of each of the wake mitigation techniques on the electrical energy productions of, and the load impacts on the turbines in simulation.

Following completion of the simulation, the data were collected from each case and post-processed as follows. First, the 1000 s of time domain data for each simulation were broken into segments (see Figure 2.13). The first 200 s of each simulation were discarded

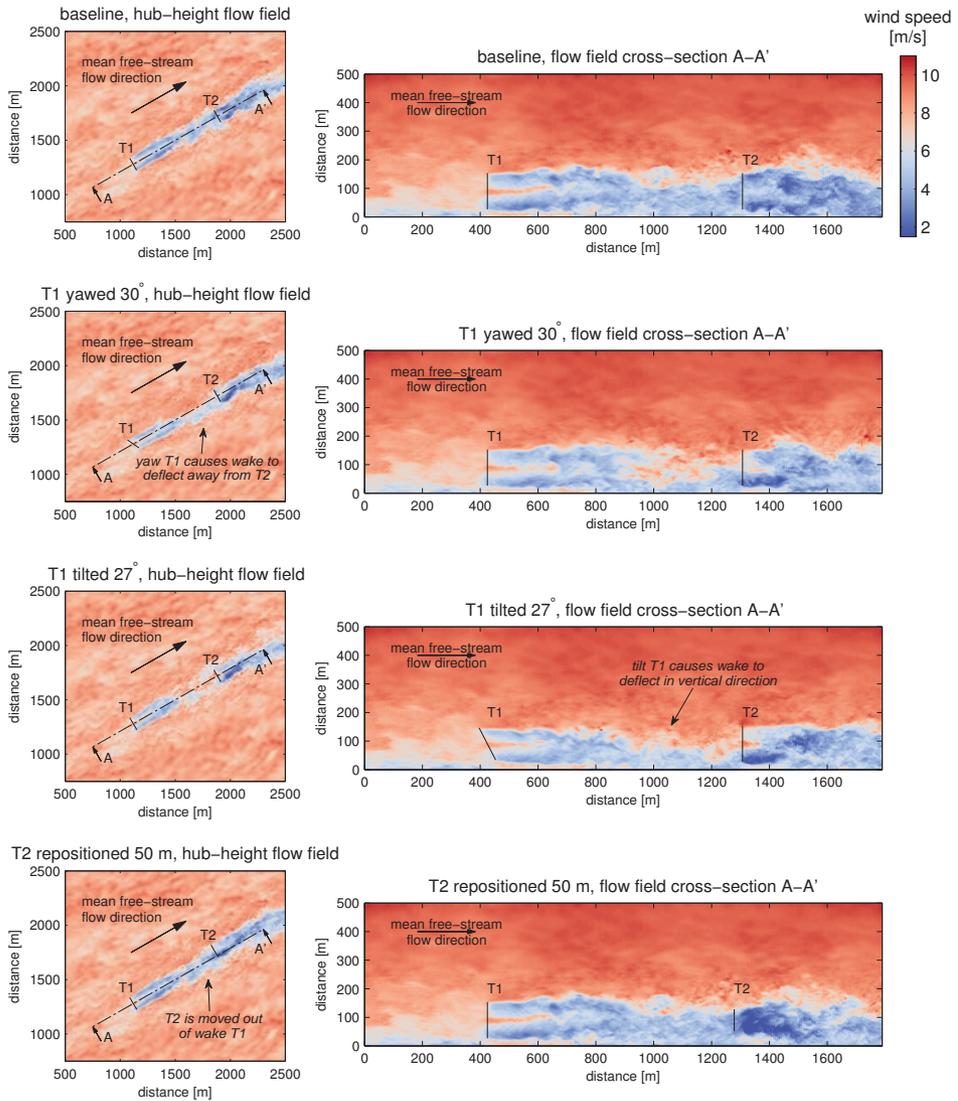


Figure 2.12: SOWFA-generated velocity fields after 1000s of simulated time, for example cases with each of the different wake mitigation methods. The rotor plane of the upstream turbine (T1) and the downstream turbine (T2) are shown as black lines.

because the wake was not fully developed. The last 100 s were also discarded because of system problems that left some files incomplete. Finally, the remaining time histories were divided into a segment from 200 s to 600 s, in which the downstream turbine is not running IPC, and a segment from 700 s to 900 s, when it is and the IPC startup transients have vanished. Although it should be possible to start IPC smoothly, because the transition was not our research focus, we start the controller rather abruptly. In the baseline case, IPC is never enabled, to provide a basis for comparison. From the wake development time and the velocity in the wake, it can be found that the average turbine-to-turbine flow-through time is approximately 200 s, and that we thus use the data of two flow-through periods for the analysis of the operation with IPC, and one flow-through period for the operation without IPC.

From these two blocks of time (200-600 s and 700-900 s), several metrics are computed. First, the average power is computed for each turbine. Next, loads are computed for blade out-of-plane (OOP) bending, drivetrain torsion, tower bending and yaw bearing moment. In the case of the tower load, a combined load is computed from the separate fore-aft and side-side loads using a root-sum-square combination. This is likewise done to combine the separate load components on the yaw bearing. For each of these load signals, a damage equivalent load (DEL) is computed. The DEL is a standard metric of fatigue damage (see Buhl Jr. (2008)). The power and DEL results are summarized in Figure 2.14.

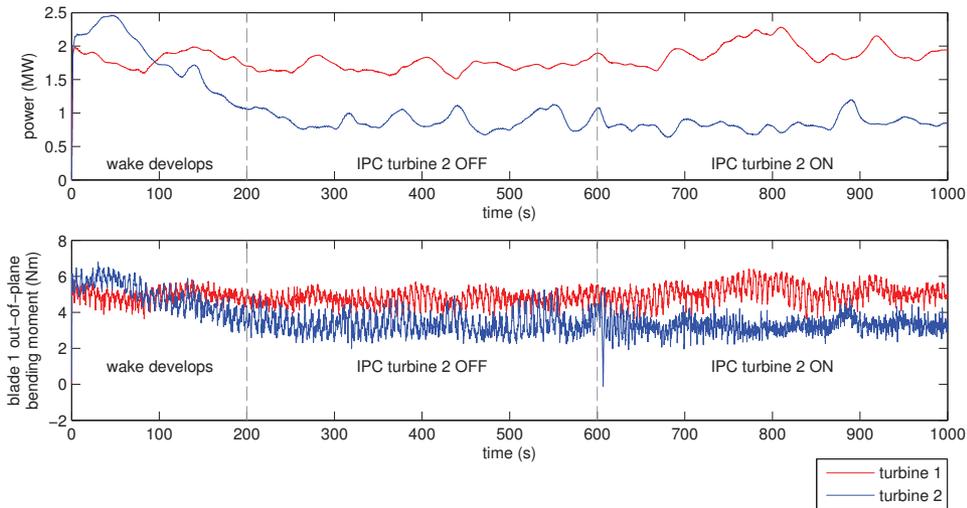


Figure 2.13: Power and blade 1 out-of-plane bending signals for the simulation case in which turbine 1 has a 5 degree yaw angle.

2.4.2.3. ANALYSIS AND DISCUSSION

In general, each of the concepts show good potential to increase electrical power production. Redirecting the wake with yaw and tilt on the upstream turbine or repositioning the downstream turbine, such that overlap of the wake of the upstream turbine with

the downstream is reduced and the wind velocity in the wake increased, indeed yields a power increase on the downstream turbine. With increasing yaw or tilt angle, the production of the upstream turbine is reduced, but the production gain on the downstream turbine is larger for a range of the misalignment angles. If we consider the simulations in which the downstream turbine does not use IPC (Figure 2.14a), in the best cases, yaw control shows an increase in total production of the setup of 4.6%, for tilt the maximum increase is 7.1%, and there is 41% improvement when the downstream turbine is moved a full rotor diameter.

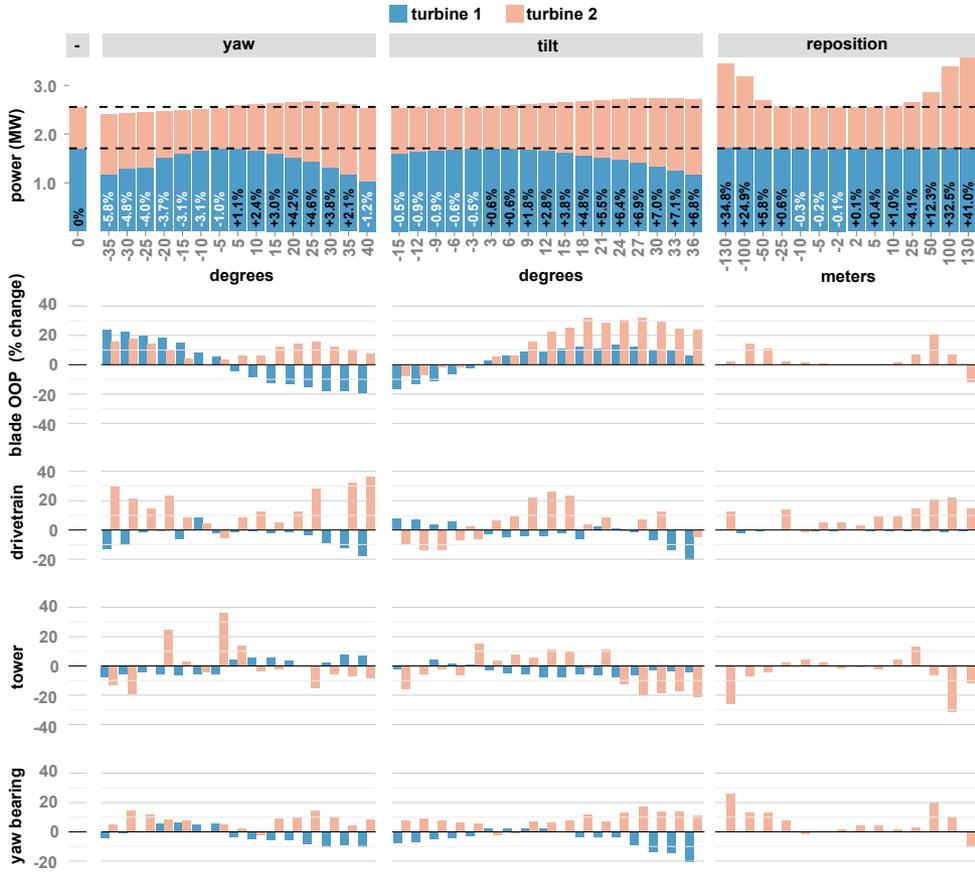
Tilt misalignment thus shows larger potential electrical power production improvements than yaw when considering large positive tilt angles. Both negative and positive tilt angles will redirect the wake away from the downstream turbine rotor, but the positive tilting has the advantage that it will redirect the wake towards the ground, allowing high velocity air from higher altitudes to flow towards the upper part of the downstream rotor, resulting in higher power production of the downstream turbine.

In the repositioning technique, for smaller displacements, there is little power production increase, and significant increase for larger displacements. This can be explained by the fact that the wake expansion is substantial, therefore the displacements need to be large for the wake overlap to be reduced. Also observing the loads for the downstream turbine, there is little change for small alterations in position, significant change for the displacements yielding partial overlap, and then no change again when the turbine is moved a full rotor diameter. The feasibility of this floating wind plant control concept thus strongly depends on the particular constraints on the displacements of the floating turbines in the wind plant.

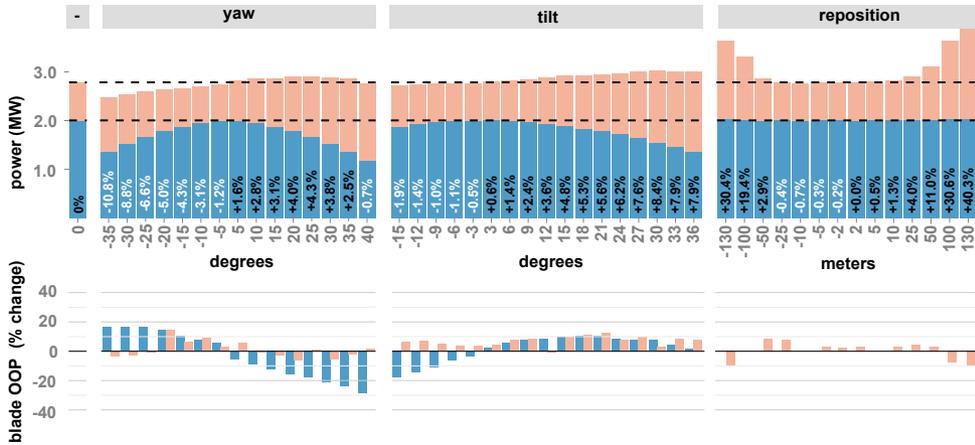
Looking at the loads across simulations without the use of IPC (also in Figure 2.14a), the upstream turbine either experiences an increase or decrease in blade OOP bending, depending on the angle chosen. Blade bending reduces with positive tilt angles and increases with negative yaw angles. This is in agreement with an other investigation (Kragh et al., 2013a) on this behavior in sheared inflow. For tilt, blade bending increases with positive offset angles and decreases with negative offset angles. The drivetrain, tower, and yaw load for the upstream turbine either decrease or minimally increase by yaw- or tilt-angle adjustments. A possible explanation for this effect is that these methods generally reduce the power capture of the upstream turbine, and derating can be considered a load mitigation strategy. For the downstream turbine, all loads generally increase somewhat, and this is most likely due to moving from full to partial wake overlap, which may increase the asymmetry of thrust forces on the rotor.

Figure 2.14b shows the results for power and blade loads for the simulations where the downstream turbine is using IPC to mitigate the effect of partial wake overlap. These results are based on 200 s of simulation versus 400 s in Figure 2.14, and are from a different point in the simulation. Overall, the results indicate a very strong motivation for the use of IPC, in general, and as a way to reduce the negative impacts of using wake-mitigation strategies, since the impact of changing from fully-, to partially-, to non-waked conditions can be reduced.

It is important to acknowledge the shortcomings of this current work:



(a) Time-averaged power and load results without using IPC.



(b) Time-averaged power and load results when using IPC on the downstream turbine.

Figure 2.14: Summary of results of two-turbine simulations. The three columns are divided by control action. The top row shows the combined power output for each case, compared to the baseline case on the far left. The remaining rows indicate the percent change in load compared to the baseline.

- First, because of computational/time constraints, results are based on simulations of one inflow case and one atmospheric condition. As in the axial-induction control cases, the results in terms of potential gain of these methods are expected to vary with the amount of wake recovery, which is dependent of the inflow and atmospheric conditions. We have studied a case with a large expected power gain, since the relatively small turbulence intensity in the inflow together with the neutral atmospheric stability, is expected to result in a relatively slow wake recovery.
- Further, the power and DEL calculations are based on lengths of time shorter than what would be standard practice for a normal wind turbine simulation regime. This is because of the large computational cost of the SOWFA simulations, and limited access to the computational resources.
- Additionally, this work is based on a baseline of almost full overlap of the wake of the upstream turbine with the rotor of the downstream turbine. This selection represents a ‘worst case’ condition to investigate the potential gain of applying these techniques. In practice, the more likely scenario will be to go from partially overlapped to less overlapped, or not overlapped, and in these cases the potential power increase on two turbines would typically be smaller. An example study is provided in Chapter 4.

2.5. CONCLUSIONS

In this chapter we have evaluated the potential of different control degrees-of-freedom of the wind turbine to control the wake and optimize the performance of the wind plant. We have subdivided the control strategies in axial-induction-based, wake redirection and repositioning techniques. In this section, we summarize our findings for each of the techniques.

In literature, studies on the axial-induction-based strategies have been reported, where simulations, wind tunnel experiments and field tests have been used to test the concept, with a large range of reported potential gains in terms of power production. In our own high-fidelity simulations, we have shown a case with a relatively low ambient turbulence (neutral atmospheric boundary layer, low surface roughness), in which axial-induction-based strategies are not successful, because there is too much wake expansion. We conclude that the potential gain is sensitive to the properties of the inflow, the atmospheric conditions, but also the particular turbine characteristics. Therefore, a data-driven control strategy is needed to successfully implement an axial-induction-based technique that is adaptive to the particular conditions in a wind plant at a certain time. A direct data-driven strategy will be presented in Chapter 3. Further developments of engineering models for axial-induction-based control using data-driven model-based control are presented in Annoni et al. (2014a).

For the wake redirection techniques we have shown in simulation examples of a particular flow case that there is significant potential to increase power production using yaw and tilt. Loads on the misaligned turbine are increasing or decreasing depending on the particular misalignment angle and direction used. Possible load increases on the downstream turbines because of partial wake overlap can be mitigated by the use

of standard load-reducing IPC on the downstream turbine. The novel concept of using of IPC to redirect the wake, by intentionally inducing an asymmetry in the forces on the flow, has also shown to work, but the resulting load increases on the turbine are substantial. As for the axial-induced-based strategies, for each of the wake redirection strategies, it is to be expected that the potential to reduce loads and increase power production is dependent on the particular atmospheric conditions, the wind turbine properties and the wind plant setup. A data-driven strategy for yaw-based wake redirection control is presented in Chapter 4.

Since both the axial-induction-based strategies, and the yaw-based wake redirection strategies were tested in the same simulated ambient turbulence conditions, the results suggest that there is more potential for yaw-based wake redirection wind plant control, than for axial-induction-based wind plant control, although further research is needed to quantify the exact sensitivity to the ambient conditions (inflow turbulence, atmospheric stability, surface roughness). Because turbulence promotes wake recovery, it seems apparent that the level of ambient turbulence is negatively correlated with the potential of wind plant control to improve wind plant performance, but the exact influence is not quantified yet.

For the repositioning techniques for floating wind plant, we have shown that if we consider the baseline case of alignment of two rotors in the wind direction, because of wake expansion, substantial displacements are needed (half a rotor diameter or more), to reduce the wake effects sufficiently to have significant effects on total power production. The potential of this concept is thus dependent on the constraints on possible displacement of the floating turbines in a wind plant setup.

3

AXIAL-INDUCTION-BASED OPTIMIZATION CONTROL FOR WIND PLANTS

This chapter presents a data-driven adaptive scheme to adjust the control settings of each wind turbine in a wind plant such that an increase in the total power production of the wind plant is achieved. This is done by taking into account the interaction between the turbines through wake effects. The optimization scheme is designed in such a way that it yields fast convergence, so that it can adapt to changing wind conditions quickly. The scheme has a distributed architecture in which each wind turbine adapts its control settings through gradient-based optimization, using information that it receives from neighbouring turbines. The novel control method is tested in a simulation of the Princess Amalia Wind Park. It is shown that the distributed gradient-based approach performs the optimization in a more time-efficient manner compared to an existing data-driven wind plant power optimization method that uses a game theoretic approach.

3.1. INTRODUCTION

The aim of control algorithms in modern wind turbines is to adjust the control degrees of freedom of the turbine, such as the generator torque and the pitch angles of the rotor blades, to changing wind conditions, with the aim of maximizing the energy capture of the wind turbine while keeping the structural loads on the turbine within acceptable limits. Nowadays, wind turbines are often placed with other turbines in wind plants to reduce use of space and costs of installation and maintenance. However, placing wind turbines in a wind plant introduces aerodynamic interaction between the turbines that affect the power production and loads on each turbine in the plant. These interactions effects are not taken into account in the current practice of wind turbine control design.

This chapter has been published in Gebraad and van Wingerden (2014b).

The aerodynamic interaction follows from the fact that a wind turbine converts the kinetic energy of an incoming wind flow into electrical energy, which results into the formation of a wake of turbulent slow moving air downstream of the rotor. As the wake travels downstream, the wake expands, and recovers to the free stream conditions because of mixing with the surrounding air. If another turbine is standing in the path of a wake that has not fully recovered, the reduced wind speed in this wake results in a lower electrical power production of the downstream turbine. Adjusting the control parameters of a wind turbine affects the extraction of power from the wind flow, and therefore the velocity deficit in the wake it produces. Therefore in wind plants, in which turbines are placed relatively close to each other, the wake effect causes a coupling between the control parameters of upstream turbines and the power productions and loads on downstream turbines. Cooperative control strategies that take into account the wake effect can be used to optimize the total power production of the wind plant. This is done by reducing the power production of the upstream turbines, in order to reduce the velocity deficit in the downstream wind flow, which increases the power production of the downstream turbines (Steinbuch et al., 1988; Schepers et al., 2007; Johnson et al., 2009). In a similar manner, it is possible use cooperative control to distribute the structural loads acting on the individual wind turbines in the wind plant more equally.

One approach to deal with wake interactions is to derive a model that describes the dynamics of the wind plant, and to use this model to synthesize control laws (Soleimanzadeh et al., 2013; Madjidian et al., 2011; Bitar et al., 2013)¹, or to directly calculate control actions using model predictive control techniques (Soleimanzadeh et al., 2012, 2011; Spudić et al., 2010; Heer et al., 2014). In this chapter we aim to develop a wind plant control method that is directly data-driven rather than model-based, in the sense that it makes direct use of measured data in order to optimize the control parameters of the wind turbines and adapt them to time-varying wind conditions, without using a predefined model that predicts the effect of each control action. Thereby the objective of the optimization is to maximize the total power production of the wind plant in below-rated wind conditions, although the method may be extended to perform load control, by including static load measures in the objective function, as was done in Soleimanzadeh et al. (2012); Madjidian et al. (2011).

The method that in this chapter is chosen to perform the data-driven adaptive control is known as Maximum Power-Point Tracking (MPPT). In previous work (Koutroulis et al., 2006), the MPPT method was used to optimize the power of a single wind turbine using a real-time closed-loop scheme, where the change of the power production of the turbine as a result of control changes are measured, and subsequently the control parameters are adapted in a direction that yields a power improvement. In this chapter, the MPPT method is extended in such a way that it optimizes the total power of a wind plant, by letting the wind turbines exchange information about their power production with other wind turbines in the wind plant.

In this work, the MPPT wind plant control method is made adaptive to time-varying wind speeds, by designing the algorithm in such a way that its objective function is the efficiency by which the wind plant converts the kinetic energy of the incoming wind

¹The wind plant model that is used in Madjidian et al. (2011) to synthesize control laws does not include interaction between turbines through the wakes.

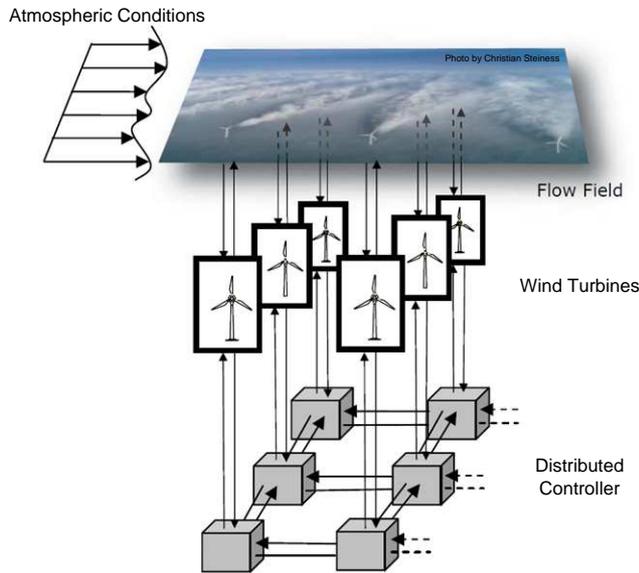


Figure 3.1: A distributed architecture for wind plant control.

into electrical energy, rather than the total power production of the wind plant. The wind plant control also needs to be adaptive to other changing wind conditions, such as a changing wind direction, or a changing turbulence intensity which will affect the amount of mixing with the surrounding air and thereby the wake recovery. To accomplish this, it is required that the optimization takes place in a time-efficient manner. To this end, gradient-based techniques are used to perform the optimization, and the algorithm is designed using a distributed architecture in which the control parameters of a wind turbine are adapted based on information from the nearest neighbouring turbines only. This architecture is illustrated in Figure 3.1. The proposed algorithm is an example of a direct data-driven method as described before in Section 1.4.

This chapter presents the new MPPT wind plant control method, and demonstrates its features through simulation examples in which the performance of the method is compared with a benchmark algorithm using the game theoretic approach of Marden et al. (2013). The simulation results are generated using the Jensen wake model, to which a delay structure is added to simulate the dynamics of the wake travelling through the wind plant.

The chapter is organized as follows. A full explanation of the MPPT wind plant control approach is given in Section 3.2. The game theoretic wind plant control approach is explained in Section 3.3. An explanation of the wind plant model used in the simulation examples is given in Section 3.4. In Section 3.5, the simulation examples are described in detail and the results are given. Finally, in Section 3.6 the conclusions are presented.

3.2. MAXIMUM POWER-POINT TRACKING CONTROL FOR WIND PLANTS

In this section, the MPPT wind plant control method is presented in two variants: the Gradient-Ascent MPPT (GA-MPPT) method and the Quasi-Newton MPPT (QN-MPPT) method. Both methods make use of gradient-based optimization techniques to find the control settings that yield a maximum total power production of the wind plant. As these gradient-based techniques can be categorized as being local optimization techniques, they may converge to a local maximum instead of a global maximum of the total power (Boyd et al., 2004). Further, in developing the MPPT methods simplifying assumptions are made in order to be able to perform the optimization using a distributed control architecture, in which each turbine uses information from the nearest neighbouring downstream turbine only. This gradient-based, distributed optimization approach is taken to improve the time-efficiency of the optimization.

The GA-MPPT and QN-MPPT methods are explained for a single row of wind turbines in Sections 3.2.1 and 3.2.2, respectively. In Section 3.2.3, the two methodologies are extended in such a way that they can be used on any wind plant configuration.

3.2.1. GRADIENT-ASCENT MPPT CONTROL OF A ROW OF WIND TURBINES

Consider a row of n wind turbines standing in the wake of each other, in a wind field with an incoming free stream speed V_∞ , as depicted in Figure 3.2. The turbines have power productions $\{P_i\}_{i=1}^n$ and certain control settings $\{a_i\}_{i=1}^n$ that influence the power production of the turbines. In the simulation examples in this chapter, it is assumed that the control variable a_i is the axial induction factor of turbine i . This is a generalization in the sense that in a modern turbine the axial induction factors can be influenced by adapting either the blade pitch angles, or by scaling the generator torque which changes the tip speed ratio of the rotor (Bianchi et al., 2007). Hence, in practice, one would use the MPPT scheme as presented here as a supervisory controller adjusting the reference signals for the blade pitch angles for each turbine, or the generator torque scaling factors used to adjust the rotor speed of each turbine. Alternatively, one could make use of knowledge of the power and thrust characteristics of a turbine rotor, if available, to find

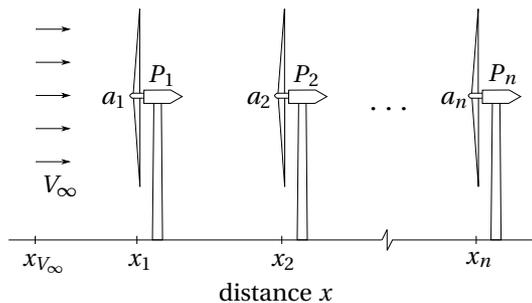


Figure 3.2: A row of n wind turbines with power productions $\{P_i\}_{i=1}^n$ and control parameters $\{a_i\}_{i=1}^n$, in a wind field with free-stream speed V_∞ .

the pitch and torque that yield a desired axial induction provided by a wind plant controller, as was shown in Schaak (2006). In both cases, each of the local turbine controllers would track the reference signals while fulfilling some requirements for safe operation (by limiting pitch rates for example).

Due to wake interaction, changing a control parameter a_i influences $\{P_j\}_{j=i}^n$, the power productions of turbine i and the turbines downstream of turbine i . Further, the power production of the turbines is dependent on the kinetic energy of the incoming wind field. The kinetic power of wind with air density ρ passing through an area A with speed V_∞ is given by Bianchi et al. (2007):

$$P_V = \frac{1}{2} \rho A V_\infty^3. \quad (3.1)$$

If constant air density is assumed, the efficiency by which the row of turbines converts the energy of the incoming wind field P_V into electrical energy can be maximized by solving the following optimization problem:

$$\max_{\{a_i\}_{i=1}^n} \sum_{i=1}^n \tilde{P}_i(a_1, a_2, \dots, a_i), \text{ with } \tilde{P}_i = \frac{P_i}{V_\infty^3}. \quad (3.2)$$

Note that in the above optimization problem, the power production of each turbine is divided by a factor that is proportional to the power of the incoming wind field. To perform the optimization in a local sense, the control parameters $\{a_i\}_{i=1}^n$ can be iteratively updated using a gradient-ascent optimization method (Boyd et al., 2004), resulting in the following Gradient-Ascent Maximum Power-Point Tracking (GA-MPPT) control update law:

$$a_i(k+1) = a_i(k) + K \sum_{j=i}^n \frac{\partial \tilde{P}_j}{\partial a_i}(k), \quad (3.3)$$

for $i = 1, \dots, n$, with index k denoting the iterations, and with a small scalar $K > 0$ being a scaling factor for the size of the steps on a_i . This design variable K can be used to tune the convergence properties of the gradient-ascent optimization. In order to perform the optimization in a real-time data-driven manner, the gradients $\frac{\partial \tilde{P}_j}{\partial a_i}$ can be approximated from the past iterations through first-order backward differencing:

$$\frac{\partial \tilde{P}_j}{\partial a_i}(k) \approx \frac{\tilde{P}_j(k) - \tilde{P}_j(k-1)}{a_i(k) - a_i(k-1)}. \quad (3.4)$$

A difficulty of the approach described above is that it takes a substantial amount of time to obtain the gradients of the objective function. For example, suppose that in iteration k , a_1 is changed by a certain step, then to find the gradient $\partial P_n / \partial a_1(k)$ using the above update rule, one would have to wait for the air in the wake of turbine 1 to travel to turbine n to find the effect of the control update on P_n , which is the power of the last turbine in the row. Because of the large distances between the turbines (typically 7 to 8 rotor diameters), the time this takes is very long for a large wind plant. During this traveling time, the speed of the incoming wind field is likely to have changed. Also other wind conditions may change over time, such as the wind direction, and the turbulence intensity in the free stream flow, which affects the amount of wake recovery in between the

turbines. To make the algorithm adaptive to time-varying wind conditions, two changes are made to the control scheme described above:

1. To overcome the problem of a changing speed of the incoming wind field, the delays related to the wind travelling from one turbine to the next are taken into account in the definition of the efficiencies \tilde{P}_i . Let $T_{V_\infty \rightarrow i}$ denote the time it takes for the wind field to travel from x_{V_∞} (the location where the incoming free-stream wind speed is measured) to x_i (the location of a turbine i). Then to compensate for the wind travelling delays, the efficiency \tilde{P}_i can be found from the power P_i at a time instant t , by:

$$\tilde{P}_i = \frac{P_i(t)}{V_{\infty,i}^{del}(t)^3}, \quad (3.5)$$

with:

$$V_{\infty,i}^{del}(t) = V_\infty(t - T_{V_\infty \rightarrow i}). \quad (3.6)$$

2. Changes in wind conditions such as wind direction and turbulence intensity change the way in which the power of each turbine is dependent on the control parameters of upstream turbines. A speed-up of the algorithm is needed for the optimization to be able to track these changes. A practical approach to speed-up the optimization scheme is to only take into account the influence of a turbine's control settings on the power of the turbine itself, and on the power of the neighbouring downstream turbine. This then results in the following control update scheme:

$$a_i(k+1) = a_i(k) + K \left[\frac{\partial \tilde{P}_i}{\partial a_i}(k) + \frac{\partial \tilde{P}_{i+1}}{\partial a_i}(k) \right]. \quad (3.7)$$

Only taking into account the effect on the downstream neighbouring turbine can be a good approximation as in practical cases there is a substantial reduction of the velocity deficit in the wake as the air travels from one turbine to the next, because of mixing with the free stream air. Thus, the effect of a control setting change is far larger on the nearest downstream neighbouring turbine than on turbines further downstream.

3.2.2. QUASI-NEWTON MPPT CONTROL OF A ROW OF WIND TURBINES

In order to further improve the convergence properties of the algorithm, one can use a Quasi-Newton optimization method to perform the local optimization of the control variables. When the vector notation $a = [a_1, \dots, a_n]^T$ is used for the set of control parameters, the Quasi-Newton update law reads:

$$a(k+1) = a(k) + KB(k)J(k). \quad (3.8)$$

As in the previous method, the scalar parameter $K > 0$ again determines the step-size. The matrix $J(k)$ represents an approximation of the gradient of the objective function, in which again only the gradients corresponding to the effect of a turbine's control parameter on its own power and on its downstream neighbouring turbine are taken into

account:

$$J(k) = \begin{bmatrix} \frac{\partial \tilde{P}_1}{\partial a_1}(k) + \frac{\partial \tilde{P}_2}{\partial a_1}(k) \\ \frac{\partial \tilde{P}_2}{\partial a_2}(k) + \frac{\partial \tilde{P}_3}{\partial a_2}(k) \\ \vdots \\ \frac{\partial \tilde{P}_{n-1}}{\partial a_{n-1}}(k) + \frac{\partial \tilde{P}_n}{\partial a_{n-1}}(k) \\ \frac{\partial \tilde{P}_n}{\partial a_n}(k) \end{bmatrix}. \quad (3.9)$$

As in the previous method, the gradients $\partial P_j / \partial a_i$ are approximated using (3.4). The matrix B represents an approximation of the inverse Hessian of the objective function. To avoid the inversion of an approximate Hessian matrix, the matrix B is directly approximated using the Davidon-Fletcher-Powell formula (Fletcher et al., 1963):

$$B(k) = B(k-1) + \frac{\Delta a(k) \Delta a(k)^T}{\Delta J(k)^T \Delta a(k)} - \frac{B(k-1) \Delta J(k) \Delta J(k)^T B(k-1)^T}{\Delta J(k)^T B(k-1) \Delta J(k)}, \quad (3.10)$$

with:

$$\Delta a(k) = a(k) - a(k-1), \quad (3.11)$$

$$\Delta J(k) = J(k) - J(k-1). \quad (3.12)$$

As a starting point of the above iterations, $B(0)$ should be set to a symmetric positive definite matrix, for example, the identity matrix.

3

3.2.3. MPPT CONTROL OF A WIND PLANT

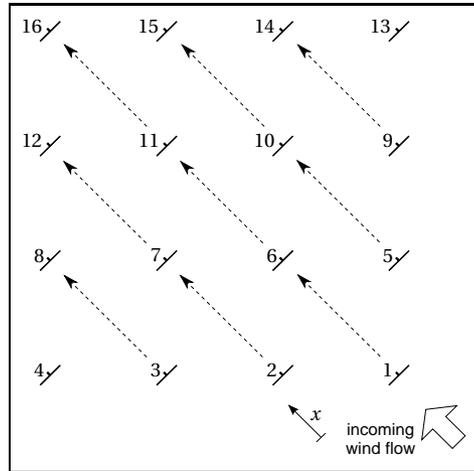
This section presents the scheme to perform the real-time closed-loop MPPT control on a wind plant of an arbitrary, but known spatial configuration, taking into account the delays between the control update and the power responses of the different turbines. Let $F = \{1, 2, \dots, N\}$ denote a set of indices that number the wind turbines in a wind plant, with N denoting the total number of turbines. Let $G \subset F$ be the set of turbines that are directly influencing neighbouring downstream turbines through wake interaction, and let $d(i)$ be the index of the nearest neighbour downstream turbine that a turbine $i \in G$ is directly influencing. Further, $L = \{i \in F | i \notin G\}$ is the set of turbines that are not influencing other turbines. In Figure 3.3 an example is given of how to define the sets F , G , and L and the mapping $i \mapsto d(i)$ for a given wind plant configuration and wind direction. It is assumed that the sets F , G , L can be updated using information of the wind plant configuration and the wind directions in the wind plant. Notice that an estimate of the wind direction is available in most wind turbines, as it is used to align the rotor axis with the wind direction through yaw control.

Using the above definitions, the GA-MPPT control update law for the wind plant is written as:

$$a_i(k+1) = a_i(k) + K \left[\frac{\partial \tilde{P}_i}{\partial a_i} + \frac{\partial \tilde{P}_{d(i)}}{\partial a_i} \right] \quad \forall i \in G, \quad (3.13)$$

and the QN-MPPT control update law for the wind plant is:

$$a_i(k+1) = a_i(k) + K s_i(k) \quad \forall i \in G, \quad (3.14)$$



turbine index $i \in G$	1	2	3	5	6	7	9	10	11
turbine $d(i)$	6	7	8	10	11	12	14	15	16
set $R(i)$	{1, 6, 11, 16}	{2, 7, 12}	{3, 8}	{5, 10, 15}	{6, 11, 16}	{7, 12}	{9, 14}	{10, 15}	{11, 16}

turbine index $j \in F$	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
turbine $\tilde{i}(j)$	1	2	3	4	5	1	2	3	9	5	1	2	13	9	5	1

Figure 3.3: The above picture shows the top view of a 4-by-4 wind plant in a south-eastern wind flow. The dotted arrows show which turbine is directly influencing which other nearest neighbour downstream turbine through wake interaction. In this case, the set of indices numbering each turbine is $F = \{1, 2, \dots, 16\}$. The indices of turbines i that are influencing other turbines are collected in the set G . The index of the neighbouring downstream turbine that a turbine $i \in G$ is directly influencing is given by $d(i)$. The mapping $i \mapsto d(i)$ is given in the table below the picture. The indices of turbines that do not influence other turbines are collected in the set $L = \{4, 8, 12, 13, 14, 15, 16\}$. The notation $R(i)$ is used for the full set of turbines that are standing in the wake of a turbine $i \in G$, and $\tilde{i}(j)$ is the most upstream turbine of the row that a turbine j is part of (both notations are used to define the wind plant simulation model in Section 3.4).

with $s_i(k)$ being the i -th element of the search direction vector $s(k)$, defined as:

$$s(k) = B(k)J(k), \quad (3.15)$$

with:

$$J(k) = \begin{bmatrix} J_1(k) \\ \vdots \\ J_N(k) \end{bmatrix} \quad \text{where } J_i(k) = \begin{cases} \frac{\partial \tilde{P}_i}{\partial a_i}(k) + \frac{\partial \tilde{P}_{d(i)}}{\partial a_i}(k) & \text{for } i \in G, \\ 0 & \text{for } i \notin G, \end{cases} \quad (3.16)$$

and $B(k)$ being generated by the update rule (3.10). In both MPPT methods, turbines in the set L are controlled to operate in such a way that the power of the turbine itself is maximized:

$$a_i = a_i^{\text{opt}} \quad \forall i \in L, \quad (3.17)$$

where a_i^{opt} is the control setting that yields maximum power production for the turbine i itself. In the case that a_i is the axial induction factor, $a_i^{\text{opt}} = 1/3$, (Bianchi et al., 2007).

After a control variable a_i is updated, there is a time delay before this change has an effect on the turbine i itself, and another time delay before the change has an effect on the neighbouring downstream turbine $d(i)$. The scalar $T_{s,t}$ denotes the largest settling time of the responses of the power P_i of each turbine $i \in G$ to the change of their own control variables a_i . The scalar $T_{s,d} \in \mathbb{R}$ is an upper bound for the time interval that it takes for each control variable $\{a_i | i \in G\}$ to have its full effect on the power production of the neighbouring downstream turbine, $P_{d(i)}$. The interval $T_{s,d}$ includes the maximum wake travelling time between a turbine $i \in G$ and its downstream neighbouring turbine $d(i)$. Therefore, the interval $T_{s,d}$ can be assumed to be larger than $T_{s,t}$. Notice that if the control variables $\{a_i, i \in G\}$ are updated simultaneously to an iteration $a_i(k)$, at a time instant denoted by t_k^{upd} , the gradient updates can be scheduled according to the following update rules:

$$\frac{\partial \tilde{P}_i}{\partial a_i}(k) = \frac{\tilde{P}_i(t_k^{\text{upd}} + T_{s,t}) - \tilde{P}_i(t_k^{\text{upd}})}{a_i(k) - a_i(k-1)}, \quad (3.18)$$

$$\frac{\partial \tilde{P}_{d(i)}}{\partial a_i}(k) = \frac{\tilde{P}_{d(i)}(t_k^{\text{upd}} + T_{s,d}) - \tilde{P}_{d(i)}(t_k^{\text{upd}} + T_{s,t})}{a_i(k) - a_i(k-1)}, \quad (3.19)$$

for all $i \in G$. In Algorithm 1 the complete GA-MPPT wind plant control scheme is given. In this algorithm, Δt denotes the interval between two samples. In line 9 of Algorithm 1, the value of Δt is employed to update a scalar time counter $\tau \in \mathbb{R}$ that schedules the updates of the gradients. The size of the initial step on the control settings $\{a_i | i \in G\}$ is $-(\Delta a)_{\text{init}}$, where $(\Delta a)_{\text{init}} > 0$ is a scalar that is to be chosen beforehand. A Boolean variable *LocGr* is used to memorize whether or not the gradients $\frac{\partial \tilde{P}_i}{\partial a_i}$ have been updated after the last control update. Estimates of the speed of the incoming wind field V_∞ are to be produced using measuring (Mikkelsen et al., 2013) or filtering techniques (Østergaard et al., 2007; Knudsen et al., 2011), and these estimates of V_∞ are to be stored for a certain time window to be able to calculate the lagged variables $V_{\infty,i}^{\text{del}}(t)$ in lines 4, 12, and 20 of Algorithm 1. In a straightforward manner, Algorithm 1 can be adapted to give the QN-MPPT method.

Algorithm 1 The pseudocode below shows the Gradient-Ascent MPPT wind plant control algorithm. The variables \bar{a}_i and \bar{P}_i are used to store past values of the control variables and the power of the turbines.

```

1: given  $T_{s,d}, T_{s,t}, (\Delta a)_{\text{init}}, K$  and sets  $F, G$ 
2:  $\tau \leftarrow 0, \text{LocGr} \leftarrow \text{False}$ 
3:  $a_i \leftarrow a_i^{\text{opt}} \forall i \in F$ 
4: measure  $P_i(t)$ , estimate  $V_{\infty,i}^{\text{del}}(t) \forall i \in G$ 
5:  $\bar{P}_i \leftarrow P_i(t) V_{\infty,i}^{\text{del}}(t)^{-3} \forall i \in G$ 
6:  $\bar{a}_i \leftarrow a_i \forall i \in G$ 
7:  $a_i \leftarrow \bar{a}_i - (\Delta a)_{\text{init}} \forall i \in G$ 
8: loop
9:    $\tau \leftarrow \tau + \Delta t$ 
10:  if  $\tau > T_{s,t}$  and  $\tau \leq T_{s,d}$  and  $\text{LocGr} = \text{False}$  then
11:    for all  $i \in G$  do
12:      measure  $P_i(t)$ , estimate  $V_{\infty,i}^{\text{del}}(t), \tilde{P}_i \leftarrow P_i(t) V_{\infty,i}^{\text{del}}(t)^{-3}$ 
13:       $\frac{\partial \tilde{P}_i}{\partial a_i} \leftarrow \frac{\tilde{P}_i - \bar{P}_i}{a_i - \bar{a}_i}$ 
14:       $\bar{P}_i \leftarrow \tilde{P}_i$ 
15:       $\bar{P}_{d(i)} \leftarrow \tilde{P}_{d(i)}$ 
16:    end for
17:     $\text{LocGr} \leftarrow \text{True}$ 
18:  else if  $\tau > T_{s,d}$  then
19:    for all  $i \in G$  do
20:      measure  $P_{d(i)}(t)$ , estimate  $V_{\infty,d(i)}^{\text{del}}(t), \tilde{P}_{d(i)} \leftarrow P_{d(i)}(t) V_{\infty,d(i)}^{\text{del}}(t)^{-3}$ 
21:       $\frac{\partial \tilde{P}_{d(i)}}{\partial a_i} \leftarrow \frac{\tilde{P}_{d(i)} - \bar{P}_{d(i)}}{a_i - \bar{a}_i}$ 
22:       $\bar{P}_i \leftarrow \tilde{P}_i$ 
23:       $\bar{a}_i \leftarrow a_i$ 
24:       $a_i \leftarrow \bar{a}_i + K \left[ \frac{\partial \tilde{P}_i}{\partial a_i} + \frac{\partial \tilde{P}_{d(i)}}{\partial a_i} \right]$ 
25:    end for
26:     $\tau \leftarrow 0$ 
27:     $\text{LocGr} \leftarrow \text{False}$ 
28:  end if
29: end loop

```

3.3. BENCHMARK WIND PLANT CONTROL ALGORITHM WITH A GAME THEORETIC APPROACH

In the simulation examples of Section 3.5, the MPPT approaches are compared with a Game Theoretic (GT) wind plant control approach with full communication between the turbines, presented in Marden et al. (2013). Like the MPPT method, the GT approach of Marden et al. (2013) is data-driven, since it only needs measurements of the power and the control parameters to track the point of maximum power. A similar game-theoretic approach was taken in Park et al. (2013), but the latter uses knowledge of the model to efficiently perform the optimization. An important difference between the MPPT method and the GT optimization approach with full communication as presented in Marden et al. (2013), is that this GT approach aims to optimize the settings of each turbine by evaluating their effect on all the turbines in the wind plant.² Furthermore, the GT approach performs this optimization by making random perturbations to the control variables and holding the settings if they yield an improvement of the wind plant total power production, so as to iteratively find the global maximum of the wind plant total power. To evaluate the effect of each control variable change on the total power production of the wind plant, the algorithm has to wait until the wake has travelled through the entire wind plant. This waiting time is denoted by $T_{s,p}$.

In Algorithm 2, the control scheme of the GT approach is given as it is implemented in our simulations. The algorithm has two parameters that are used to set the exploration rate of the randomized optimization:

- a scalar $E \in [0, 1]$ that defines the probability of using a new random setting for a_i , instead of keeping the settings that yielded the largest total power so far,
- a scalar $K \in [0, 1]$ that defines the size of the interval in which the random steps on the control settings are chosen.

The range $[a^{\min}, a^{\max}]$ is the set of allowable values for the control settings, which for the axial induction factor is given by $[0, 1/3]$. The algorithm is somewhat different than the one presented in Marden et al. (2013), since in the exploration it makes small random perturbations in each iteration, rather than taking random values in the full range $[a^{\min}, a^{\max}]$. This change is made to improve the convergence speed of the algorithm, and reduce oscillations of the power signal.

3.4. WIND PLANT SIMULATION MODEL

In Section 3.5, the control methods are evaluated in simulations of a wind plant. The simulation model used in this case study is the Jensen model, that was first introduced in Katić et al. (1986), extended with a delay model to include the wake travelling dynamics. The Jensen model is a relatively simple engineering model that gives an estimate of

²Marden et al. (2013) also presents a form of the Game-Theoretic method with limited communication between the turbines, but it is shown in the same paper that this particular method has slower convergence than the full communication GT approach.

Algorithm 2 The pseudocode below shows a wind plant control algorithm similar to the Game Theoretic approach of Marden et al. (2013). The values of \mathcal{R}_1 and \mathcal{R}_2 are drawn randomly using a uniform distribution. The variables \bar{a}_i and \bar{P}_i are used to store past values of the control variables and the power of the turbines.

```

1: given  $T_{s,p}, K \in [0, 1], E \in [0, 1]$  and set  $F$ 
2:  $\tau \leftarrow 0$ 
3:  $a_i \leftarrow a_i^{\text{opt}} \forall i \in F$ 
4:  $\bar{P} \leftarrow \sum_{i=1}^N P_i(t)$ 
5:  $\bar{a}_i \leftarrow a_i$ 
6: loop
7:    $\tau \leftarrow \tau + \Delta t$ 
8:   if  $\tau > T_{s,p}$  then
9:     if  $\sum_{i=1}^N P_i(t) > \bar{P}$  then
10:       $\bar{a}_i \leftarrow a_i \forall i \in F$ 
11:       $\bar{P} \leftarrow \sum_{i=1}^N P_i(t)$ 
12:     end if
13:     for all  $i \in F$  do
14:        $\mathcal{R}_1 \leftarrow \text{random value} \in [0, 1]$ 
15:       if  $\mathcal{R}_1 < E$  then
16:          $\mathcal{R}_2 \leftarrow \text{random value} \in [-a^{\text{max}}, a^{\text{max}}]$ 
17:          $a_i \leftarrow \min(\max(\bar{a}_i + K\mathcal{R}_2, a^{\text{min}}), a^{\text{max}})$ 
18:       else
19:          $a_i \leftarrow \bar{a}_i$ 
20:       end if
21:     end for
22:      $\tau \leftarrow 0$ 
23:   end if
24: end loop

```

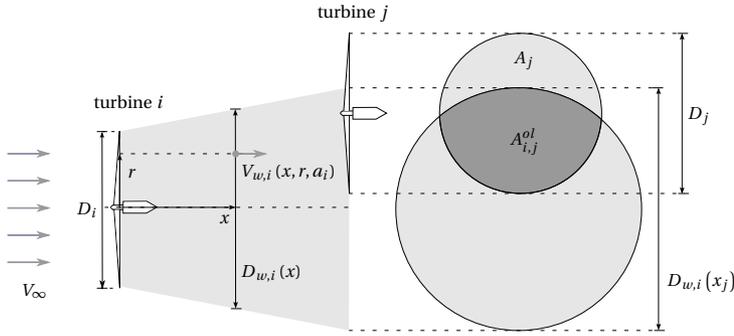


Figure 3.4: The wake expansion parameters in the Jensen model. The value of $A_{i,j}^{ol}$ is set to a maximum value of 1 when the turbines are in each full wake, i.e., when A_i and A_j fully overlap.

the velocity profile in the wind plant as a function of the incoming wind field and the set of axial induction factors of each turbine $\{a_i | i \in F\}$.

Consider a single turbine i with a rotor diameter D_i , with its rotor axis aligned with the wind direction. Assume an incoming uniform wind field with a free-stream speed V_{∞} . Let (x, r) be a point in the wake of the turbine, where x is the distance to the rotor disk plane of the turbine, and r is the distance to the centerline of the wind turbine rotor axis (see Figure 3.4). The Jensen model estimates the wind speed in the point (x, r) to be:

$$V_{w,i}(x, r, a_i) = V_{\infty} [1 - \delta V_{w,i}(x, r, a_i)], \quad (3.20)$$

with the fractional velocity deficit $\delta V_{w,i}(x, r, a_i)$ given by:

$$\delta V_{w,i}(x, r, a_i) = \begin{cases} 2a_i \left[\frac{D_i}{D_{w,i}(x)} \right]^2 & \text{for } r \leq \frac{D_{w,i}(x)}{2}, \\ 0 & \text{for } r > \frac{D_{w,i}(x)}{2}, \end{cases} \quad (3.21)$$

where $D_{w,i}$ is the diameter of the wake, which is assumed to have a circular cross-section. The diameter is assumed to expand proportional to the distance x :

$$D_{w,i}(x) = D_i + 2\kappa x, \quad (3.22)$$

where parameter κ represents a tunable wake expansion coefficient. In the simulation examples of Section 3.5, this parameter is set to $\kappa = 0.084$, to fit the offshore wind plant power data provided in Barthelmie et al. (2009).

The model is extended to include multiple turbines with interacting wakes. Then the effective wind speed V_j for a turbine $j \in F$ is found by combining the estimated wind velocity deficits created by each upstream turbine:

$$V_j = V_{\infty} [1 - \delta V_j], \quad (3.23)$$

with:

$$\delta V_j = 2 \sqrt{\sum_{i \in F: x_i < x_j} \left[a_i \left[\frac{D_i}{D_{w,i}(x_j - x_i)} \right]^2 \frac{A_{i,j}^{ol}}{A_j} \right]^2}, \quad (3.24)$$

where A_j is the rotor swept area of a turbine j , and $A_{i,j}^{ol}$ is the overlapping area of the rotor swept disk of a turbine j , and the wake generated by an upstream turbine i at the rotor plane of turbine j (see Figure 3.4), which are calculated using basic trigonometric relationships.

When the effective wind speed at each turbine is known, the power of each turbine is calculated as Bianchi et al. (2007):

$$P_j = \frac{1}{2} \rho A_j C_P(a_j) V_j^3, \quad (3.25)$$

where ρ is the air density and C_P is the power efficiency coefficient, which is expressed as a function of the axial induction factor:

$$C_P(a_j) = 4a_j [1 - a_j]^2. \quad (3.26)$$

In the above form, the Jensen model is a static model, in which a change in the axial induction factor has an immediate effect on the total power. To be able to evaluate the time-efficiency of the different wind plant control algorithms described in this chapter, simplified wake travelling dynamics are added to the model. This is done by including estimated delays corresponding to the wake travelling from one turbine to the next in the Jensen model, following an approach similar to that presented in González-Longatt et al. (2012). In this approach, an estimate of the wake travel time between a turbine i and its nearest downstream neighbour $d(i)$, denoted by $T_{i \rightarrow d(i)}$, is made by assuming a constant speed in between the turbines that is equal to the average of the wind speed just behind the turbine i and the wind speed just in front of the downstream turbine $d(i)$:

$$T_{i \rightarrow d(i)} = \frac{x_{d(i)} - x_i}{\frac{1}{2} [V_i [1 - 2a_i] + V_{d(i)}]}. \quad (3.27)$$

Before a change in the incoming wind field at location x_{V_∞} has its effect on a turbine i in the plant, the wind has to travel from location x_{V_∞} to the turbine location x_i , see also Figure 3.3. To incorporate this effect, the delays approximated with equation (3.27) are incorporated in the model in such a way that a change in the incoming wind field has a delayed effect on the turbines. In this model, a uniform incoming free stream wind field with a speed $V_\infty(t)$ is prescribed as the wind speed just in front of a wind turbine f , i.e. $x_{V_\infty} = x_f$, where f is the turbine that is standing upstream of all other turbines, i.e.:

$$f = \underset{i \in F}{\operatorname{arg\,min}}(x_i). \quad (3.28)$$

For example, for the wind plant configuration of Figure 3.3, $f = 1$. The notation $\tilde{R}(j)$ is used for the complete row of turbines $R(i)$ that a certain turbine j is part of, i.e. $\tilde{R}(j)$ is the largest set $R(i)$ for which it holds that $j \in R(i)$. Further, let $\tilde{i}(j)$ be the first member of the set $\tilde{R}(j)$ (see Figure 3.3 for an example of the mapping $j \mapsto \tilde{i}(j)$), then by summing the different delays an expression for $T_{V_\infty \rightarrow j}$, being the total wind travelling delay for a turbine j , is found:

$$T_{V_\infty \rightarrow j} = \frac{x_{\tilde{i}(j)} - x_f}{V_\infty (t - T_{\tilde{i}(j) \rightarrow j})} + T_{\tilde{i}(j) \rightarrow j}, \quad (3.29)$$

with:

$$T_{\tilde{i}(j) \rightarrow j} = \sum_{u \in \tilde{R}_j: x_u < x_j} T_{u \rightarrow d(u)}. \quad (3.30)$$

Substituting expression (3.29) in equation (3.6) yields:

$$V_{\infty, j}^{del}(t) = V_{\infty} \left(t - \frac{x_{\tilde{i}(j)} - x_f}{V_{\infty} (t - T_{\tilde{i}(j) \rightarrow j})} - T_{\tilde{i}(j) \rightarrow j} \right). \quad (3.31)$$

Moreover, for a change in control variable a_i to have an effect on the downstream turbine $d(i)$, the wake has to travel from turbine i to turbine $d(i)$. To incorporate this effect, a delay structure is added to the model for the wake velocity deficit, which for the MPPT method is given by:

$$\delta V_j^{del}(k) = 2 \sqrt{\sum_{i \in F: x_i < x_j} \left(\delta V_{w, i, j}^{del}(k) \frac{A_{i, j}^{pl}}{A_i} \right)^2}, \quad (3.32)$$

with:

$$\delta V_{w, i, j}^{del}(k) = a_i (k - \Delta_{i, j}) \left[\frac{D_i}{D_{w, i} (x_j - x_i)} \right]^2, \quad (3.33)$$

where $\Delta_{i, j}$ is the discrete delay as a consequence of the wake travelling from a turbine i to a turbine j . For example, $\Delta_{i, j} = 1$ if $j = d(i)$, $\Delta_{i, j} = 2$ if $j = d(d(i))$, and so on. Then at time $t_{upd}^k + T_{s, t}$, the power of each turbine changes in response to the change in its own control variable. Hence, the power estimate is updated by:

$$V_j(t_{upd}^k + T_{s, t}) = V_{\infty, j}^{del}(t_{upd}^k + T_{s, t}) \left[1 - \delta V_j^{del}(k) \right], \quad (3.34)$$

$$P_j(t_{upd}^k + T_{s, t}) = \frac{1}{2} \rho A_j C_P (a_j(k)) \left[V_j(t_{upd}^k + T_{s, t}) \right]^3. \quad (3.35)$$

The settling time of the turbines with respect to a change in their own control variables is assumed to be $T_{s, t} = 5s$.

In the MPPT method, the gradients are updated simultaneously after each wake has reached the next downstream turbine, which results in the following settling time $T_{s, d}$ used in scheduling the gradient updates:

$$T_{s, d} \approx \max_{i \in G} (T_{i \rightarrow d(i)}). \quad (3.36)$$

At time $t_{upd}^k + T_{s, d}$ the wakes have travelled from one turbine to the next, and the velocity deficit in front of each turbine changes. Therefore, the wind velocities and powers are updated by:

$$V_j(t_{upd}^k + T_{s, d}) = V_{\infty, j}^{del}(t_{upd}^k + T_{s, d}) \left[1 - \delta V_j^{del}(k+1) \right], \quad (3.37)$$

$$P_j(t_{upd}^k + T_{s, d}) = \frac{1}{2} \rho A_j C_P (a_j(k)) \left[V_j(t_{upd}^k + T_{s, d}) \right]^3. \quad (3.38)$$

Then in the MPPT method, after the wake has travelled to the next turbine, the control variables are updated, hence:

$$t_{upd}^{k+1} = t_{upd}^k + T_{s,d}. \quad (3.39)$$

The wake travelling modelling results in the fact that after each control update in the MPPT optimization, initially the total power will decrease as a consequence of the fact that some turbines will decrease their own power extraction, but when the wakes of those turbines travel to the next row of turbines, the total power will increase as a consequence of the reduced velocity deficits in the wakes.

In the model used to evaluate the GT method, a similar behaviour is incorporated, but part of the delay structure is omitted because in the GT method the wake will have travelled through the full wind plant before a control update takes place. Under the same assumptions as used above to estimate $T_{s,d}$, an estimate of $T_{s,p}$ is obtained, which denotes the largest time it takes for a change in a control variable a_i of a turbine $i \in G$ to have its effect on the power of all of its downstream turbines, by summing each of the turbine-to-turbine wake travel times. This time interval $T_{s,p}$ is used to schedule the control updates in the GT method. To find an expression for $T_{s,p}$, the notation $R(i)$ is used for a set that includes the index of a turbine $i \in G$ and the indices of the full row of turbines in the set G that are affected by that turbine i , i.e., $R(i) = \{i, d(i), d(d(i)), \dots\}$ (see Figure 3.3 for an example of the mapping $i \mapsto R(i)$). Using this notation, the approximation is given by:

$$T_{s,p} \approx \max_{i \in G} \left(\sum_{j \in R(i)} T_{j \rightarrow d(j)} \right). \quad (3.40)$$

The estimates for the delays in the model are fairly rough, but since the underlying assumptions are similar for the estimation of $T_{s,d}$ and $T_{s,p}$, these estimates can be used to make a relative comparison of the time-efficiency of each optimization method.

3.5. SIMULATION EXAMPLES

This section presents the results of simulation experiments that compare the performance of the different wind plant control methods presented in this chapter in terms of the time-efficiency of the power optimization, the power production increases that can be achieved, and the adaptability of each method to varying wind conditions.

In the simulation examples, the MPPT approaches presented in 3.2 are compared to the Game-Theoretic (GT) approach described in Section 3.3, which is a global optimization approach for maximizing the total power production of the wind plant. It is shown that as a consequence of the simplifying assumptions that are taken in the MPPT optimization scheme (in which the effect on nearest neighbouring turbines is taken into account only), the power production increases that can be obtained with the MPPT approaches may be somewhat smaller than those that can be obtained with the GT approach, although the differences are small. The distributed gradient-based optimization approach of the MPPT wind plant control method yields a large improvement of the time-efficiency however, which is important for making real-time implementation possible.

In each of the simulation experiments, the wind plant model as presented in Section 3.4 is used to simulate the Princess Amalia Wind Park, an offshore wind plant that is located 23 km off the coast of The Netherlands, under different wind conditions. The Princess Amalia Wind Park consists of 60 wind turbines with a rotor diameter of 80 m and a rated power of 2 MW. The locations of the turbines in the The Princess Amalia Wind Park are shown in Figure 3.5, together with a compass rose that defines the different wind directions as mentioned further on in the simulation examples. In each of the cases, a constant air density $\rho = 1.225 \text{ kg} \cdot \text{m}^{-3}$ is assumed.

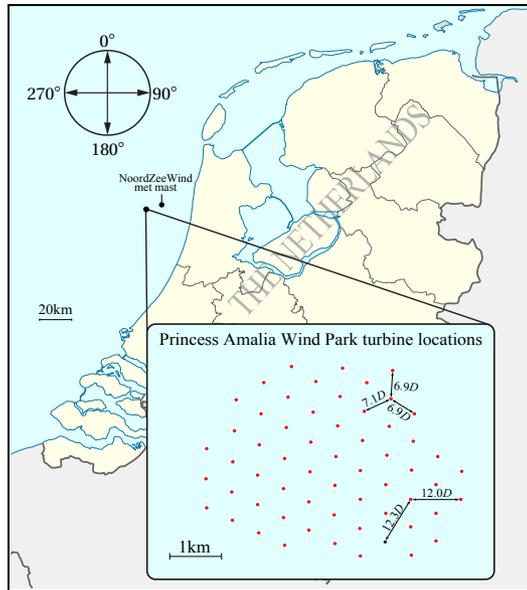
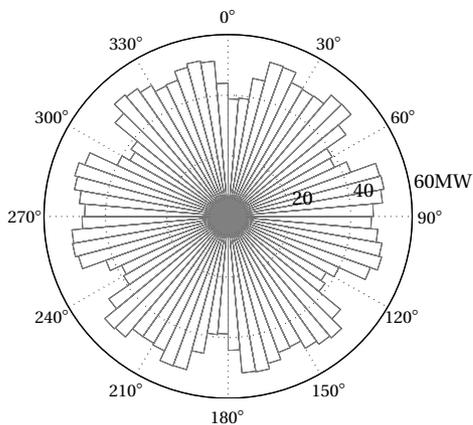
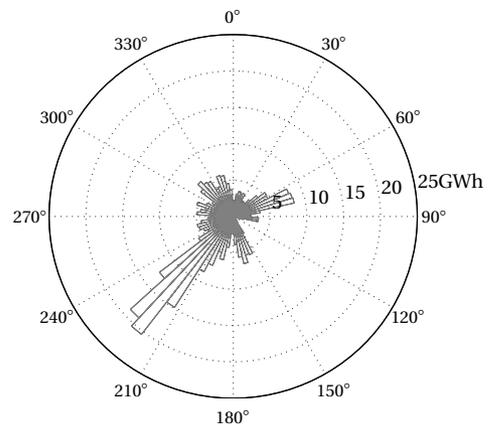
In Section 3.5.1, the time-efficiency and gain of the MPPT power optimization methods is compared to that of the GT approach in a simulation of the wind plant with a constant incoming wind speed and a single wind direction. Thereby each of the parameters of the optimization algorithms are set to deliver fast convergence properties for this specific case. In Section 3.5.2, it is evaluated whether the same settings will also result in good convergence properties for other wind directions, by repeating the simulation experiment of Section 3.5.1 for a range of wind directions. These results are then used in Section 3.5.3 to make an estimate of the energy production increase that can be achieved using these methods. In the final experiment in Section 3.5.4, it is shown that the proposed MPPT algorithms are able to deal with a time-varying incoming wind speed.

3.5.1. COMPARATIVE SIMULATION STUDY OF THE WIND PLANT CONTROL APPROACHES FOR A CONSTANT INCOMING WIND SPEED

3

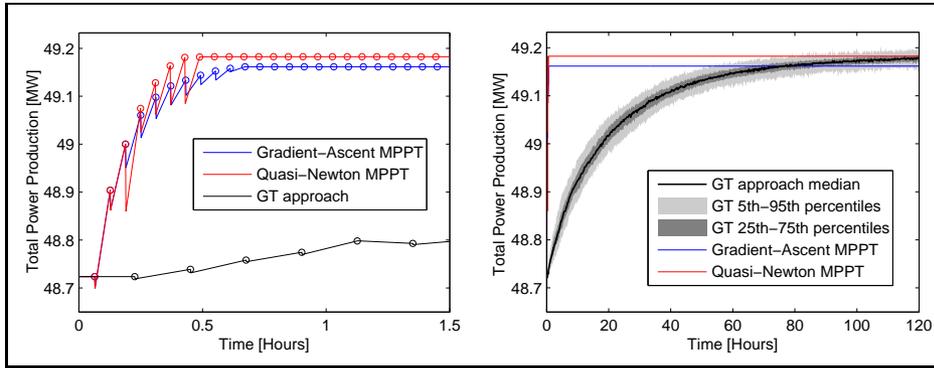
In the first simulation example, a uniform wind field with a direction of 25° is assumed, with a constant, below-rated speed of $V_\infty = 8 \text{ ms}^{-1}$. The parameters determining the iteration step-size of each of the methods are tuned to yield fast convergence towards the optimum. With the turbine powers P_i expressed in the megawatt unit, this resulted in setting $K = 7.7 \text{ MW}^{-1}$ for the GA-MPPT approach, $K = 5.12 \text{ MW}^{-1}$ for the QN-MPPT approach, and $K = 0.06$ and $E = 0.1$ for the GT approach. In both MPPT methods, the size of the initial step on the control settings is set to $(\Delta a)_{\text{init}} = 0.01$.

The simulation results are given in Figure 3.6a. After the control updates take place, first a drop in total power production is observed, since the production of the upstream turbines decreases, and then after the wake has travelled to the downstream turbines, an increase of the total power is achieved. The results show that each of the control methods will iteratively improve the total electrical power production. The GA-MPPT approaches converge to a slightly lower total power than the GT and the QN-MPPT approach. This is because the GT is guaranteed to converge to a global optimum of the wind plant total power, and the QN-MPPT approach also finds this optimum in this case, but the GA-MPPT converges to a local optimum that is close to this global optimum. The MPPT approaches increase power much quicker than the GT approach. This is because the MPPT methods use gradient information to converge to the local optimum in a faster way. Also, it is because the distributed optimization approach of the MPPT methods consider the wake effect on the nearest neighbouring turbines only, which makes that the MPPT algorithms is able to update the control settings more frequently, since the turbine-to-turbine settling time $T_{s,d}$ is much shorter than the total wind plant settling time $T_{s,p}$.

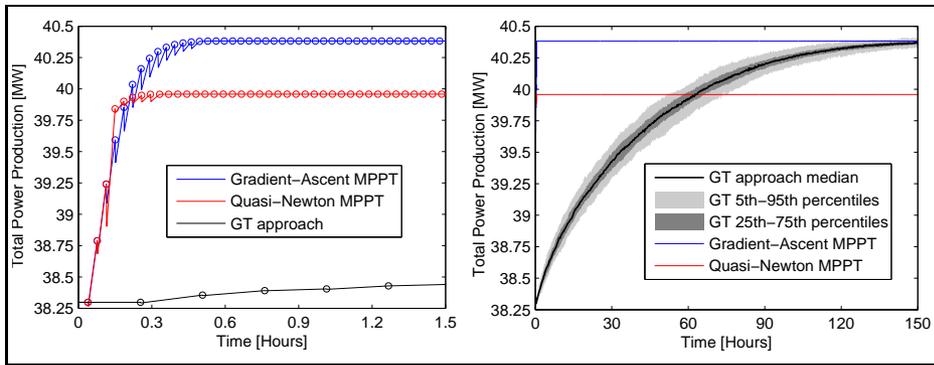
(a) Locations of the wind turbines (with rotor diameter $D = 80\text{m}$)(b) Wind plant electrical power production for different wind directions, with $V_\infty = 8\text{ms}^{-1}$ without optimization, as predicted by the model presented in Section 3.4.

(c) Total annual wind plant electrical energy produced in below-rated wind conditions for different wind directions, without optimization, as predicted by the model presented in Section 3.4, using the wind data of the NoordZeeWind met mast (NoordzeeWind B.V., 2013).

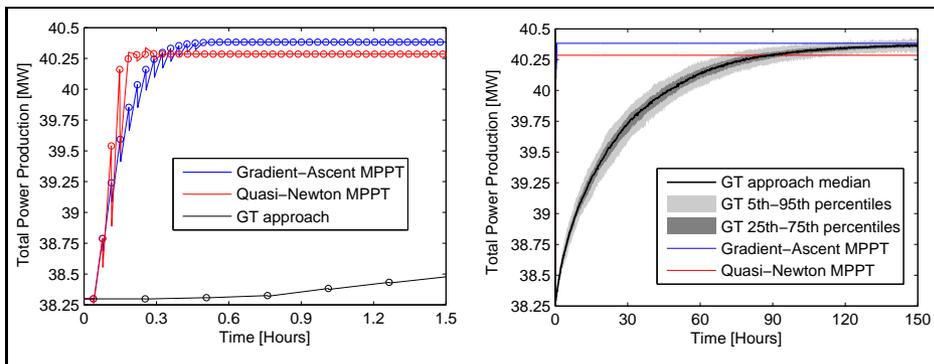
Figure 3.5: Properties of the Princess Amalia Wind Park offshore wind plant.



(a) Results for wind direction 25° .



(b) Results for wind direction 60° , parameters of MPPT and GT not returned (i.e. kept the same as for 25° wind direction).



(c) Results for wind direction 60° , with parameters of QN-MPPT and GT returned.

Figure 3.6: Results of power optimization control with the GA-MPPT, QN-MPPT, and the GT approach in the wind plant simulation described in Section 3.5.1, where the incoming wind speed is kept constant at $V_\infty = 8\text{ms}^{-1}$. On the left, the 'o'-markers on the total power curves correspond to the time instances at which the control updates take place. On the right, the results are shown on a larger time range to show the convergence of the GT approach, and to show the effect of the randomization in this GT method, the distribution of the results of 100 experiments is shown.

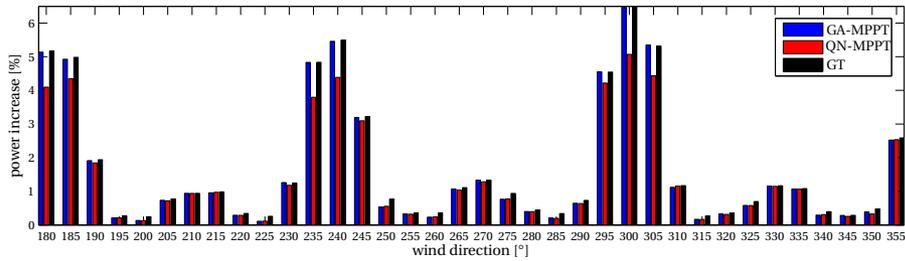
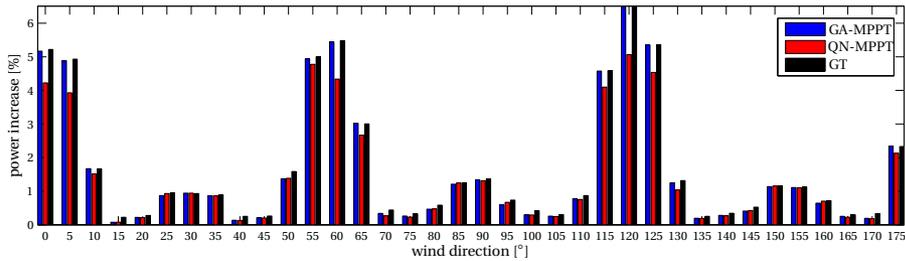
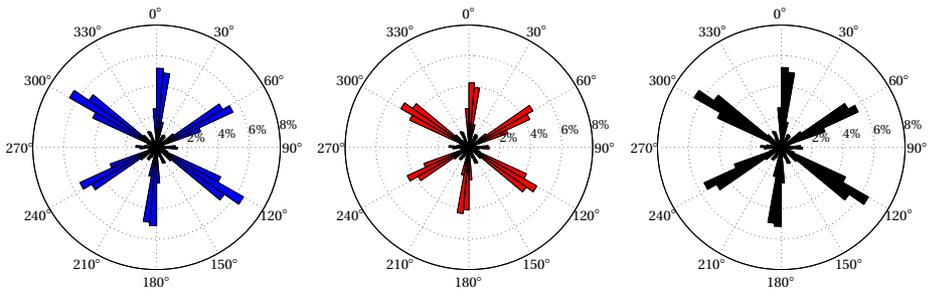
3.5.2. EVALUATION OF THE POWER GAIN OF THE WIND PLANT CONTROL APPROACHES FOR DIFFERENT WIND DIRECTIONS

In the second case study, it is evaluated how the power increase that is achieved with the MPPT and GT control methods changes with the direction of the incoming wind. The simulation experiment of Section 3.5.1 is repeated for the $0^\circ, 5^\circ, 10^\circ, \dots, 355^\circ$ wind directions. In each simulation the incoming wind speed is kept constant at $V_\infty = 8 \text{ ms}^{-1}$. The power increase that is achieved after each of the control methods have converged is shown in Figure 3.7. It can be seen that the power increase that is achieved is highly dependent on the wind direction, as the spatial configuration of the turbines in the Princess Amalia Wind Park is optimized for more frequently occurring wind directions. As in the previous case study, the MPPT methods converge much faster than the GT method. This is shown in Figure 3.7e, in which the convergence times of the methods are given for each wind direction.

The optimization results for different wind directions are obtained without adjusting the parameter K for each direction (the same settings are used as in the example of Section 3.5.1). In Figure 3.7d it can be seen that while the QN-MPPT may yield a slightly higher power increase for the 25° wind direction (for which the control parameters are tuned), for other wind directions the GA-MPPT yields a higher power increase. This is because the convergence properties of the QN-MPPT approach are sensitive to the tuning of the K parameter, and the QN-MPPT approach may need retuning of the step-size scaling parameter K to different wind directions to have good convergence properties, while on the other hand, the GA-MPPT is more easy to use, in the sense that it does not need adjustment of the K parameter to have good convergence properties for different wind directions. This is also illustrated in Figure 3.6b and 3.6c. In Figure 3.6b the power time-series for each of the optimization methods is shown for the wind direction 60° , where for each method the same tuning is used as previously for the 25° wind direction. It can be seen that this results in bad convergence properties for the QN-MPPT method, as it does not come as close to the global optimum found by the GT method as the GA-MPPT method does. When the QN-MPPT method is retuned to $K = 0.044$, better convergence properties can be obtained, the results are shown in Figure 3.6c. Also, in Figure 3.6c, the GT method is retuned such that faster convergence occurs, by setting $E = 0.2$ and keeping $K = 0.06$.

3.5.3. ESTIMATION OF THE ANNUAL ENERGY GAIN OF THE WIND PLANT CONTROL METHODS

In the wind plant model of Section 3.4, the fractional velocity deficit in the wake (i.e. the relative amount of wake recovery) is independent of the incoming wind speed. Therefore, the settings finally found by the GT and MPPT optimization algorithms yield the same power increase for different below-rated wind speeds. This can also be confirmed by rerunning the experiments of Section 3.5.2 with different below-rated incoming wind speeds (the results are omitted for brevity). By using wind measurements at a nearby location, an estimate is made of the increase of the energy annually produced in below-rated conditions that can be achieved using the different optimization methods. These wind measurements were made by the NoordzeeWind meteorological mast at a nearby



(d) Comparison of power production increase for GA-MPPT, QN-MPPT and GT control for different wind directions.

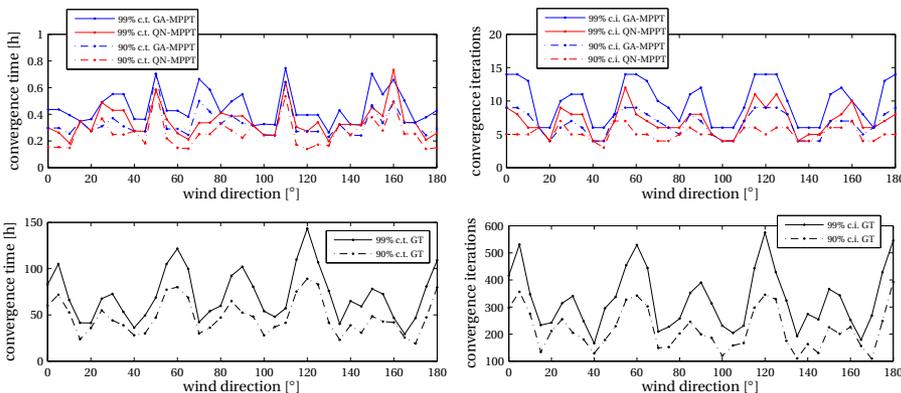


Figure 3.7: Results of the power optimization the Princess Amalia Wind Park for different wind directions.

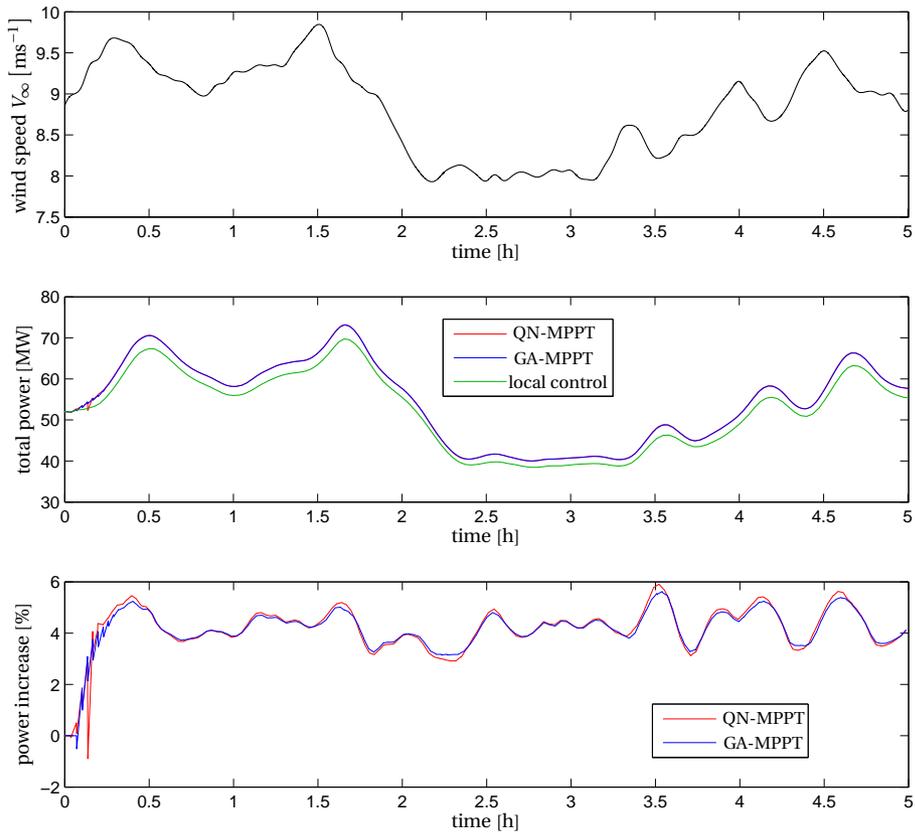
location in the North Sea during the period from July 1st, 2005 to June 30th, 2006 (Brand et al., 2013). The measurements are available at NoordzeeWind B.V. (2013). The measurements consist of 10 minute averages of the wind direction and the free stream wind speeds. With the wind plant model of Section 3.4, the energy production of the Princess Amalia Wind Park is calculated for each of the wind directions, for the case in which each turbine is controlled individually, and the results are shown in Figure 3.5c. When the final power increase that is obtained with each of the control methods is added to these productions, and the results are summed over the year, a rough estimate can be made of the annual production with each of the methods. In this way it is estimated that on a yearly basis, the energy produced in below-rated wind conditions can be increased with 1.36% using the GA-MPPT method, with 1.19% using the QN-MPPT method, and with 1.42% using the GT method. Notice that in these calculations, the convergence time of each of the methods is not taken into account, and that it is thus assumed that each of the control methods is able to quickly track the changing wind conditions. Given the large convergence times of the GT method when compared to the MPPT method, it is less likely that the GT method is able to perform this tracking, and thus it is less likely that the estimated energy production increase can be achieved in practice using the GT method as described in Section 3.3 in an online implementation.

3.5.4. SIMULATION OF THE MPPT APPROACHES WITH A VARYING WIND SPEED

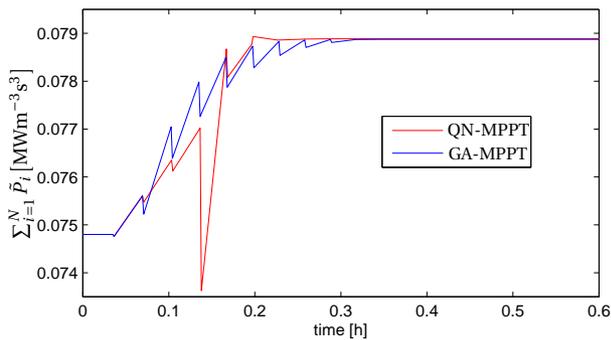
In the final case study, it is shown in simulation that indeed the total power can be optimized under varying wind speeds using the MPPT methods. The results, shown in Figure 3.8, are obtained by simulating the Princess Amalia Wind Park with a varying incoming wind speed signal, obtained from a part of the December 2010 wind speed measurements of the NoordzeeWind meteorological mast in the North Sea (NoordzeeWind B.V., 2013), shown in the top plot of Figure 3.8a. It should be noted that this wind speed is smoothed, as it is an interpolation of 10 minute averages of the measured wind speed. In the lower two plots in Figure 3.8a, it can be seen that a power increase of about 4% can be obtained with the MPPT techniques, using $K = 0.25$ for the QN-MPPT approach, and $K = 13$ for the GA-MPPT approach. In both MPPT approaches, the initial step size is set to $(\Delta a)_{\text{init}} = 0.01$. Notice that in this time-varying wind speed case, the algorithm is able to continue the optimization under changing wind velocity, making use of the fact that the objective function is defined as the sum of the efficiencies \tilde{P}_i as defined in (3.6), rather than the sum of turbine power productions. In Figure 3.8b it is shown how this objective function is optimized within the first 20 minutes of the simulations. After this convergence period, the algorithm does not make any adjustments of control settings anymore, and the wind plant optimally captures power as long as the wind direction does not change.

3.6. CONCLUSIONS

In this chapter, two data-driven MPPT control algorithms for wind plants are presented that optimize the control settings of each turbine in the plant in a real-time closed-



(a) Results of wind plant simulation experiments with a varying incoming wind speed. The varying incoming wind speed V_∞ is shown in the top plot. In the middle plot, it is shown how the different MPPT control methods yield an increase of the produced power when compared to the case where locally optimal control is used. The relative increase in power for each time instant is shown in the bottom plot.



(b) The optimization of the objective function using the GA-MPPT and QN-MPPT methods as it takes place within the first 20 minutes of the wind plant simulation experiments with varying wind speeds as described in Section 3.5.4.

Figure 3.8: Varying wind speed simulation results

loop manner. The control algorithms achieve a power production increase of the wind plant by taking into account the interaction between the turbines through the wake effect. A speed-up of the optimization is achieved by using gradient-based optimization techniques with a distributed approach in which we take into account the effect on neighbouring turbines only. Using information on the spatial configuration of the wind plant in this way, results in a much faster convergence of the power optimization than is achieved with the existing game-theoretic method with full communication between the turbines presented in Marden et al. (2013). This is demonstrated in the first simulation example in Section 3.5.1.

As the gradient information used in the GA-MPPT and QN-MPPT is calculated from measured data, the method is adaptive to changing wind conditions such as a changing wind direction. For the optimization algorithms to be able to track time-varying wind conditions it is needed that the optimization takes place in a time-efficient manner. Therefore, the gradient-based distributed MPPT approaches may be a more likely candidate for practical application than the GT approach.

It is shown in the second simulation example in Section 3.5.2, that the GA-MPPT method is more robust to changing wind conditions than the QN-MPPT approach, as the latter optimization method may need an adjustment of the iteration step-sizes to adapt to changing wind conditions.

The final simulation example in Section 3.5.4 showed that by letting the control algorithm optimize the energy conversion efficiency of the wind plant rather than the total power, the optimization scheme is made adaptive to time-varying incoming wind speeds. To define this efficiency, the algorithm needs an estimate of the effective incoming wind speed, and an estimate of the delays related to the wind field travelling from one turbine to the next. In the simulation experiments with the Jensen model, these wind speeds and delays are assumed to be exactly known. Future research aims at applying the MPPT wind plant methods on a more advanced simulation model that describes the wake and turbine dynamics in more detail, such as the SOWFA model presented in Churchfield et al. (2012b); Fleming et al. (2013b). To be able to do this, measuring (Mikkelsen et al., 2013) or filtering techniques (Østergaard et al., 2007; Knudsen et al., 2011) are to be applied that produce estimates of the speed of the incoming wind field, and models are to be developed that give estimates of the wake traveling delays (a more advanced delay model than used in this work, is presented in Chapter 5). Also, filtering techniques are to be applied within the optimization gradient-based frame to deal with the effect of small-scale turbulence on power production, which may form a problem in determining the gradients.

The potential to increase wind plant power production with axial-induction-based control is sensitive to the wind direction, as is shown in Figure 3.7 and was also pointed out in Knudsen et al. (2014). Therefore, evaluating the methods in realistic wind scenarios with changing wind conditions, will be informative to further evaluate the potential of the control techniques. Moreover, the (more recent) SOWFA simulation results in Chapter 2 and Annoni et al. (2014a) show that axial-induction-based control may only have a potential for particular inflow conditions and turbine power and thrust characteristics, and finding these this range of conditions, and the influence of the turbine

characteristics, is ongoing research. The results of aforementioned recommended research, can also impact the predictions on annual wind plant yield increase as a result of axial-induction-based, as presented in Section 3.5.2, and thus we have to remain critical towards these results.

Since the publication of this work, further research has been performed towards direct data-driven axial-induction-based wind plant control methodologies that take into account the wake propagation delays: Ahmad et al. (2014), that uses the wake delay model as presented in this Chapter and shows a faster convergence with a simultaneous perturbation stochastic approximation method than the MPPT methods, and the extremum-seeking control method in Yang et al. (2013b). This shows that there are alternative ways of performing the data-driven optimization, but in each case the time-efficiency has to be taken into account. In future work, the optimization methodologies can be extended to above-rated wind conditions by including constraints in the optimization problem that give upper bounds on the power production of each turbine in the wind plant. Also, the optimization methodology may be used to perform balancing of the loads on each turbine, using an approach similar to the one presented in Soleimanzadeh et al. (2012), in which the objective function is extended with static turbine load measures.

4

YAW-BASED OPTIMIZATION CONTROL FOR WIND PLANTS

This chapter presents a wind plant control strategy that optimizes the yaw settings of wind turbines for improved electrical energy production of the whole wind plant by taking into account wake effects. The optimization controller is based on a novel internal parametric model for wake effects, called the FLOW Redirection and Induction in Steady-state (FLORIS) model. The FLORIS model predicts the steady-state wake locations and the effective flow velocities at each turbine, and the resulting turbine electrical energy productions, as a function of the axial induction and the yaw angle of the different rotors. The FLORIS model has a limited number of parameters that are estimated based on turbine electrical power production data. In high-fidelity computational fluid dynamics simulations of a small wind plant, we demonstrate that the optimization control based on the FLORIS model increases the electrical energy production of the wind plant, with a reduction of loads on the turbines as an additional effect.

4.1. INTRODUCTION

Each wind turbine in a cluster of wind turbines (a wind power plant) can influence the performance of other turbines through the wake that forms downstream of its rotor. The wake is a flow structure that is characterized by a reduced wind speed because the turbine rotor extracts kinetic energy from the incoming flow, and an increased turbulence because the turbine obstructs the flow. If another turbine is standing in the path of a wake at a location where the flow has not yet fully recovered to free stream conditions, the reduced wind speed results in a lower electrical energy production of that turbine. In addition, the increased turbulence and shear in the wake may induce an increase in dynamic loads on the downstream turbine. These wake interaction effects have been stud-

This chapter will appear as a journal publication in Gebraad, Teeuwisse, van Wingerden, Fleming, Ruben, Marden, and Pao (2014d).

ied extensively; see Vermeer et al. (2003); Crespo et al. (1999); Sanderse et al. (2011) for reviews of the literature. The topology and amount of the wake interaction depends on time-varying atmospheric conditions (e.g., wind direction, wind speed, turbulence, and atmospheric stability), and on the operating point of each turbine that can be adjusted by changing their control settings (generator torque, pitch angles of the blades Ainslie (1988), or yaw angle Dahlberg et al. (2003); Wagenaar et al. (2012); Jiménez et al. (2010)).

In current industrial practice, wind turbines in wind plants are still controlled to maximize their own individual performance, ignoring the effect that the turbines have on other turbines through their wakes Johnson et al. (2009). Recently, the wake interaction effects have become a more significant field of study in the research on wind turbine control algorithms, as wind plants have grown in size, and more knowledge has become available on the loss of efficiency due to the wake interaction effect. The study in Barthelmie et al. (2009), for example, reports an energy production loss in an offshore wind plant due to the wake effects, of 12%, averaged over different wind directions.

Previous work on wind plant control has mainly focused on reducing wake interaction by adjusting the axial induction of turbines to improve the overall wind plant performance, which can be achieved by adjusting pitch and torque. This concept was first proposed in the late 1980s in Steinbuch et al. (1988). Static model-based optimization strategies, based on simplified parametric wake models are tested in Johnson et al. (2009); Bitar et al. (2013); Heer et al. (2014) (based on the Jensen wake model, Jensen (1984); Katić et al. (1986)) and in Gonzalez et al. (2013) (based on the Frandsen wake model, Frandsen et al. (2006)). In Schepers et al. (2007) a similar static optimization is performed based on a computational fluid dynamics (CFD) model. Axial-induction-based wind plant control strategies that adjust the control settings to changing inflow conditions were developed in Marden et al. (2013) and Gebraad et al. (2014b) (using model-free data-driven approaches), and in Soleimanzadeh et al. (2012, 2011, 2013) (using a model-based control approach using simplified CFD models).

The goal of the work presented here, is to optimize the yaw angles of the wind turbines for increased total electrical power production of the wind plant. By changing the yaw angle of a turbine, not only the axial induction of the rotor, but also the flow direction of the wake is changed. By controlling the deflection of the wake through yawing, the wake can be directed away from downstream turbines. This approach was shown to have great potential in CFD simulations in Jiménez et al. (2010); Fleming et al. (2014d,b). The concept was also tested in wind tunnel experiments with scaled turbines Dahlberg et al. (2003) and on a small wind plant in Wagenaar et al. (2012). In these tests it was confirmed that the wake can be redirected using yaw, but since only a limited amount of data could be gathered, no quantitative analysis could be made. Further, an interesting work in this context is Kragh et al. (2013a), where it is shown that misaligning the rotor yaw of a turbine with the wind direction can be used to reduce the loads on that turbine.

The response of the complete wind plant system with respect to control setting changes is slow, because of large delays that are associated with the flow in the wake traveling from one turbine to the next González-Longatt et al. (2012); Choi et al. (2013). This is a disadvantage for model-free global optimization approaches such as the one proposed in Marden et al. (2013): the long time it would take a global optimization to iteratively test

control settings on the real system and converge, forms a problem if the controller has to adapt to time-varying conditions such as wind direction and inflow velocity. In Gebraad et al. (2014b) the problem of having large delays in the system is addressed, and an alternative gradient-based ‘localized’ approach for model-free optimization is proposed to improve the time-efficiency of the wind plant control. The ‘localized’ optimization algorithm only takes into account the effect of control settings changes on the nearest downstream neighbouring wind turbines. This approach was used in Gebraad et al. (2014b) for pitch- and torque-based wind plant control, but it is less suited for yaw control if the goal is to not only deflect the wake away from the nearest downstream turbine, but also to avoid the wake hitting turbines farther downstream. Therefore in this work, we propose a model-based control scheme, in which an optimization algorithm can test a large number of possible control settings on the model, in order to iteratively find the optimal settings based on the model predictions, before applying them to the real system.

The supervisory wind plant control scheme proposed in this chapter increases the total electrical energy production of the wind plant by model-based optimization of the yaw control settings. An overview of the proposed control scheme is shown in Figure 4.1. An important part of the work presented in this chapter, is the development of the ‘internal model’ for the wind plant controller, that predicts the wake effects in the wind plant.¹ High-fidelity CFD-based models, based on a coupling of detailed turbine dynamics models with accurate wind flow models, such as the ones presented in Yang et al. (2013a); Larsen et al. (2007); Churchfield et al. (2012b); Schepers et al. (2007), have an important role in wind plant controls development, as they allow the algorithms to be tested in a controlled environment. However, because of their computational complexity, accurate CFD-based models are less suited to be used as internal models for real-time controllers. The simplified parametric Jensen and Frandsen models, that were used in axial-induction-based wind plant control strategies mentioned before Johnson et al. (2009); Bitar et al. (2013); Gonzalez et al. (2013), do not include the ability to predict the effect of yaw control on wake redirection. Therefore, we have developed a novel control-oriented model that is able to predict the steady-state effects of yaw control on the wakes, and the resulting effects on the turbine electrical power productions. The model has parameters that can be identified by fitting the predictions of the model to turbine power measurements, an approach that can be referred to as ‘gray-box’ system identification. Also, the model uses measurements from the wind plant to estimate relevant properties of the inflow into the wind plant. The combination of the model identification and the model-based optimization steps in the control scheme, is illustrated in the overview in Figure 4.1. The fact that we use measurements to identify the model parameters and the inflow conditions, motivates why later on we refer to the control scheme as ‘data-driven’. Furthermore, the model has a relatively simple structure allowing for quick computation, such that it is suited for real-time control based on model-based optimization of the control settings.

Because we do not have access to a real-world wind plant to perform yaw control experiments with, in this work, a high-fidelity CFD wind plant model is used to generate the data needed to develop, and identify the parameters of the simplified parametric model. Then the model is implemented in a wind plant control scheme that performs

¹A preliminary version of this work appeared in Gebraad et al. (2014c).

model-based optimization of the yaw settings of each turbine using a game-theoretic approach. Finally this model-based optimization control strategy is tested in the high-fidelity wind plant simulation, in which the effects on power production and loads are calculated. Hence, the high-fidelity simulation is used to provide a proof of concept for the data-driven optimization control scheme based on the simplified parametric model. Previous work on yaw optimization for wind plants, Park et al. (2013), did not include validation of the optimized settings using high-fidelity numerical simulations.

The remainder of this chapter is organized as follows. The simulation experiments performed in the high-fidelity CFD simulator to obtain identification data for the parametric model are described in section 4.2. The simplified parametric model is presented in section 4.3. In section 4.4, the game-theoretic approach to calculate optimal yaw control settings based on the simplified model is explained. In section 4.5, simulation studies are presented to validate the data-driven model-based optimization approach in the high-fidelity CFD wind plant simulation. Finally, in Section 4.6 the conclusions are presented.

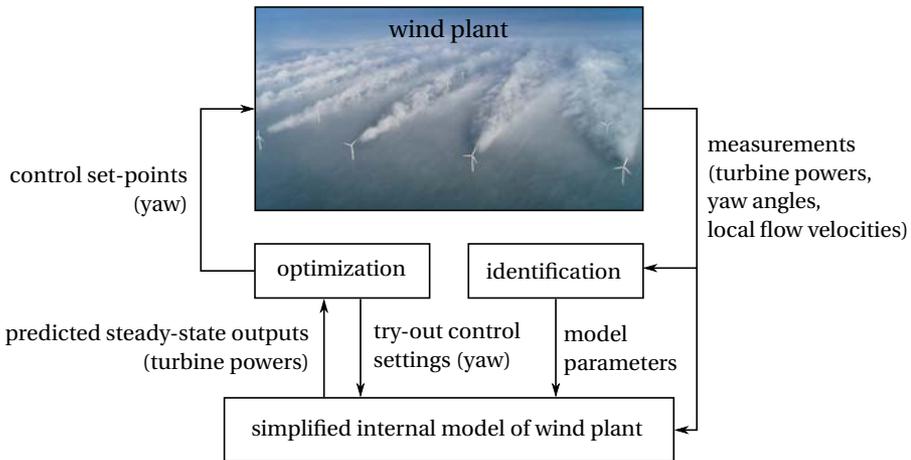


Figure 4.1: Overview of the data-driven model-based wind plant control optimization approach (wind plant photo courtesy: Vattenfall, C. Steiness).

4.2. CHARACTERIZING WAKE EFFECTS THROUGH SIMULATIONS IN SOWFA, A HIGH-FIDELITY CFD WIND PLANT SIMULATOR

In this section, we describe the simulations we performed in a CFD simulator to obtain identification data for the parametric model. We use the Simulator for Onshore/Offshore Wind Farm Applications (SOWFA), which is a Large-Eddy Simulation (LES) of the three-dimensional wind flow around one or more turbine rotors in the atmospheric boundary layer. The rotating rotor blades are modeled through an actuator line approach Sanderse et al. (2011). The actuator lines are coupled with the FAST turbine aeroelastics simulator

tool Jonkman et al. (2005) that calculates the loads, power, and rotor speed of each turbine, in addition to the forces that each turbine blade exerts on the flow. Each turbine in the simulation can be controlled using an individual control algorithm implemented in FAST, but also through a supervisory or distributed plant-wide controller. More details on the CFD calculations in SOWFA can be found in Churchfield et al. (2012b), and in Fleming et al. (2013b,a) more explanations on controls implementation in SOWFA are provided. Further, Churchfield et al. (2012a) presents a validation of SOWFA with time-averaged turbine powers measured at the Lillgrund wind plant.

In Fleming et al. (2014d,b), SOWFA simulation results were presented which show:

- how effective the yaw techniques are at wake redirection,
- the effect of yaw wake redirection techniques on the electrical energy production and loads of downstream turbines that are standing in the wake of the yawing turbine,
- the effect of repositioning a turbine such that the overlap with a wake of an upstream turbine is reduced, on electrical energy production and loads on that turbine.

More in particular, in Fleming et al. (2014b), the results of SOWFA simulations of a setup of two NREL 5-MW baseline turbines Jonkman et al. (2009). These turbines have a rotor diameter $D = 126.4$ m. In this setup the turbines are aligned in the wind direction with a downwind spacing of 7 rotor diameters ($7D$). The turbines are placed in a domain that is 3 km (horizontal length) by 3 km (horizontal width) by 1 km (height). The turbulent inflow into the domain has a mean hub-height free-stream wind speed U_∞ of 8 m/s and a turbulence intensity of 6%. This turbulent inflow is generated by a precursor simulation of the neutral boundary layer in the same domain, with an aerodynamic surface roughness that has a low value of 0.001 m, which is typical for offshore conditions.

The data of the following two series of simulations performed in Fleming et al. (2014b) are used in this chapter (see Figure 4.2a):

- in *SOWFA Simulation Series 1*, the upstream turbine (turbine 1) is yawed to redirect its wake away from the downwind turbine (turbine 2), resulting in an electrical power production decrease on turbine 1 caused by a loss of rotor efficiency, but an electrical power production increase of turbine 2 caused by an increase of the velocity of the inflow into turbine 2,
- in *SOWFA Simulation Series 2*, turbine 2 is moved in the crosswind direction to reduce the overlap of its rotor with the wake of turbine 1, also causing an increase in the electrical power production of turbine 2.

For each yaw setting and position, a 600-second simulation was run. The wakes were allowed to develop during the first 200 seconds of the simulation, and then 400 seconds of simulated data was collected. By averaging of the power signals of the turbines over these 400 seconds, the results presented in Figure 4.2b were generated. In each case, the turbines use the baseline pitch and torque controllers defined in Jonkman et al. (2009). For the simulated flow conditions, both the upstream and the downstream turbine operate in a below-rated operating region (region 2) and thus use constant pitch, variable torque

control to maximize power production Jonkman et al. (2009); Burton et al. (2002a). For most cases, the downwind turbine produces less electrical energy than the upwind one because it is subjected to the low-speed wake of the upwind turbine.

SOWFA high-fidelity CFD simulations are typically run for a few days on a cluster with a few hundred processors Fleming et al. (2014d,b). Due to the complexity and computational costs of the SOWFA model, it is not suitable as an internal model for a wind plant controller. However, the data generated by SOWFA can be used to develop simplified models that can be directly used by the controller. In Section 4.3, we describe how the power data from *SOWFA Simulation Series 1* and *2* are used to identify such a simplified control-oriented model. In Section 4.5, SOWFA is used to evaluate the control techniques based on the simplified internal model in high-fidelity simulations.

4.3. DATA-DRIVEN PARAMETRIC WIND PLANT MODEL: FLORIS

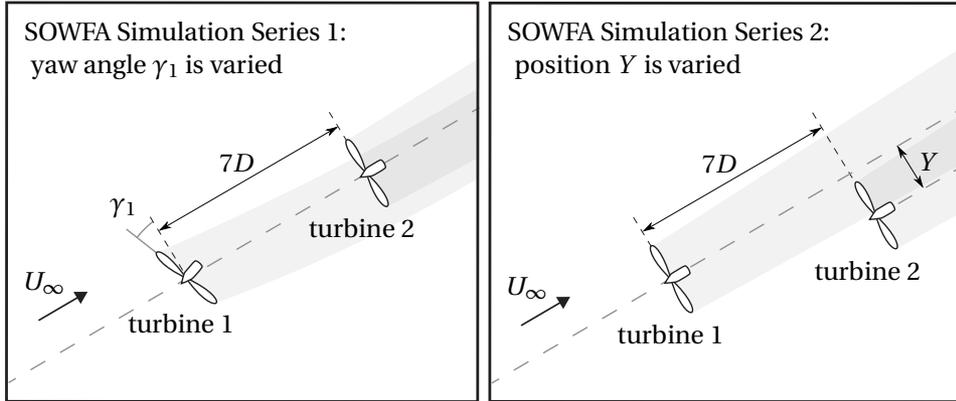
In this section, we explain the structure of a parametric model predicting the steady-state effects of yaw misalignment of different turbines on the electrical energy productions of wind turbines in a wind plant. It captures the effects of the yaw control on both the redirection of the wake behind the turbine, and on the velocity in the wake. This is important for predicting the electrical energy productions on downstream turbines, as is also pointed out in Choi et al. (2013). Since it includes both effects, for the remainder of the chapter we refer to the model as the FLORIS (FLOW Redirection and Induction in Steady-state) model.

4

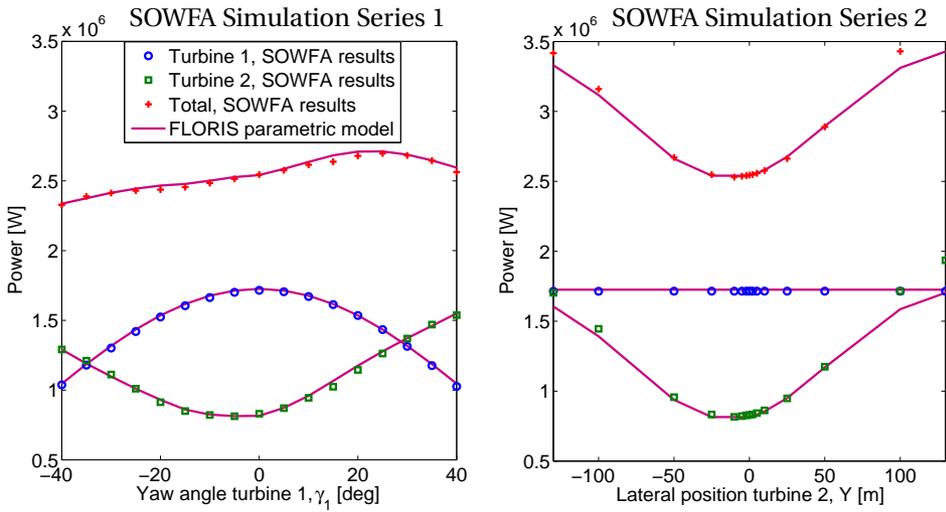
The FLORIS model is a combination of the Jensen model Jensen (1984); Katić et al. (1986), and a model for wake deflection through yaw first presented in Jiménez et al. (2010). Further, augmentations were made to the Jensen model in order to better model situations with partial wake overlap, and to the wake deflection model in order to include wake position offsets caused by rotor rotational effects. These augmentations also make that we can better fit the model with the power measurements obtained in *SOWFA Simulation Series 1* and *2*.

Figure 4.3 gives an overview of the different parts of the model, and of how it interacts with the yaw optimization algorithm of the wind plant controller. It also shows that *measured* power and yaw setting of turbines, as well as wind direction measurements at each turbine, are fed into the model. The measurements are used to estimate certain atmospheric conditions, being the current direction and free-stream velocity of the inflow into the wind plant. These yaw measurements should be distinguished from the *try-out yaw settings* that the optimization algorithm feeds into the model, and the corresponding *predicted turbine power outputs* that the model generates on the basis of those try-out yaw settings and the estimated inflow properties, and feeds back to the optimization algorithm.

In this section the different parts of the FLORIS model are presented. First, in Section 4.3.1 it will be explained how the electrical power productions of the turbines are calculated (the turbine power model in Figure 4.3). To calculate these power levels, estimates of the effective inflow speeds are used. These inflow speed estimates follow from the wake model. In the wake model we use a specific down-/crosswind coordinate



(a) Experimental setups



(b) Time-averaged power data

Figure 4.2: Setup and results for the *SOWFA Simulation Series 1* and *2*, as described in Section 4.2. The power data was used to find the parameters of the FLORIS model, see Section 4.3.

frame. Figure 4.3 shows that the turbine coordinates are transformed to these coordinates using measured wind directions at the turbines. This step is further explained in Section 4.3.2. Submodels for different wake properties are the wake decay, deflection and expansion models, also shown in Figure 4.3. These submodels are explained in Sections 4.3.3 to 4.3.5. Finally, 4.3 shows that there is the wake combination submodel, that defines how the wake effects of the different turbines are combined in order to find the effective inflow speeds at each turbine. This submodel is explained in Section 4.3.6. In explaining the model, different coefficients are introduced that serve as model parameters that are to be tuned to measurements from a wind plant. In this work, we use the power measurements from SOWFA to find the FLORIS model parameters, as discussed in Section 4.3.7.

4.3.1. TURBINE POWER

Let $\mathcal{F} = \{1, 2, \dots, N\}$ denote a set of indices that number the wind turbines in a wind plant, with N denoting the total number of turbines in the plant. The steady-state electrical power of a turbine $i \in \mathcal{F}$, denoted as P_i , is calculated as follows Bianchi et al. (2007):

$$P_i = \frac{1}{2} \rho A_i C_P(a_i, \gamma_i) U_i^3 \quad \forall i \in \mathcal{F}, \quad (4.1)$$

where ρ is the air density, A_i is the rotor swept area, and C_P is the power coefficient of the turbine, and U_i is the effective wind speed at the turbine. In non-yawed idealized conditions, the power coefficient is related to the axial induction factor of each turbine, defined as $a_i = 1 - U_{i,D}/U_i$, with $U_{i,D}$ being the wind speed at the rotor, and U_i the free-stream wind speed in front of turbine i , as $C_P(a_i) = 4a_i[1 - a_i]^2$ Bianchi et al. (2007). In the model presented here, we apply a correction on this relationship to account for the effect of the yaw misalignment angle γ_i on the rotor power coefficient, following the example of the experimental studies in Medici (2005). In addition, we use a constant scaling of the C_P value, η , to account for other losses. This results in

$$C_P(a_i, \gamma_i) = 4a_i[1 - a_i]^2 \eta \cos(\gamma_i)^{p_P}. \quad (4.2)$$

To match the maximum $C_P = 0.482$ and 94.4% generator efficiency reported in Jonkman et al. (2009) for the NREL 5-MW turbine that is used in *SOWFA Simulation Series 1* and 2, we use a loss factor $\eta = 0.768$. In Medici (2005), a parameter value $p_P = 2$ was found to fit data from wind tunnel tests with yawing turbines, but the parameter settings listed in Table 4.1 are found to fit the yaw-power relationship of the upstream turbine (turbine 1) in *SOWFA Simulation Series 1* (see Figure 4.2b). To find these parameters, we assume an idealized axial induction of $a_i = 1/3$ for below-rated conditions. The model can be extended to above-rated operation by make corrections on a_i as a function of inflow speed U_i (based on a maximum rated power), or by making a_i a function of pitch and tip-speed ratio, using knowledge of the C_P -curve.

In the remainder of this section, it will be described how the effective inflow speeds U_i at each turbine are estimated by the model, by predicting the steady-state wake characteristics as a function of the yaw angles.

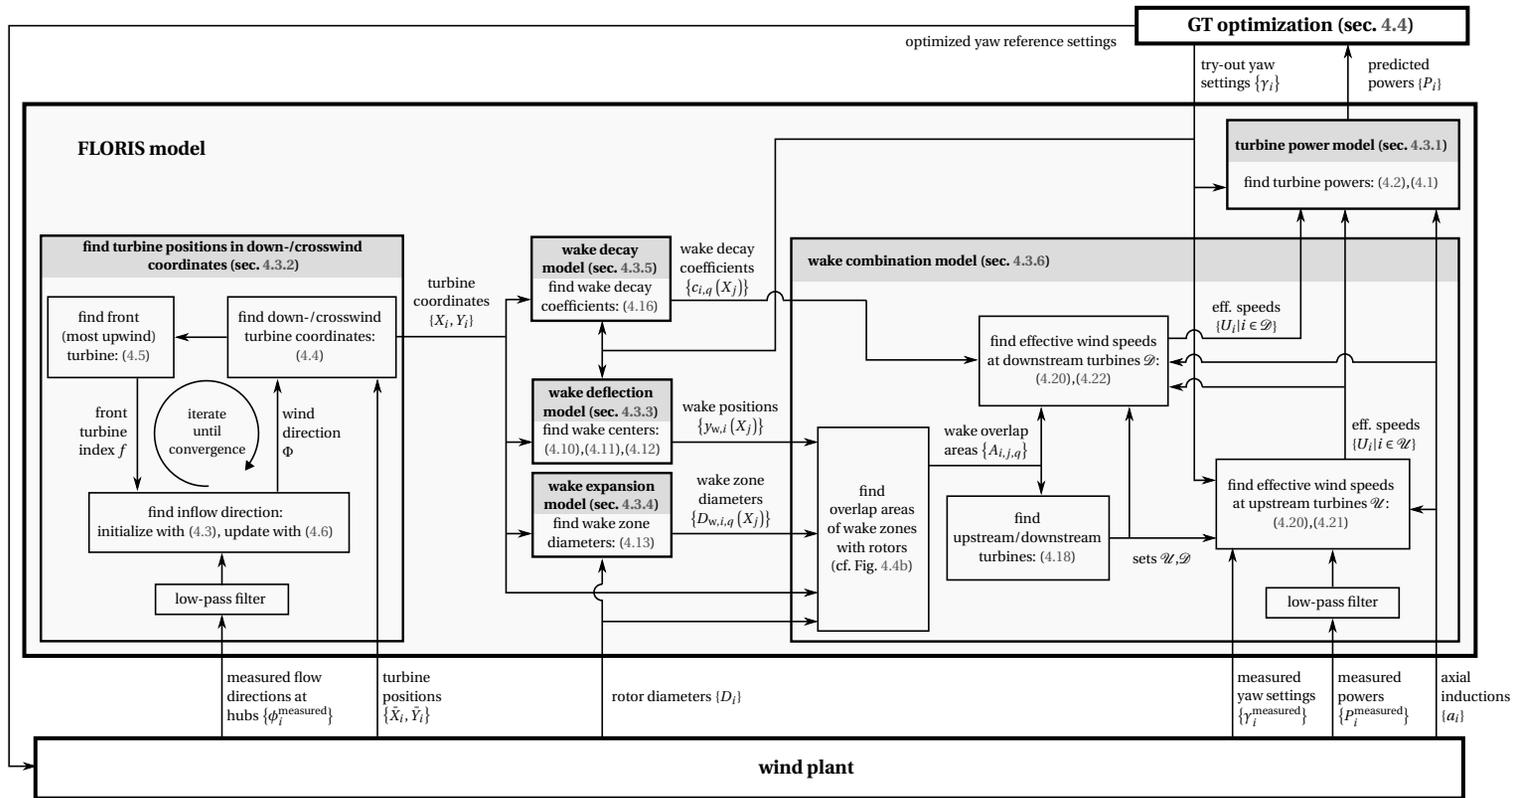


Figure 4.3: Overview of the FLORIS data-driven parametric model as it is implemented in the wind plant controller. Figure 4.1 shows the same basic control scheme, but in this figure the different submodels of the FLORIS simplified wind plant model are shown, and the identification block is omitted. The FLORIS model uses some measurements from the wind plant (shown below) to estimate in the inflow properties (speed and direction). The GT optimization algorithm (right-top) uses the FLORIS model to test yaw settings for these particular inflow conditions, and finally sends optimized yaw settings as reference signals to the wind plant. In the scheme shown here, the shorthand notation $\{\theta_i\}$ is used for the set $\{\theta_i | i \in \mathcal{F}\}$, where θ_i is a certain property of a wind turbine i in the wind plant.

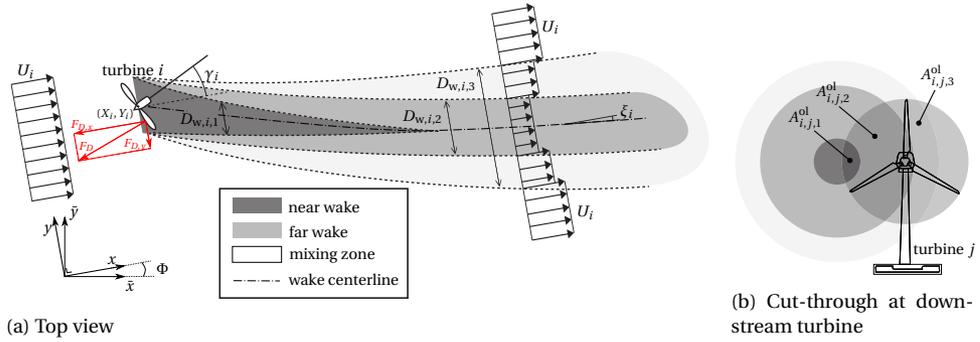


Figure 4.4: The three different wake zones of the FLORIS parametric model. The free-stream wind vectors are indicated as arrows with length U_i (the free-stream velocity). Inside the wake zones the wind vectors have a reduced velocity (see Section 4.3.5). The areas overlapping with a downstream rotor j , $A_{i,j,q}^{ol}$, are used to calculate the effective wind speed at turbine j , see (4.20),(4.22).

4.3.2. INFLOW DIRECTION AND THE DOWNWIND-CROSSWIND COORDINATE FRAME

In order to describe the spatial properties of the wakes behind the turbines, we adopt a Cartesian coordinate framework (x, y) in which the x -axis is pointing downwind along an estimated mean inflow direction in the wind plant, and the y -axis is pointing orthogonal to the x -axis in the horizontal direction, i.e., along the crosswind direction (as illustrated in Figure 4.4). Then naturally the z -axis represents the altitude. In this work, we assume that each turbine has the same hub-height, and the turbine locations in this downwind-crosswind coordinate frame are denoted as $(X_i, Y_i) \forall i \in \mathcal{F}$.

The mean inflow direction, denoted by Φ , can be estimated in several ways. In the model as used in the CFD simulation examples in Section 4.5, it is found using wind direction measurements at the most upwind turbine. However, to determine which is the front turbine, we need some initial guess of the free-stream wind direction, therefore we use the iterative procedure described below. The steps of this procedure are also illustrated in the left-most block in the FLORIS model scheme in Figure 4.3. The steps are as follows:

1. We average the flow direction measurements at the hub of each turbine $i \in \mathcal{F}$ ², denoted as ϕ_i^{measured} , to provide a first estimate of the inflow direction:

$$\Phi = \frac{1}{N} \sum_{i=1}^N \phi_i^{\text{measured}}. \quad (4.3)$$

2. Then, the turbine positions in downwind/crosswind coordinates are calculated according to the estimated wind direction. If $\{\bar{X}_i, \bar{Y}_i\}$ are the turbine coordinates relative to the same Cartesian coordinates (\bar{x}, \bar{y}) to which the wind direction Φ is

²In the CFD simulation examples, the wind direction at the turbine hubs is estimated by sampling the horizontal velocity components (\bar{u}_i, \bar{v}_i) at the hub location of each turbine $i \in \mathcal{F}$ from the flow field calculated by the CFD simulator, and calculating the direction in the horizontal plane as $\phi_i = \tan^{-1}(\bar{v}_i / \bar{u}_i)$.

measured (see also Figure 4.4), the downwind-crosswind turbine coordinates are:

$$\begin{bmatrix} X_i \\ Y_i \end{bmatrix} = \begin{bmatrix} \cos(-\Phi) & -\sin(-\Phi) \\ \sin(-\Phi) & \cos(-\Phi) \end{bmatrix} \begin{bmatrix} \bar{X}_i \\ \bar{Y}_i \end{bmatrix}. \quad (4.4)$$

3. Then, it is established which turbine is the front (most upwind) turbine, and it is assumed that the mean inflow direction is equal to the wind direction measured at that turbine (i.e., we assume a uniform direction of the free-stream inflow to the wind plant):

$$f = \underset{i \in \mathcal{F}}{\operatorname{argmin}} X_i \quad (4.5)$$

$$\Phi = \phi_f^{\text{measured}}. \quad (4.6)$$

4. We repeat step (2) and (3), until convergence (i.e., no change in estimated wind direction Φ).

The wind direction estimation iterative procedure will generally converge to a wind direction measured at a certain turbine within two or three iterations, in our simulation examples. Note that in our implementation of the model as illustrated in Figure 4.3, the wind direction measurements at the hub of the turbines, defined relative to the mesh coordinates, are low-pass filtered, in order to filter out small-scale turbulence effects.

4.3.3. WAKE DEFLECTION

Yawing a turbine rotor causes the thrust force that the rotor exerts on the flow, F_D , to rotate in such a way that a crosswind component is induced Jiménez et al. (2010), which causes the wind flow to deflect in the direction opposite to the yaw rotation (see Figure 4.4a). Because the wake deflection is induced by the thrust force, the amount of deflection is a function of the thrust coefficient of the turbine $C_T = 2F_D/(\rho A_i U_i^2)$. When the yaw is not misaligned with respect to the wind direction (i.e., $\gamma_i = 0$), the thrust coefficient is related to the axial induction factor a_i of the rotor of a turbine i , as follows Bianchi et al. (2007):

$$C_T(a_i) = 4a_i[1 - a_i]. \quad (4.7)$$

The following heuristic relationship between the yaw angle of a turbine i , γ_i , the thrust factor C_T of the turbine in non-yawed conditions, and the angle of the centerline of its wake ξ_i at a downstream location $x > X_i$, was derived in Jiménez et al. (2010):

$$\xi_i(x) \approx \frac{\xi_{\text{init}}(a_i, \gamma_i)}{\left[1 + 2k_d \frac{x - X_i}{D_i}\right]^2} \text{ with } \xi_{\text{init}}(a_i, \gamma_i) = \frac{1}{2} \cos^2(\gamma_i) \sin(\gamma_i) C_T(a_i), \quad (4.8)$$

where D_i is the rotor diameter of turbine i , k_d is a model parameter that defines the sensitivity of the wake deflection to yaw, and where ξ_{init} is the initial angle of the wake at the rotor. Relation (4.8) is elegant in the sense that it only has one unknown parameter, k_d , to be tuned. This k_d parameter defines the recovery of the wake flow direction towards the main inflow direction Φ . By integrating the tangent of the wake centerline angle over

x , the yaw-induced lateral offset of the wake center with respect to the hub of a turbine i , denoted as $\delta y_{w,yaw,i}$, can be found:

$$\delta y_{w,yaw,i}(x) = \int_0^{x-X_i} \tan(\xi_i(x)) dx. \quad (4.9)$$

This integral can be approximated by integrating the second-order Taylor series approximation of $\xi(x)$, yielding

$$\delta y_{w,yaw,i}(x) \approx \frac{\xi_{\text{init}}(a_i, \gamma_i) \left[15 \left[\frac{2k_d [x-X_i]}{D_i} + 1 \right]^4 + \xi_{\text{init}}(a_i, \gamma_i)^2 \right]}{\frac{30k_d}{D_i} \left[\frac{2k_d [x-X_i]}{D_i} + 1 \right]^5} - \frac{\xi_{\text{init}}(a_i, \gamma_i) D_i \left[15 + \xi_{\text{init}}(a_i, \gamma_i)^2 \right]}{30k_d}. \quad (4.10)$$

Further, in the simulations described in Fleming et al. (2014d), it was found that a small lateral wake deflection occurs when the turbine is not yawed (i.e., $\gamma_i = 0$). This deflection can be explained by vertical shear in the boundary layer and wake rotation: in reaction to the rotor rotating clockwise, the wake will rotate counterclockwise, and therefore low speed flow in the lower part of the boundary layer will be rotated up and to the right, and high speed flow in the upper part of the boundary layer will be rotated down and to the left. As a result the velocity deficit at the right part of the wake (looking downstream) increases, so the wake deflects to the right. Since in *SOWFA Simulation Series 1* and *2* the wake behavior was tested for a single mean wind velocity with a limited velocity variation caused by turbulence, the exact dependence of the wake deflection on the rotor speed could not be derived from the power data obtained. Therefore, this rotation induced wake lateral offset is parameterized through a simple linear function of the downstream distance front the rotor:

$$\delta y_{w,rotation,i}(x) = a_d + b_d [x - X_i]. \quad (4.11)$$

Combining the rotation induced and yaw induced components, the position of the wake center of a turbine i at a downstream location $x > X_i$ is given by:

$$y_{w,i}(x) = Y_i + \delta y_{w,rotation,i}(x) + \delta y_{w,yaw,i}(x). \quad (4.12)$$

4.3.4. WAKE EXPANSION

The Jensen model Jensen (1984); Katić et al. (1986) assumes a wake that is expanding proportionally to the axial downstream distance from the rotor, and a wind velocity in the wake that is uniform in the lateral direction. In reality, the velocity will be closer to the free-stream velocity towards the edges of the wake, due to mixing Sanderse et al. (2011). Therefore, we expand the Jensen model in order to better model partial wake situations, by dividing the wake in three zones that also expand proportionally with the distance from the rotor, but each with their own expansion factor (see Figure 4.4a). The diameters of the wake zones behind a turbine i are given by:

$$D_{w,i,q}(x) = \max(D_i + 2k_e m_{e,q} [x - X_i], 0) \quad (4.13)$$

for $x > X_i$, and index $q = 1, 2, 3$ numbering the different zones, D_i being the rotor diameter of turbine i , and with parameters $m_{e,q}$, k_e being coefficients defining the expansion of the zones. The different wake zones can be referred to as the ‘near wake’ ($q = 1$), ‘far wake’ ($q = 2$), and ‘mixing zone’ ($q = 3$), in accordance with the terms that are commonly used in literature to describe wake characteristics Vermeer et al. (2003); Crespo et al. (1999); Sanderse et al. (2011). The scaling parameter for the expansion of the near wake, $m_{e,1}$, is typically set to a negative value, which prescribes that the cross-section of the near wake zone is decreasing to zero with the distance to the rotor. The extension to different wake zones makes that we can better match the data from *SOWFA Simulation Series 1* and 2, as is shown in Section 4.3.7.

4.3.5. WIND VELOCITY IN A SINGLE WAKE

By definition, the axial induction is the relative amount of velocity drop at the rotor with respect to the inflow velocity. From actuator disk theory, it follows that the relative rotor-induced drop of the velocity behind the rotor is two times the axial induction factor Bianchi et al. (2007). In the wake behind the rotor, the velocity will gradually recover to the free-stream velocity by turbulence-induced mixing. The Jensen model assumes that the time-averaged velocity deficit in the far wake decays quadratically with the expansion of the wake. In Annoni et al. (2014b) it is shown that the parameters of the Jensen model can be tuned to obtain a good fit of the time-averaged velocity profile in the far wake as predicted by the SOWFA model for non-yawed conditions. An extension made in the FLORIS model, in order to better fit the power data from *SOWFA Simulation Series 1* and 2, is that the wake is divided into three zones, as described in the previous section, and that the velocity deficit decays quadratically with the distance from the rotor, rather than being directly related to the wake expansion. Hence, the velocity profile behind a turbine i is modeled as:

$$U_{w,i}(x, y) = U_i [1 - 2a_i c_i(x, y)] \quad (4.14)$$

for $x > X_i$, with U_i again denoting the free-stream speed in front of the turbine, and with the wake decay coefficient $c_i(x, y)$ being a piecewise constant function of the lateral offset of the location y with respect to the wake center of turbine i :

$$c_i(x, y) = \begin{cases} c_{i,1} & \text{if } |r| \leq D_{w,i,1}(x)/2 \\ c_{i,2} & \text{if } D_{w,i,1}(x)/2 \leq |r| \leq D_{w,i,2}(x)/2 \\ c_{i,3} & \text{if } D_{w,i,2}(x)/2 \leq |r| \leq D_{w,i,3}(x)/2 \\ 0 & \text{if } |r| > D_{w,i,3}(x)/2 \end{cases} \quad (4.15)$$

with $r = y - y_{w,i}(x)$

and with the local wake decay coefficient for each zone given by:

$$c_{i,q}(x) = \left[\frac{D_i}{D_i + 2k_e m_{U,q}(\gamma_i) [x - X_i]} \right]^2. \quad (4.16)$$

The coefficients $m_{U,q}$ are parameters defining how quickly the different wake zones decay. Following a similar approach to that in Section 4.3.1, the wake decay rates are adjusted for the rotor yaw angle in order to fit the data from *SOWFA Simulation Series 1*

and 2, by empirically deriving the following relationship between the coefficient $m_{U,q}$ and the yaw angle γ_i :

$$m_{U,q}(\gamma_i) = \frac{M_{U,q}}{\cos(a_U + b_U \gamma_i)} \quad (4.17)$$

for $q = 1, 2, 3$ with model parameters $M_{U,q}$, a_U , b_U .

In order to better predict the effect of axial-induction-based control on wind plant performance, modifications to the wake velocity model in FLORIS were made by Annoni et al. (2014a). These modifications are not included in the model presented here, that is aimed at predicting the effect of yaw-based wake redirection control.

4.3.6. COMBINING WAKES TO FIND THE TURBINE EFFECTIVE WIND VELOCITIES

In the submodel described in this section we combine the effects of the wake zones of different turbines, in order to estimate the effective inflow velocity at each turbine. The different parts of this submodel are illustrated in the lower-right block in the FLORIS model scheme in Figure 4.3. The submodel described in Section 4.3.5 estimates the velocity deficits in the wake with respect to a certain free-stream inflow speed. Generally, we cannot assume the free-stream inflow velocity into the wind plant to be known, but by inverting relationship (4.1), we can estimate the effective wind speeds at the front turbines from turbine power and yaw angle measurements, denoted as P_i^{measured} and $\gamma_i^{\text{measured}}$, respectively. Then, the wind speeds at the downstream turbines are estimated by combining the effect of the wakes, weighting the wake zones by their overlap with the rotors using the root-sum-square method of Katić et al. (1986). This results in the following formulations.

First, the overlapping areas between turbine rotors and the different zones of the wakes are calculated from the wake center and wake diameter predictions described in Section 4.3.3 and 4.3.4, using basic geometry. We denote the overlapping area between a wake zone q of a turbine i and a rotor of a downstream turbine j , by $A_{i,j,q}^{\text{ol}}$, see Figure 4.4b. Then, let $\mathcal{U} \subset \mathcal{F}$ denote the set of front, upstream turbines that are not influenced by other turbines through wake interaction, because their rotors do no overlap with any wakes, and \mathcal{D} the set of turbines that *are* influenced by other turbines, i.e.:

$$\begin{aligned} \mathcal{U} &= \left\{ j \in \mathcal{F} \mid A_{i,j,q}^{\text{ol}} = 0 \forall i \in \mathcal{F}, q \in \{1, 2, 3\} \right\} \\ \mathcal{D} &= \left\{ j \in \mathcal{F} \mid j \notin \mathcal{U} \right\} \end{aligned} \quad (4.18)$$

Further, $u(j)$ denotes the index of a turbine in the set \mathcal{U} that has the largest overlap area with a turbine $j \in \mathcal{D}$ when compared to other turbines in the set \mathcal{U} :

$$u(j) = \arg \max_{i \in \mathcal{U}} \left(\sum_{q=1}^3 A_{i,j,q}^{\text{ol}} \right) \forall j \in \mathcal{D}. \quad (4.19)$$

Figure 3.3 gives an illustrative example of how the sets \mathcal{U} and \mathcal{D} and the mapping $j \rightarrow u(j)$ are defined in a particular case. The effective inflow speed at a turbine $j \in \mathcal{D}$ is assumed to be the velocity at turbine $\rightarrow u(j)$ multiplied with a factor that represents

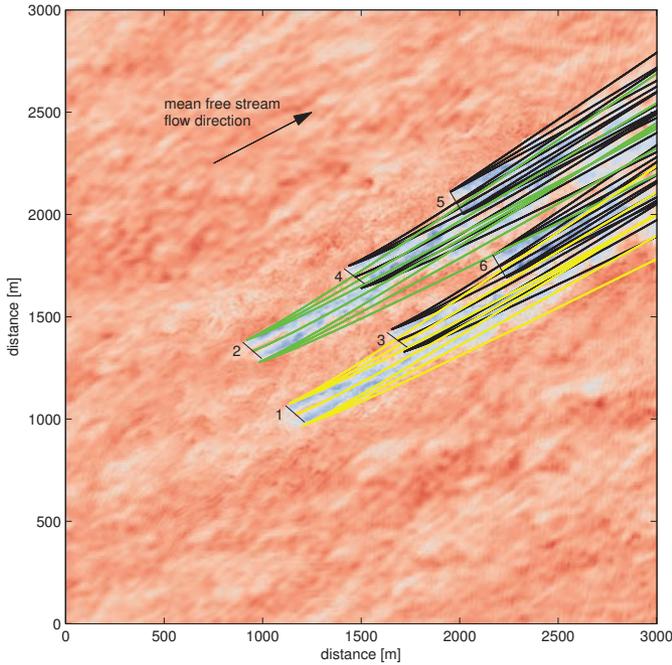
the effects of the different wake zones, which are weighted by the overlap of these wake zones with the rotor. This results in the following relations for estimating the effective wind speeds U_j at each turbine $j \in \mathcal{F}$:

$$U_j = \begin{cases} f_1(j) & \forall j \in \mathcal{U} \\ f_2(j) & \forall j \in \mathcal{D} \end{cases} \quad (4.20)$$

with functions:

$$f_1(j) = \left[\frac{2P_j^{\text{measured}}}{\rho A_j C_P(a_j, \gamma_j^{\text{measured}})} \right]^{1/3} \quad (4.21)$$

$$f_2(j) = U_{u(j)} \left[1 - 2 \sqrt{\sum_{i \in \mathcal{F}: X_i < X_j} \left[a_i \sum_{q=1}^3 c_{i,q}(X_j) \min\left(\frac{A_{i,j,q}^{\text{ol}}}{A_j}, 1\right) \right]^2} \right]. \quad (4.22)$$



set \mathcal{F}	{1, 2, 3, 4, 5, 6}
set \mathcal{U}	{1, 2}
set \mathcal{D}	{3, 4, 5, 6}

turbine index $i \in \mathcal{F}$	1	2	3	4	5	6
turbine $u(i)$	n/a	n/a	1	2	2	1

Figure 4.5: The above picture shows the hub-height flow field of a 3-by-2 wind plant, with the centerline and boundaries of the wake zones as predicted by the FLORIS model shown as solid lines. The set definitions for this case, $\mathcal{F}, \mathcal{U}, \mathcal{D}$, and the mapping $i \rightarrow u(i)$ are shown in the table underneath the picture. The mapping $i \rightarrow u(i)$ follows from the fact that the rotors of turbines 3 and 6 have the largest overlap with the model-predicted wake of turbine 1 (shown in yellow), when compared to other turbines in the set \mathcal{U} , and the rotors of turbines 4 and 5 have the largest overlap with the wake of turbine 2 (shown in green).

4.3.7. FITTING THE WAKE MODEL PARAMETERS

By fitting to the power data from turbine 2 in *SOWFA Simulation Series 1* and *2*, the parameters for wake deflection, expansion and decay were tuned ‘manually’ (see Figure 4.2b and Table 4.1). The results were validated by comparing the resulting wake velocity profiles for a single yawed turbine with the corresponding data generated by SOWFA in the simulation experiments described in Fleming et al. (2014d); see Figure 4.6 for this comparison. It can be seen that by dividing the wake in different zones, as described in Section 4.3.4, and introducing a rotation-induced wake position offset (Section 4.3.3), we are able to better match the wake velocity profile.

turb. power	wake							
	deflection		expansion		velocity			
η 0.768	k_d 0.15	k_e 0.065	$M_{U,1}$ 0.5	a_U 5				
p_P 1.88	a_d -4.5	$m_{e,1}$ -0.5	$M_{U,2}$ 1	b_U 1.66				
	b_d -0.01	$m_{e,2}$ 0.22	$M_{U,3}$ 5.5					
		$m_{e,3}$ 1						

Table 4.1: FLORIS model parameters

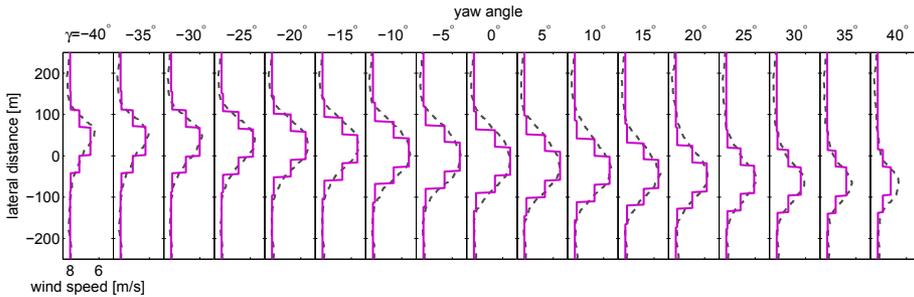


Figure 4.6: Time-averaged wake velocity profiles at turbine hub-height at a 7 rotor diameter distance from the rotor, for different yaw angles of the rotor γ , as calculated with SOWFA (dashed) in the simulation described in Fleming et al. (2014d), and with the FLORIS model (solid).

4.4. WIND PLANT YAW OPTIMIZATION USING A GAME-THEORETIC APPROACH

In this work, we use the game-theoretic (GT) approach of Marden et al. (2013) to perform an optimization of the yaw settings of the turbines in a wind plant with the objective of maximization of the total wind plant power production. The GT approach performs the optimization by making random perturbations to the yaw settings and holds the settings as a baseline setting if they yield an improvement of the wind plant total power, so as to iteratively find the global maximum of the wind plant total power. Following the control scheme of Figure 4.1 (shown in more detail in Figure 4.3), the optimization is based on the turbine power predictions of the FLORIS model, hence the randomized search of

the GT approach is performed using the simplified model, and once the optimized settings are found, they are applied on the wind plant. In the remainder of this section, we will explain the optimization algorithm in more detail. A specification of the different parameters of the optimization algorithm for a particular case study follows in Section 5.2.

In Algorithm 3, the (simplified) optimization scheme of the GT approach is given as it is implemented in our simulation examples. The randomized perturbation on the yaw settings of the turbines in the wind plant model takes place in lines 9-15 of the algorithm. In lines 10-11, a randomized process determines whether the yaw setting of a specific turbine is updated; E is the probability of using a new random setting for γ_i , instead of keeping the baseline setting, denoted by $\bar{\gamma}_i$. By setting this search rate E , we can control the amount of updates on turbines taking place at the same time, and as such tune the convergence properties of the GT algorithm. If the yaw setting of a turbine is indeed updated, it is randomly selected from the range $[\gamma_{\min}, \gamma_{\max}]$ (in lines 12-13). This range of possible yaw settings is discretized with an interval $\Delta\gamma$ in order to improve the convergence speed. Testing whether the updated yaw settings improve the total wind plant power, based on an evaluation of the FLORIS model, takes place in lines 18-22. If they yield an improvement, the updated yaw settings are stored as a baseline setting $\bar{\gamma}_i$ (see line 20). After that evaluation, a new iteration of the GT algorithm will follow.

4.5. WIND PLANT YAW OPTIMIZATION SIMULATION EXAMPLES

We perform an evaluation of the online yaw optimization wind plant control strategy based on the FLORIS parametric model, by using it in SOWFA simulations of a small wind plant.

4

4.5.1. SIMULATION SET-UPS

The simulated wind plant consists of two rows with three NREL 5-MW baseline turbines Jonkman et al. (2009) each, with a 5 rotor diameter spacing in the downwind direction, and 3 rotor diameters in the crosswind direction. We simulate a constant wind direction, and 3 configurations of the wind plant, in which the setup is rotated 0° , 5° and 10° with respect to the wind direction. For each of these configurations, we run simulations of two cases:

- a case with the model-based control performing plant-wide optimization of the yaw settings enabled, and
- a case in which the yaw settings that yield maximum power for each individual turbine are used, i.e., each rotor is aligned perpendicular to the mean wind direction. Since these yaw settings result in maximization of electrical power of the turbine itself, but not in production maximization on a plant-level, we refer to these settings as ‘greedy’.

The turbine positions, and the SOWFA-calculated flow fields for each of the cases are shown in Figure 4.7. In this figure it can be seen that yaw misalignment indeed leads

Algorithm 3 The pseudocode below shows a Game Theoretic approach for wind plant control, performing optimization of the yaw angles for increased electrical energy production. Index k denotes the iterations of the optimization. The variables $\bar{\gamma}_i$ and \bar{P}_i are used to store baseline values of the control variables and the corresponding turbine powers (yielding the maximum wind plant power found so far). \mathcal{U} denotes a uniform distribution, \mathbb{Z}_n denotes the set of integers $[0, 1, \dots, n]$.

```

1:  $\gamma_i \leftarrow 0 \forall i \in \mathcal{F}$ 
2:  $k \leftarrow 0$ 
3:  $n \leftarrow \frac{\gamma_{\max} - \gamma_{\min}}{\Delta\gamma}$ 
4: update  $P_i(\gamma_i) \forall i \in \mathcal{F}$  using (4.20),(4.1)
5:  $\bar{P} \leftarrow \sum_{i=1}^N P_i(t)$ 
6:  $\bar{\gamma}_i \leftarrow \gamma_i$ 
7: loop
8:    $k \leftarrow k + 1$ 
9:   for all  $i \in \mathcal{F}$  do
10:      $\mathcal{R}_1 \leftarrow$  random value  $\sim \mathcal{U}(0, 1)$ 
11:     if  $\mathcal{R}_1 < E$  then
12:        $\mathcal{R}_2 \leftarrow$  random value from the set  $\mathbb{Z}_n$ 
13:        $\gamma_i \leftarrow \gamma_{\min} + \mathcal{R}_2 \Delta\gamma$ 
14:     else
15:        $\gamma_i \leftarrow \bar{\gamma}_i$ 
16:     end if
17:   end for
18:   update  $P_i \forall i \in \mathcal{F}$  using (4.20),(4.1)
19:   if  $\sum_{i=1}^N P_i(\gamma_i) > \bar{P}$  then
20:      $\bar{\gamma}_i \leftarrow \gamma_i \forall i \in \mathcal{F}$ 
21:      $\bar{P} \leftarrow \sum_{i=1}^N P_i(t)$ 
22:   end if
23: end loop

```

to a redirection of the wake, and an increase of flow velocity in the wake caused by a reduction of axial induction.

In these SOWFA simulations, an inflow with a 6% turbulence intensity and an 8 m/s mean velocity is used, which is the same inflow condition as in *SOWFA Simulation Series 1* and *2* described in Section 4.2. Note that a different spacing between the turbines in the wind direction is used in *SOWFA Simulation Series 1* and *2* to obtain the parameters of the FLORIS model, namely 7 rotor diameters, and in that sense we use the model for extrapolation. The wind plant setup is placed in a 3km (horizontal length) by 3km (horizontal width) by 1km (height) mesh, see Figure 4.7. The smallest mesh cells for the CFD calculation, which contain the turbine rotors, the axial induction zones of the rotor and the wakes between the turbines, have a size of $3\text{m} \times 3\text{m} \times 3\text{m}$. Further away from the turbines the mesh is coarsened to $6\text{m} \times 6\text{m} \times 6\text{m}$ cells, and then to $12\text{m} \times 12\text{m} \times 12\text{m}$ cells, resulting in a total of $32 \cdot 10^6$ cells. Using a time-step of 0.02s, a 1000s simulation is performed for each of the six cases. Because of the high mesh and time resolution required, the computational cost of the CFD simulations is high: 59 hours of distributed computation on 512 processors for each 1000s simulation.

4.5.2. SPECIFICATIONS OF THE PLANT-WIDE OPTIMIZATION CONTROLLER USED IN THE SIMULATION EXAMPLES

In accordance with the scheme in Figure 4.3, the electrical power production, yaw, and local wind direction measurements as calculated by SOWFA are directly fed into the internal FLORIS model. The FLORIS model uses these signals to estimate the inflow properties (effective wind speed and direction). A first-order low-pass filter with a -3dB cut-off frequency of 2 mHz is used on the measured power and wind direction signals in order to make sure that the wind plant yaw controller only responds to the slower trends in the changes of the inflow conditions. When the plant-wide controller is switched on, in each 0.02s timestep of the simulation the yaw optimization is performed. The GT algorithm parameters are set to $\gamma_{\min} = 0^\circ$, $\gamma_{\max} = 40^\circ$ (offset relative to wind direction), $\Delta\gamma = 0.05^\circ$, $E = 0.2$, and in each optimization 1000 iterations of the yaw settings are tested on the internal FLORIS model, which is sufficient for convergence to a maximum in the predicted power production. In each iteration of the optimization it takes about 0.04ms to evaluate the C-implementation of the FLORIS model on a single CPU. The search is restricted to positive yaw angles since previous simulation studies Fleming et al. (2014d) showed that for the given inflow conditions, yawing in the positive direction yields a reduction of relevant structural loads on the turbine, while negative yaw increases blade loads. After the optimization procedure, the baseline yaw setting for each turbine as calculated by the optimization algorithm is set as a reference to which the turbines respond with a maximum yaw rate of $1^\circ/\text{s}$. In the following timestep, the optimization is initialized with the previous yaw reference setting.

In this case study, an unrealistically high update rate for the yaw reference settings was used when considering the maximum yaw rate of $1^\circ/\text{s}$. In practice, one would run the optimization algorithm far less frequently, and the 0.04s needed to run the optimization would be sufficiently short to allow for online optimization. The time between yaw reference updates on a real wind turbine may be in the order of tens of seconds to min-

utes (examples are discussed in Hau (2013a) and Kragh et al. (2013b)). Since both a conventional ‘greedy’ yaw control system and the proposed plant-level optimized yaw control system respond to changes of the wind direction, the yaw reference update rate can be adjusted to make a trade-off between additional yaw actuator usage and electrical power production, where in principle the plant-level optimized scheme does not need more actuator usage than the ‘greedy’ scheme.

4.5.3. DETAILED RESULTS OF SIMULATIONS WITH THE 5 DEGREE ROTATED WIND PLANT

For the cases with the 5° rotated wind plant, relevant signals are shown in Figure 4.8. We compare the SOWFA simulation with the ‘greedy’ settings, with a SOWFA simulation in which the plant-wide optimization controller is switched on after 400s. Figure 4.8a compares the electrical power output calculated by SOWFA and predicted by the FLORIS model for both cases. Figure 4.8b shows the yaw angles for both cases, and the wind direction estimated by the controller.

Figure 4.8a shows that the predictions given by FLORIS are not always closely matching the SOWFA results for each turbine, since the FLORIS model does not include the transient effects related to the wakes taking some time to propagate through the wind plant, and also higher-frequency variations related to turbulence are not included in the FLORIS model. Still, the FLORIS-model predictions are accurate enough to be used by the GT optimization algorithm in the controller for the calculation of yaw settings that yield a significant increase in total power production of the wind plant, when compared to the case with the greedy settings (see lower-left plot in Figure 4.8a). In the SOWFA results, over the period of 50s to 500s after plant-wide optimization control is switched ON, on average 13% more electrical energy is produced compared to the ‘greedy’ yaw case. In Figure 4.8b it can be seen that the online implementation of the GT model-based optimization responds to changes in the inflow conditions (direction and velocity), which are predicted based on the measured signals.

4.5.4. OVERVIEW OF ELECTRICAL POWER PRODUCTION AND LOAD RESULTS OF SIMULATIONS WITH DIFFERENT WIND PLANT CONFIGURATIONS

Table 4.2 shows the time-averaged results for the total wind plant electrical power production calculated with SOWFA and estimated with FLORIS, for both ‘greedy’ and plant-wide optimized settings for the 0°, 5° and 10° rotated configurations of the wind plant. Figure 4.9 shows the power productions for each individual turbine. When comparing the FLORIS and SOWFA predictions on the power production, we again see that there is not a perfect match, but the FLORIS predictions are accurate enough to enable a power production increase when using them for optimization of the yaw angles. In each of the cases with the plant-wide optimization controller, it can be seen that the loss of power through yawing on upstream turbines is compensated by a larger power gain in a downstream turbine. When comparing simulations with different configurations of the wind plant, we can see the effect of changes in the mean wind direction on the wake interaction in the wind plant: for a 10° rotation, the power production gain from the yaw

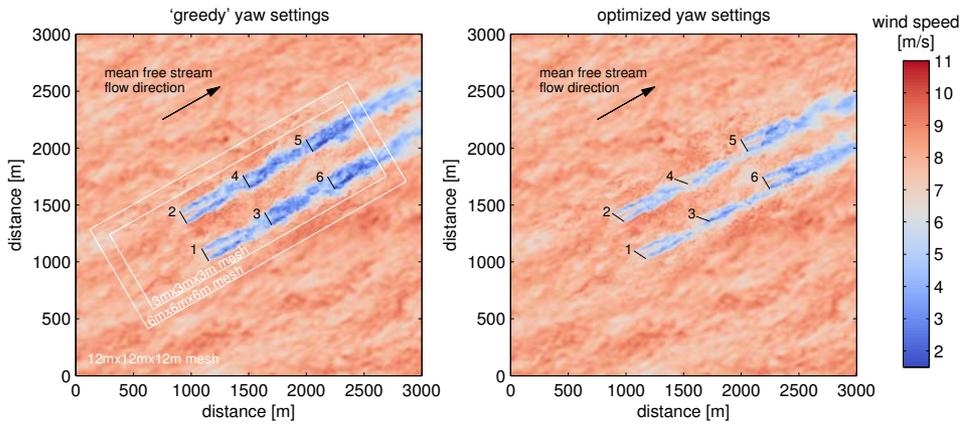
optimization is much smaller. This is because in that case also with the ‘greedy’ settings there is little overlap of the wakes with the downstream turbines (see also the flow fields in Figure 4.7). Table 4.3 lists the optimal yaw angles as calculated by the GT optimization, which shows that in the case with the 10° rotated wind plant, the optimized yaw angles are much smaller.

Also included in Figure 4.9 are relevant structural loads on each turbine for each case, as estimated by the FAST dynamic model of each turbine. It has to be remarked that the validation of load predictions from SOWFA is still ongoing work. Loads are computed for blade out-of-plane (OOP) bending moments, drive-train low-speed shaft (LSS) torsion, yaw bearing moments and tower bending moments. For each of these load signals, a damage equivalent load (DEL) is computed, which is a standard measure of fatigue damage Buhl Jr. (2008). The comparison shows that for most of the investigated loads, a reduction of DEL is observed when the yaw settings are optimized by the plant-wide controller, even though this was not the objective of the optimization. The observation that yaw misalignment on the upstream turbines can reduce the loads on these turbines themselves, is consistent with the findings in Kragh et al. (2013a), and load reductions on downstream turbines can be attributed to a reduction of wake overlap. There are increases in some DELs on some of the downstream turbines, which can be explained by an increase of imbalance through partial wake overlap. In Fleming et al. (2014b) it is suggested that these loads can be mitigated using individual pitch control. Also, a drive-train damper can mitigate the increase of loads on the LSS observed at turbine 5 for the 0° rotated wind plant case. There are substantial differences in the DELs when comparing the two 3-turbine rows in the wind plant, even though turbine spacing in each row is the same. This can be explained by the fact that the turbulence in the inflow is different for each of these rows, and that the data is averaged over a relatively short period.

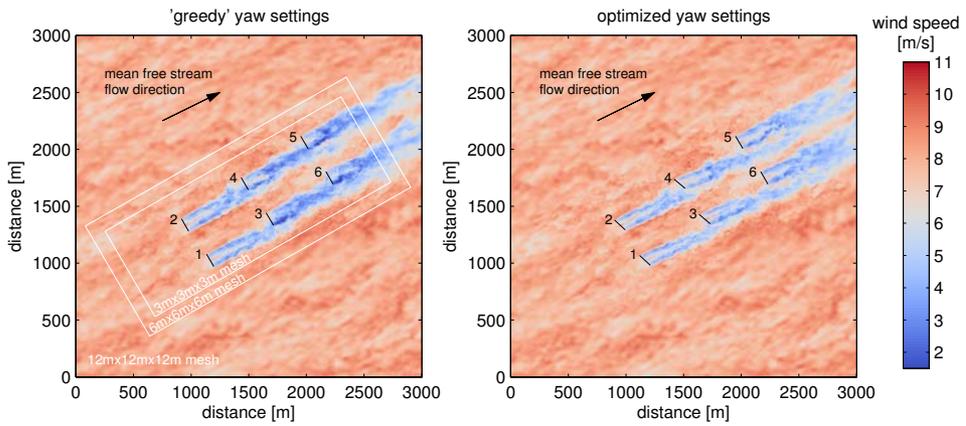
Finally, it should be remarked that the expected beneficial effects of the proposed yaw optimization control on a real large wind plant are smaller than in this case study. In this case study there is a relatively small inter-turbine spacing in the flow direction, and a relatively slow wake recovery through mixing because of the low turbulence-level in the inflow, and the neutral boundary layer conditions. These are conditions that make that the wake losses are relatively large in this wind plant, and thus also we see a large effect of mitigating these losses through control.

	greedy yaw			optimized yaw			increase
	SOWFA	FLORIS	error	SOWFA	FLORIS	error	
wind plant rotated 0°	6.68MW	6.34MW	5.27%	7.55MW	7.66MW	1.55%	13.03%
wind plant rotated 5°	8.75MW	8.75MW	0.08%	9.91MW	9.99MW	0.79%	13.19%
wind plant rotated 10°	10.80MW	11.04MW	2.23%	10.91MW	11.22MW	2.76%	1.04%

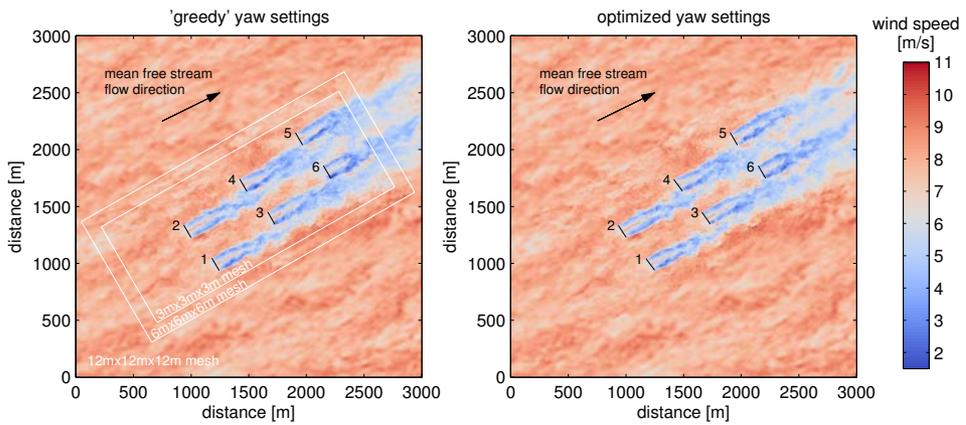
Table 4.2: Total wind plant electrical power production in the SOWFA simulation results, and as predicted by the FLORIS model, for both ‘greedy’ control and plant-wide optimized yaw control. The data is averaged over a period of 50s to 500s after plant-wide optimization control is switched ON (and over the same period for the ‘greedy’ case). Also listed are the errors of the FLORIS prediction relative to the SOWFA results, and the increase obtained when using optimized yaw relative to the greedy yaw control case in the SOWFA results.



(a) 3x2 wind plant rotated 0° w.r.t. wind direction



(b) 3x2 wind plant rotated 5° w.r.t. wind direction



(c) 3x2 wind plant rotated 10° w.r.t. wind direction

Figure 4.7: Hub-height wind field at 800s simulated time, as calculated by SOWFA for 3 different configurations of the wind plant. Black lines indicate rotor positions and yaw orientation of each turbine (the yaw angles are listed in Table 4.3).

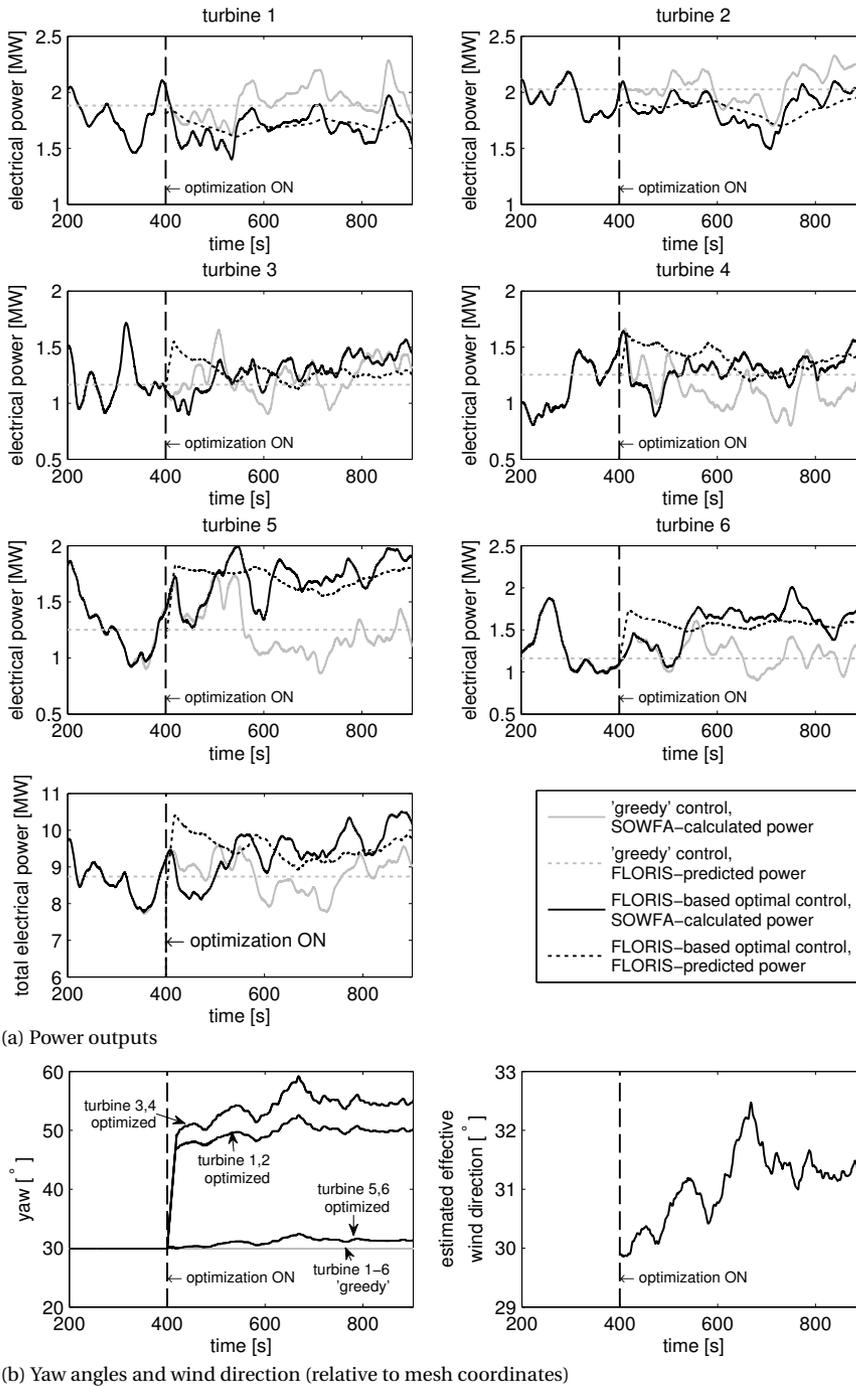


Figure 4.8: SOWFA simulation results and FLORIS predictions of power productions with and without FLORIS-based optimization control (a), and the yaw angles used in these simulations (b), for the 5° rotated wind plant configuration (see also Figure 4.7b).

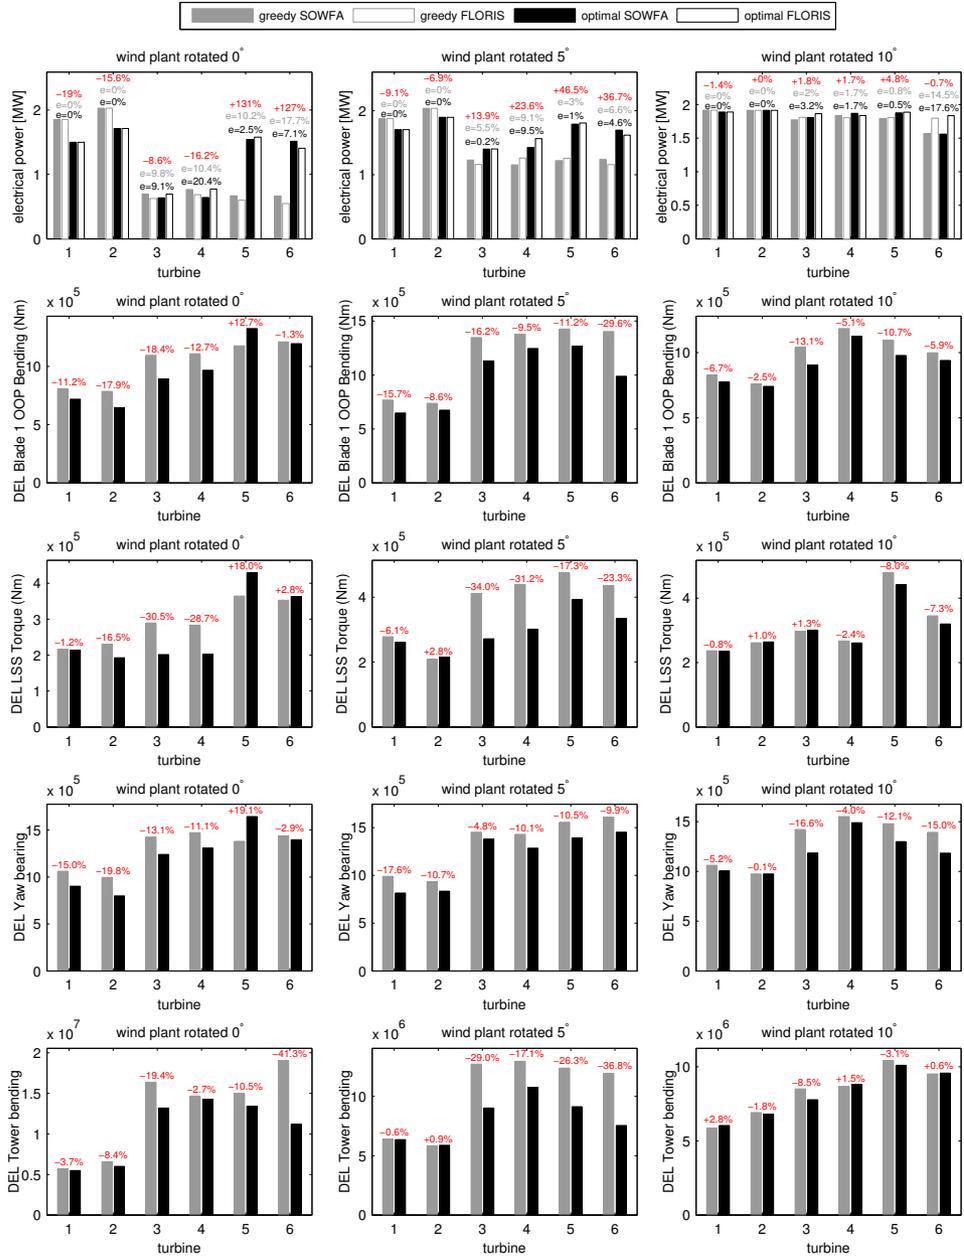


Figure 4.9: Time-averaged SOWFA simulation results (solid bars) for turbine electrical power production and structural loads for both turbine-level optimal ('greedy') and plant-wide optimal yaw settings ('optimal'), for 3 different wind plant configurations (see Figure 4.7 for these configurations). The data is averaged over a period of 50s to 500s after plant-wide optimization control is switched ON (and over the same period for the 'greedy' case). Listed are increases of power and loads of the optimized case relative to the greedy case as calculated by SOWFA (in %, red). Also included are the FLORIS estimates of the power productions (clear bars), and the associated errors ϵ relative to the SOWFA-calculated values (in %, black and gray). Table 4.2 shows total wind plant electrical power production. Table 4.3 lists the yaw angles used in the 'optimal' case.

	turbine 1	turbine 2	turbine 3	turbine 4	turbine 5	turbine 6
wind plant rotated 0°	25.85°	25.15°	39.80°	39.75°	0.45°	0.35°
wind plant rotated 5°	19.00°	19.00°	23.80°	23.80°	0.05°	0°
wind plant rotated 10°	6.25°	3.45°	4.75°	6.35°	-0.05°	0.05°

Table 4.3: Optimized turbine yaw angles (relative to the mean 30° wind direction) after 800s of simulation for different wind plant configurations.

4.5.5. ESTIMATION OF ENERGY GAIN OF YAW-BASED WIND PLANT CONTROL ON THE PRINCESS AMALIA WINDPARK

In Fleming et al. (2014c), the annual energy production gain that can be obtained through yaw-based wake redirection control for the Princess Amalia Windpark offshore wind plant, is estimated to be 1.1%. This estimate is obtained using the same approach as taken in Section 3.5.3: for each wind direction the annual energy gain is evaluated of using optimized yaw instead of greedy control, using the FLORIS model, and then the annual estimated gain is based on the frequency of those wind directions over the year, based on measured wind data from a nearby location in the North Sea NoordzeeWind B.V. (2013); Brand et al. (2013) .

4.6. CONCLUSIONS

In this chapter, a wind plant control strategy for the optimization of the total electrical energy production of the wind plant by changing the yaw control set-points of each turbine, was presented. The yaw control is used to change the direction and velocity of the wake forming behind each turbine in the wind plant. The optimization is based on predictions provided by the FLORIS model, a novel simplified parametric model for these wake effects in the wind plant.

In high-fidelity CFD simulation examples, it was shown that the control strategy could be applied successfully on a small wind plant. The FLORIS model was found to be able to predict the time-averaged turbine powers for both the ‘greedy’ and plant-wide optimized settings with sufficient accuracy to indeed yield a significant power production increase, for different configurations of the wind plant. The CFD simulations also predict that a reduction of loads can be achieved through the yaw control. The CFD simulation examples provide a first proof of concept for the data-driven optimization scheme based on the FLORIS model.

Ongoing work is aimed at the further development of the control scheme such that it can be applied on a real wind plant, under changing atmospheric conditions. When for example the turbulence intensity of the inflow changes, the wake properties are affected, and the model parameters should be updated online. The FLORIS model has a relatively simple formulation, with a small number of parameters that can be identified using power measurements of the different turbines in the wind plant. This enables the development of such a fully data-driven approach for adaptive wind plant optimization control.

In further ongoing work, the FLORIS model is extended with a more accurate descrip-

tion of the wake effects, in order to give better predictions of the wind plant performance. This work includes the extensions proposed by Annoni et al. (2014a). These extensions consist of a better description of wake decay through turbulent mixing in overlapping wakes, and a better prediction of the effects of changing axial induction using blade pitch and generator torque. Further, in the next Chapter 5, the FLORIS model is extended with dynamic effects, by including the delays associated with the wake effects propagating through the wind field.

Certainly when considering varying atmospheric conditions, the FLORIS model needs to be further validated. Gathering validation data from relevant and realistic scenarios with changing wind conditions from high-fidelity models like SOWFA, or with a real wind plant, is a substantial task, given the computational costs of high-fidelity models, and the uncontrollable nature of the conditions in real wind plants. With proper validation of the model under varying wind conditions, also a description of the uncertainty of the FLORIS model should be formulated, with which a robust yaw control strategy can be developed, possibly with some conservatism in the amount of yaw offset used.

5

A CONTROL-ORIENTED DYNAMIC MODEL FOR WAKES IN WIND PLANTS

In this chapter, a novel control-oriented dynamic model for predicting wake effects in wind plants, called the FLOW Redirection and Induction Dynamics (FLORIDYN) model, is presented. The model predicts the wake locations and the effective flow velocities at each turbine, and the resulting wind turbine electrical energy productions, as a function of the control degrees-of-freedom of the turbines (the axial induction and the yaw angle of the different rotors). The model is an extension of the previously presented static model (FLORIS). It includes the dynamic wake propagation effects that cause time delays between control setting changes and the response of downstream turbines. These delays are associated with a mass of air in the wake taking a certain amount of time to travel from one turbine to the next, and the delays are dependent on the spatially- and time-varying state of the wake. The extended model has a state-space structure combined with a non-linear feedback term. While including the control-relevant dynamics of the wind plant, it still has a relatively small amount of parameters. A Kalman filter is developed for the model that corrects the flow field predictions of the model using wind turbine power production measurements. The computational complexity of the model is small enough such that it has the potential to be used for dynamic optimization of the control reference signals for improved wind plant control, as is demonstrated in a case study.

5.1. INTRODUCTION

Because a wind turbine extracts energy from the flow, it has a wake of turbulent flow downstream of its rotor, in which the wind velocity in the wake is reduced with respect to the free-stream velocity. Downstream of the rotor the wake expands, and turbulent mixing and diffusion causes the wake velocity to recover towards the free-stream velocity further downstream. In a cluster of wind turbines (a wind plant), the wake of one turbine

This chapter has been published in Gebraad and van Wingerden (2014a).

can overlap with another turbine rotor, which affects the electrical power production and loads on that turbine. The topology and amount of the wake interaction depends on time-varying atmospheric conditions (e.g., inflow direction, speed and turbulence, and atmospheric stability), and on the control settings of each turbine: the rotor speed and pitch angles of the blades affect the rotor axial induction and thus the wake velocity deficit (Ainslie, 1988), and the rotor yaw angle affects both the velocity deficit and flow direction in the wake (Burton et al., 2002b; Dahlberg et al., 2003; Wagenaar et al., 2012; Adaramola et al., 2011; Jiménez et al., 2010; Fleming et al., 2014b,d).

In Chapter 4, the FLOW Redirection and Induction in Steady-state (FLORIS) model was developed, a simplified control-oriented model that predicts the steady-state characteristics of wakes in a wind plant as a function of the axial inductions and yaw angles of the wind turbine rotors. The complete flow field in a plant does not respond instantaneously to a change in turbine control settings however, since the flow takes some time to move downstream, resulting in a delay of the response of the downstream turbine. In Knudsen et al. (2012), system identification experiments are described with a setup with two megawatt-scale wind turbines aligned in the wind direction in which the axial induction of the front turbine is varied, and the response of the downwind turbine is measured. In these system identification experiments, the wake propagation delay is identified from measured data, as a constant time-delay that is incorporated in the transfer function form of the dynamic model of the interaction between the upstream and downstream turbine. The lengths of the wake propagation delays are dependent on the spatially- and time-varying flow velocity profile in the wake however, therefore, it is useful to parameterize them as such. In an early attempt to incorporate wake velocity dependent delays in a control-oriented model, in Koch et al. (2005), the wake travel time was estimated by dividing turbine-to-turbine distance with the flow speed at the most downstream turbine. Since the velocity deficit in the wake of a turbine is at its maximum just behind the rotor and recovers to the free-stream velocity through mixing further downstream, the delay is underestimated with this method. In Section 3.4 of this thesis we made an extension of the Jensen model (Jensen, 1984; Katić et al., 1986) with a delay model that calculated the wake travel time between two turbines based on the average of the flow velocities at the upstream and the downstream turbine, similar to the approach in Gonzalez et al. (2013). In Choi et al. (2013) a more accurate estimation of the delay time was made by further segmentation of the wake velocity profile in the downstream direction, and summing the delay for each segment. Each of the delay models do not take into account the dynamic response of the wake velocity profile to control setting changes, and do not include the effect of the yaw degree-of-freedom on the wake deflection. Therefore, in this work, we extend the FLORIS model with a simplified dynamic model for the propagation of the effects of control settings changes through the wake. We refer to the model as the FLOW Redirection and Induction Dynamics (FLORIDYN) model.

The outline of this chapter is as follows. In Section 5.2, a description of the FLORIDYN model is given. Then, in Section 5.3 the results of a case study are presented, in which the predictions given by the FLORIDYN model are compared with those of a high-fidelity computational fluid dynamics (CFD) simulation. In Section 5.4, a Kalman filter is developed that makes corrections on the flow field predicted by the FLORIDYN model,

on the basis of data measured at the turbine. In Section 5.5, a case study is provided in which the FLORIDYN model is used to enable plant-level optimized dynamic control. Finally, conclusions and recommendations are given in Section 5.6.

5.2. FLORIDYN MODEL DESCRIPTION

A detailed overview of the FLORIDYN model is provided in the scheme in Figure 5.1 (the gray parts). In brief, the FLORIDYN model is a combination of a static nonlinear mapping describing the wake velocity profile, based on an augmented Jensen model (Jensen, 1984; Katić et al., 1986), and the wake deflection through yaw (based on Jiménez et al. (2010)), extended with a state-space model describing the propagation of control settings changes through the wake.

In this section, the model will be described part-by-part, starting with a model for the turbine power generation from the wind with a certain effective wind speed (Section 5.2.1). Then, in Section 5.2.2 a subdivision is made between front and downstream turbines, because only for the downstream turbines the wake properties are to be estimated to find the effective wind velocities. Further prerequisite definitions follow in Section 5.2.3, where a special Cartesian coordinate system is explained, and in Section 5.2.4, where a subdivision of the wake in different zones is introduced. The wake propagation model is defined in Section 5.2.5, and the model to find the local wake characteristics is given in Section 5.2.6.

An overview of notations that are used throughout the model description in this section, is given in Table 5.1. These notations will also be introduced step-by-step in the next subsections.

5.2.1. TURBINE ELECTRICAL POWER PRODUCTIONS

The index t is used to count the different wind turbines in a wind plant. The index k denotes the discrete time steps. When the effective wind speed at a time k at a turbine t , denoted as $u_T(t, k)$, is known, the electrical power of a turbine t is calculated as:

$$P_T(t, k) = \frac{1}{2} \rho A_T(t) C_P(a_T(t, k), \gamma_T(t, k)) u_T(t, k)^3 \quad (5.1)$$

where ρ is the air density, and $A_T(t)$ is the rotor swept area, and C_P is the power coefficient of the turbine t . In Section 4.3.1 we derived the following heuristic relationship between the axial induction factor a_T and the yaw angle γ_T of the rotor and the power coefficient C_P , based on the work in Medici (2005):

$$C_P(a_T(t, k), \gamma_T(t, k)) = 4a_T(t, k)[1 - a_T(t, k)]^2 \eta \cos(\gamma_T(t, k))^{p_P}. \quad (5.2)$$

Note that a_T in this relation corresponds to the axial induction factor when the rotor is not yawed, which can be found from the blade pitch angle and rotor speed using knowledge of the turbine characteristics. Scalars p_P and η are turbine-specific model parameters: η is a efficiency factor used to represent possible losses with respect to the theoretical maximum efficiency in non-yawed operation, parameter p_P is used to fit the

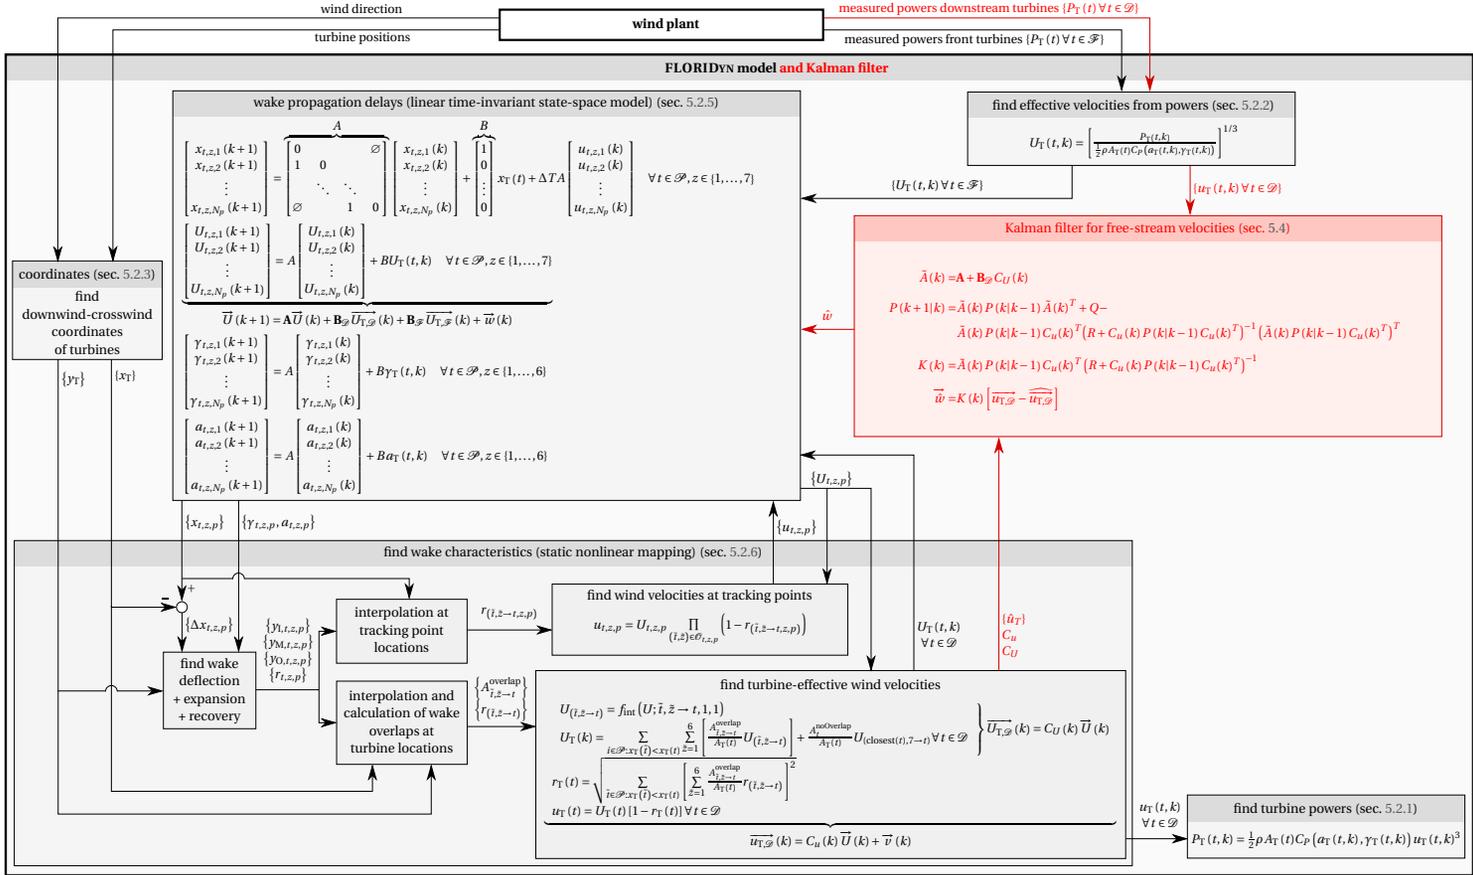


Figure 5.1: FLORIDYN model scheme, with Kalman filter (in red). Measurements from the wind plant (top) are fed into the model to estimate the inflow properties.

Time stepping

k	discrete time-step index
ΔT	sample time, interval between two discrete time steps k and $k + 1$
A, B	state-space matrices, cf. eq. 5.4 (definition), and eq. 5.6-5.7

Turbine and wake indices and sets

t	turbine index
N_t	number of turbines in wind plant
\mathcal{P}	set of all turbines in wind plant
\mathcal{F}	set of front turbines
\mathcal{D}	set of downstream turbines
z	wake zone index
p	wake tracking point (TrP) index
N_p	number of TrPs in a wake zone

Turbine properties

A_T	turbine rotor area
D_T	turbine rotor diameter
x_T	turbine x-position in downwind-crosswind coordinates
y_T	turbine y-position in downwind-crosswind coordinates
P_T	turbine power production
C_P	turbine power coefficient, cf. eq. 5.1
pp	coefficient for adjusting power coefficient C_P to yaw offset, cf. eq. 5.1
η	coefficient for adjusting power coefficient C_P for losses, cf. eq. 5.1
a_T	turbine axial induction (non-yawed)
γ_T	turbine rotor yaw angle
u_T	effective wind speed at turbine
U_T	effective free-stream wind speed at turbine
r_T	effective wind speed reduction factor at turbine

Properties of a Tracking Point (TrP) p in the wake zone z of a turbine t

$x_{t,z,p}$	TrP x-position in downwind-crosswind coordinates
$\Delta x_{t,z,p}$	TrP x-distance to turbine rotor
$y_{c,t,z,p}$	y-position of wake center effective at TrP
$y_{I,t,z,p}$	y-position of wake zone inner-point
$y_{M,t,z,p}$	y-position of wake zone mid-point
$y_{O,t,z,p}$	y-position of wake zone outer-point
$a_{t,z,p}$	axial induction effective at TrP
$\gamma_{t,z,p}$	yaw angle effective at TrP
$U_{t,z,p}$	free-stream speed effective at TrP
$u_{t,z,p}$	wake speed at TrP
$r_{t,z,p}$	wake speed reduction factor effective at TrP
$\mathcal{O}_{t,z,p}$	set of wake zones that is overlapping with a TrP

Other wake zone properties

$D_{w,n}, D_{w,f}, D_{w,m}$	diameters of main wake zones (near, far and mixing zone), cf. Fig. 5.2, eq. 5.10
-----------------------------	--

Flow parameters

ρ	air density
k_d	wake deflection coefficient, cf. eq. 5.8-5.9
α_d, β_d	coefficients for wake deflection caused by rotor rotation, cf. eq. 5.8-5.9
k_e	wake expansion coefficient, cf. eq. 5.8-5.9
$m_{e,n}, m_{e,f}, m_{e,m}$	wake zones expansion coefficients, cf. eq. 5.10
$M_{U,n}, M_{U,f}, M_{U,m}$	wake zone velocity recovery coefficients, cf. eq. 5.12-5.14
α_U, β_U	coefficients for adjusting wake velocity deficit to yaw offset, cf. eq. 5.12-5.14

Table 5.1: Notations repeatedly used in the FLORIDYN model description, Section 5.2

power loss when a turbine is yawed. In Chapter 4, the values $\eta = 0.768$, $p_p = 1.88$ were found to provide a good match with experiments in which the yaw angle of the NREL 5-MW baseline turbine was varied.

In Figure 5.1 the turbine power model is used in the lower-right sub-block of the FLORIDYN model to estimate the turbine electrical power productions from estimated effective wind velocities. This is only done for the set of downstream turbines \mathcal{D} , as is explained in the next section.

5.2.2. FRONT AND DOWNSTREAM TURBINES

The set $\mathcal{P} = \{1, 2, \dots, N_T\}$ denotes the set of indexes t of all turbines in the plant, with N_T denoting the total number of turbines in the plant. Given a certain inflow direction, we can distinguish some front turbines in the wind plant, for which the rotor is not overlapping with the wake of any other upstream turbines. Those front turbines are in the set $\mathcal{F} \subset \mathcal{P}$. From a controls perspective, the velocity of the inflow to these turbines, is a given input (disturbance) to the wind plant system. From measurements at these front turbines (power, yaw and axial induction) we can estimate the local free-stream velocities, denoted by U_T , by inverting relation (5.1):

$$U_T(t, k) = \left[\frac{P_T(t, k)}{\frac{1}{2} \rho A_T(t) C_P(a_T(t, k), \gamma_T(t, k))} \right]^{1/3} \quad \forall t \in \mathcal{F} \quad (5.3)$$

as is also shown in the top-right sub-block of the model scheme in Figure 5.1. The above inversion assumes that the C_P factor remains constant with changing wind speed, which is a valid assumption for below-rated operation with constant tip-speed-ratio, and also above-rated this is expected to be a workable assumption, as long as the C_P (or the axial induction) is updated with the most recent tip-speed-ratio estimate, and the wind speed changes slowly enough. An alternative is to use more advanced methods for estimating effective wind speed from power using filtering techniques based on a dynamic turbine model, such as those presented in Østergaard et al. (2007); Knudsen et al. (2011).

5

Adjustments of the axial induction and yaw of the turbine rotors affect the wake effects on the turbines that are standing downstream of the front turbines, that are included in the set $\mathcal{D} = \{t \in \mathcal{P} | t \notin \mathcal{F}\}$. The inflow velocities for those turbines can be considered as variables that can be controlled to have values within a certain range. In the remainder of this section, it will be described how the effective inflow speeds u_T at each turbine $t \in \mathcal{D}$ are estimated through a dynamic model of the wake characteristics. Based on those estimated effective inflow speeds u_T , the FLORIDYN model estimates the power production of those turbines, using eq. 5.1.

5.2.3. DOWNWIND-CROSSWIND COORDINATE SYSTEM

In order to describe the spatial properties of the wakes in the wind plant, a Cartesian coordinate system (x, y) is adopted, in which the x -axis is pointing downwind along an estimated mean inflow direction in the plant, and the y -axis is pointing orthogonal to the x -axis in the horizontal direction, i.e., along the cross-wind direction (see Figure 5.2). In

this work, we assume that each turbine has the same hub-height, and the turbine locations in this downwind-crosswind coordinate system are denoted as $(x_T(t), y_T(t))$. In the upper-left sub-block of the model scheme in Figure 5.1, the turbine positions are transformed to the downwind-crosswind coordinate system using information of the (free-stream) wind direction. A detailed description of a method to perform this transformation based on turbine measurement data, is provided in Section 4.3.2.

5.2.4. WAKE ZONES

Following the approach of the FLORIS model, the wake is divided in different zones (see Figure 5.2), each with their own expansion and recovery properties, so that the crosswind velocity profile can be fitted more precisely than with the standard Jensen model that has a wake velocity profile that is constant in the crosswind direction. In the FLORIS model there are three zones: the (inner) near wake, the (middle) far wake and the (outer) mixing zone, we will refer to them later on as the ‘main zones’. A difference with the FLORIS model described in Section 4.3, is that these main zones are further divided in a left and right part, and we add a zone to describe the free-stream, resulting in a total of 7 wake zones. This further subdivision is needed because in the FLORIDYN model, in order to calculate the delays, an estimate of the velocity profile over the full wake is needed, rather than just the effective velocities at the turbines, as in the FLORIS model. With partial overlap of wakes, the propagation delays may be different in the left and right part of the wake.

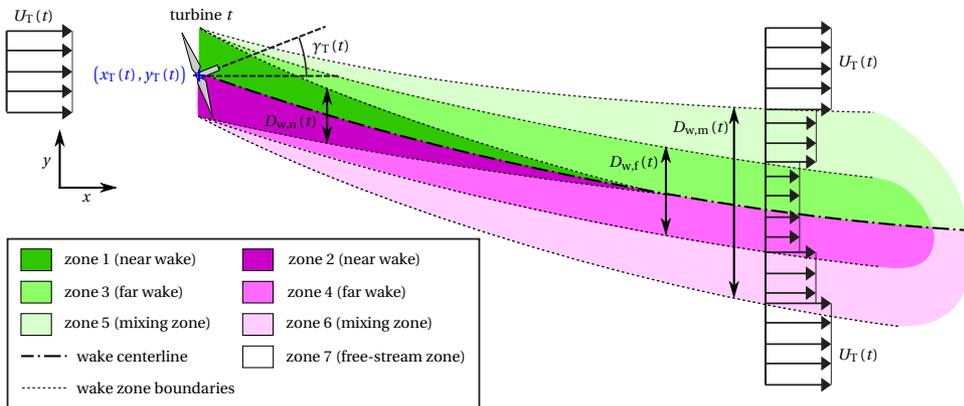


Figure 5.2: The different zones of the wake in the FLORIDYN model

5.2.5. WAKE PROPAGATION DYNAMICS

In this subsection, we present the submodel that describes the delays in the response of the wake velocity profile to changes in the control settings, which is the top-center sub-block in the FLORIDYN model scheme in Figure 5.1. The delay mechanism is implemented by defining in each wake zone of each turbine, a finite number of so-called Tracking Points (TrPs), as shown in Figure 5.3. The first TrP is located in the turbine rotor

plane. It is assumed that at each discrete time step, a mass of air will move from one TrP to the next downstream TrP. The axial distances between the TrPs are adjusted accordingly, based on the (to be estimated) wind velocity in between the TrPs. Variables that are measured at the turbines, are passed on between the TrPs, using the time update laws given below, and then these variables are used to calculate the local wake characteristics, using the submodel described in the next section (lower-left block in Figure 5.1).

The index $z \in \{1, \dots, 7\}$ numbers the zones in a turbine wake. The total number of TrPs in one wake zone of a turbine is N_p . The index $p \in \{1, \dots, N_p\}$ numbers the TrPs in a wake zone. We use a notation in which $x_{t,z,p}(k)$ and $u_{t,z,p}(k)$ are respectively the downwind position and the velocity of a TrP p in a wake zone z of a turbine t at a time-step k . For adjusting the downwind positions of the TrPs in each time-step, it is assumed that the velocity in between two TrPs is constant over the downwind distance, resulting in the following update law:

$$\begin{bmatrix} x_{t,z,1}(k+1) \\ x_{t,z,2}(k+1) \\ \vdots \\ x_{t,z,N_p}(k+1) \end{bmatrix} = \overbrace{\begin{bmatrix} 0 & & & \emptyset \\ 1 & 0 & & \\ & \ddots & \ddots & \\ \emptyset & & 1 & 0 \end{bmatrix}}^A \begin{bmatrix} x_{t,z,1}(k) \\ x_{t,z,2}(k) \\ \vdots \\ x_{t,z,N_p}(k) \end{bmatrix} + \overbrace{\begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}}^B x_T(t) + \Delta T A \begin{bmatrix} u_{t,z,1}(k) \\ u_{t,z,2}(k) \\ \vdots \\ u_{t,z,N_p}(k) \end{bmatrix} \quad (5.4)$$

$\forall t \in \mathcal{D}, z \in \{1, \dots, 7\}$

where ΔT is the sampling time, i.e. the time interval between two discrete time-steps k and $k+1$.

The mechanism of modeling the wake propagation dynamics is illustrated in Figure 5.3. In each time-step k , the following variables that are measured at the turbines or estimated from measurements at the turbines, are 'passed down' the stream from one TrP to the next downstream TrP in a zone: the turbine yaw angle γ_T , and axial induction a_T , and the free-stream velocities U_T . From these quantities, local wake properties (lateral position and velocity) at the TrPs are calculated. This makes that at the first TrP, the effects on the wake of a change in yaw are observed after one time-step, and in the second TrP the effects of the yaw properties two time-steps ago is observed, etc., and likewise for the rotor axial induction and the free-stream velocity. The free-stream velocity, axial induction, and yaw effective at a TrP p in zone z of turbine t at time-step k , are denoted respectively as $U_{t,z,p}(k)$, $a_{t,z,p}(k)$ and $\gamma_{t,z,p}(k)$. The yaw and axial induction property is only passed on in this way in zones 1 to 6, and not in free-stream zone 7, since by definition the free-stream flow is not affected by the axial inductions and yaw angles of turbines. Using a similar vector notation as above, the update laws for passing on the measurements between the TrPs in each time-step are:

$$\begin{bmatrix} U_{t,z,1}(k+1) \\ U_{t,z,2}(k+1) \\ \vdots \\ U_{t,z,N_p}(k+1) \end{bmatrix} = A \begin{bmatrix} U_{t,z,1}(k) \\ U_{t,z,2}(k) \\ \vdots \\ U_{t,z,N_p}(k) \end{bmatrix} + B U_T(t, k) \quad \forall t \in \mathcal{D}, z \in \{1, \dots, 7\} \quad (5.5)$$

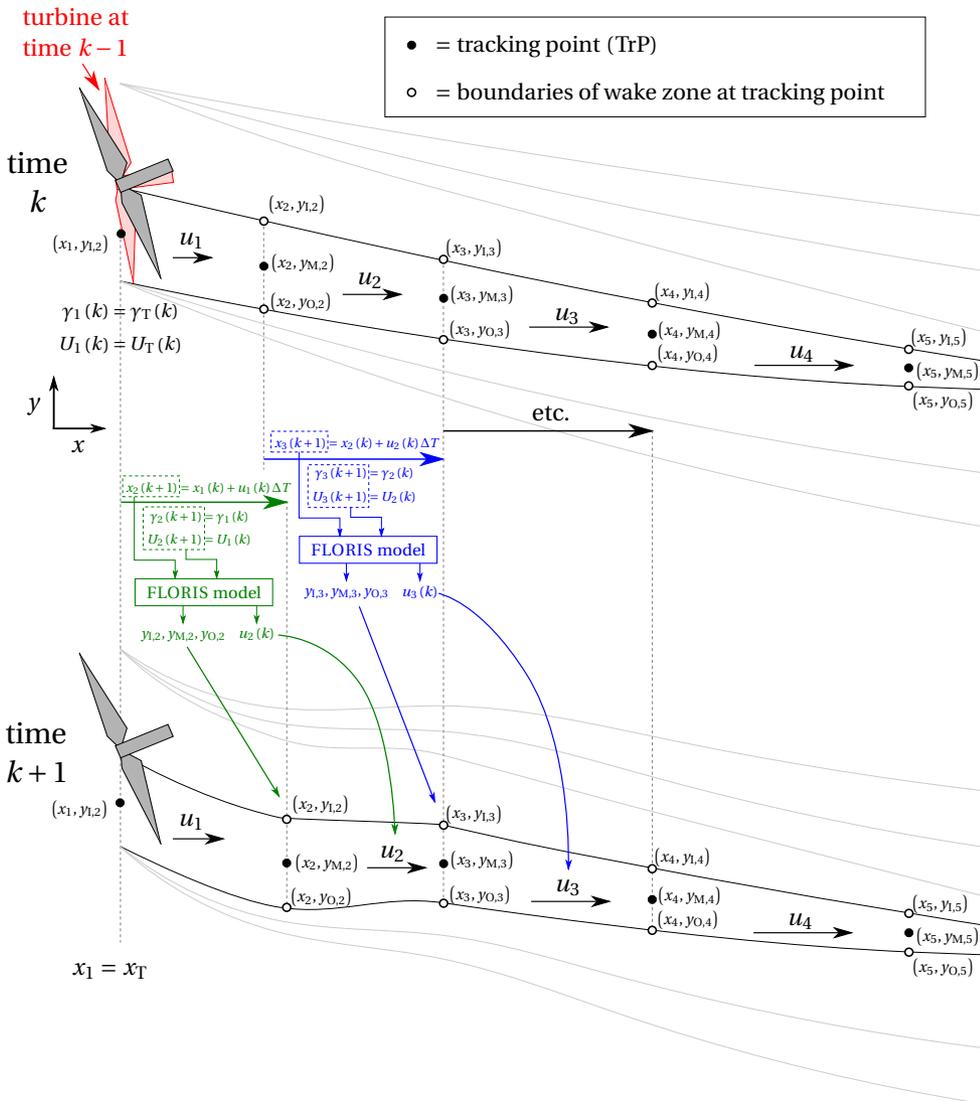


Figure 5.3: Illustration of the state update mechanism in the FLORIDYN model. The turbine changes its yaw angle between time step $k-1$ and k , and we see that the velocities, as well as the x -position, and y -positions of the wake zone mid- and boundary points at the TrP, are updated using an extended FLORIS model. The effect of the yaw change is seen as a wake deflection in the left-most TrPs first. Since we only consider one turbine and wake zone, the t and z indices are omitted in the notation.

$$\begin{bmatrix} \gamma_{t,z,1}(k+1) \\ \gamma_{t,z,2}(k+1) \\ \vdots \\ \gamma_{t,z,N_p}(k+1) \end{bmatrix} = A \begin{bmatrix} \gamma_{t,z,1}(k) \\ \gamma_{t,z,2}(k) \\ \vdots \\ \gamma_{t,z,N_p}(k) \end{bmatrix} + B\gamma_T(t,k) \quad \forall t \in \mathcal{P}, z \in \{1, \dots, 6\} \quad (5.6)$$

$$\begin{bmatrix} a_{t,z,1}(k+1) \\ a_{t,z,2}(k+1) \\ \vdots \\ a_{t,z,N_p}(k+1) \end{bmatrix} = A \begin{bmatrix} a_{t,z,1}(k) \\ a_{t,z,2}(k) \\ \vdots \\ a_{t,z,N_p}(k) \end{bmatrix} + Ba_T(t,k) \quad \forall t \in \mathcal{P}, z \in \{1, \dots, 6\} \quad (5.7)$$

with A and B as in equation (5.4).

5.2.6. CALCULATION OF WAKE CHARACTERISTICS

In this section, it is explained how in the FLORIDYN model, certain control-relevant wake characteristics are estimated at each time-step. This part of the FLORIDYN model is shown in the lower-left block in the FLORIDYN model scheme in Figure 5.1. The several sub-blocks of this part of the model are each explained in subsections. First, in order to find the dimensions and positions of the wake zones, the deflection and expansion properties are estimated (Section 5.2.6.1). Then, the velocity profile in the wake is estimated in several steps (Section 5.2.6.2), which results in the estimated velocity $u_{t,z,p}$ at each TrP, which is an input to the wake propagation model (cf. eq. 5.4). Finally, from the wake characteristics, the effective velocities at the downstream turbines are estimated, cf. Section 5.2.6.3.

Each of the blocks is based on the static FLORIS model previously presented in Chapter 4. Therefore, all the relations in this part of the model are static (nonlinear) mappings, i.e. the estimated wake properties at a TrP at a time k are calculated from the locally effective input variables ($U_{t,z,p}$, $a_{t,z,p}$ and $\gamma_{t,z,p}$) at k only. Therefore, in this section we will omit the time index k in the notations.

5

5.2.6.1. DEFLECTION AND EXPANSION OF THE WAKE ZONES

Through eq. 5.4, each time-step the x -positions of the TrPs are updated. From the locally effective yaw angle $\gamma_{t,z,p}$ and axial induction $a_{t,z,p}$ at a TrP p in wake zone z of turbine t , the model then calculates the spatial lay-out of the wake zones at that TrP in terms of its deflection and expansion relative to the turbine rotor disk. The wake deflection is described as the cross-wind translation of a wake zone centerline, and then the expansion is described relative to that centerline. From this, the y -positions of two wake zone boundary points and a wake zone mid-point are defined at each TrP x -location.

Wake centerline deflection The crosswind position of the centerpoint of the total wake at TrP p , denoted as $y_{C,t,z,p}$, is calculated as:

$$y_{C,t,z,p} = y_T(t) + \Delta y_{w,rotation}(\Delta x_{t,z,p}) + \Delta y_{w,yaw}(\Delta x_{t,z,p}, \gamma_{t,z,p}, a_{t,z,p}, D_T(t)) \quad (5.8)$$

where $y_T(t)$ is the cross-wind position of the turbine t , and where the following functions define the rotor rotation-induced and yaw-induced wake deflection, as a function of $\Delta x_{t,z,p} = x_{t,z,p} - x_T(t)$, the down-wind distance from the turbine, and $D_T(t)$ the rotor diameter of the turbine:

$$\begin{aligned} \Delta y_{w,\text{rotation}}(\Delta x) &= \alpha_d + \beta_d \Delta x \\ \Delta y_{w,\text{yaw}}(\Delta x, \gamma, a, D_T) &= \frac{\xi_{\text{init}}(a, \gamma) \left[15 \left[\frac{2k_d \Delta x}{D_T} + 1 \right]^4 + \xi_{\text{init}}(a, \gamma)^2 \right]}{\frac{30k_d}{D_T} \left[\frac{2k_d \Delta x}{D_T} + 1 \right]^5} - \frac{\xi_{\text{init}}(a, \gamma) D_T \left[15 + \xi_{\text{init}}(a, \gamma)^2 \right]}{30k_d} \quad (5.9) \\ \text{where } \xi_{\text{init}}(a, \gamma) &= \frac{1}{2} \cos^2(\gamma) \sin(\gamma) [4a[1-a]] \end{aligned}$$

with coefficients α_d , β_d and k_d as model parameters. We refer to the FLORIS model description in the previous chapter, Section 4.3.3, for the derivation of the above functions.

Main wake zone expansion The diameters of the three main wake zones (near wake, far wake and the mixing zone), are denoted as respectively $D_{w,n}$, $D_{w,f}$, $D_{w,m}$, cf. Figure 5.2. Similar to the FLORIS model (cf. Section 4.3.4), the zone diameter at a TrP with a downstream distance from the turbine $\Delta x_{t,z,p}$, is given by:

$$D_{w,\bullet}(\Delta x_{t,z,p}) = \max(D_T(t) + 2k_e m_{e,\bullet} \Delta x_{t,z,p}, 0) \text{ with } \bullet = n, f, \text{ or } m \quad (5.10)$$

where parameters $m_{e,n}$, $m_{e,f}$, $m_{e,m}$, k_e are coefficients defining the expansion of the main zones, relative to the total wake expansion rate k_e . As in the FLORIS model, the parameter $m_{e,n}$ is typically set to a negative value, such that the near wake zone contracts over distance, and the other expansion factors are set $0 < m_{e,f} < 1$ and $m_{e,m} = 1$, such that the far wake zone expands, and k_e defines the expansion of the outer mixing zone.

Wake zone boundary- and mid-points Using the main wake zone diameters at each TrP, the model finds the local position of an inner boundary point $(x_{t,z,p}, y_{I,t,z,p})$, an outer boundary point $(x_{t,z,p}, y_{O,t,z,p})$, and a mid-point $(x_{t,z,p}, y_{M,t,z,p})$, of the wake zones 1 to 6 (cf. Figure 5.3), as follows:

$$\begin{aligned} y_{I,t,1,p} &= y_{C,t,1,p}, & y_{O,t,1,p} &= y_{C,t,1,p} + \frac{1}{2} D_{w,n}(\Delta x_{t,z,p}), \\ y_{I,t,2,p} &= y_{C,t,2,p}, & y_{O,t,2,p} &= y_{C,t,2,p} - \frac{1}{2} D_{w,n}(\Delta x_{t,z,p}), \\ y_{I,t,3,p} &= y_{C,t,3,p} + \frac{1}{2} D_{w,n}(\Delta x_{t,z,p}), & y_{O,t,3,p} &= y_{C,t,3,p} + \frac{1}{2} D_{w,f}(\Delta x_{t,z,p}), \\ y_{I,t,4,p} &= y_{C,t,4,p} - \frac{1}{2} D_{w,n}(\Delta x_{t,z,p}), & y_{O,t,4,p} &= y_{C,t,4,p} - \frac{1}{2} D_{w,f}(\Delta x_{t,z,p}), \\ y_{I,t,5,p} &= y_{C,t,5,p} + \frac{1}{2} D_{w,f}(\Delta x_{t,z,p}), & y_{O,t,5,p} &= y_{C,t,5,p} + \frac{1}{2} D_{w,m}(\Delta x_{t,z,p}), \\ y_{I,t,6,p} &= y_{C,t,6,p} - \frac{1}{2} D_{w,f}(\Delta x_{t,z,p}), & y_{O,t,6,p} &= y_{C,t,6,p} - \frac{1}{2} D_{w,m}(\Delta x_{t,z,p}), \\ y_{M,t,z,p} &= \frac{1}{2} (y_{I,t,z,p} + y_{O,t,z,p}) \text{ for all } z \in \{1, \dots, 6\} \end{aligned} \quad (5.11)$$

These different points can then be used to test how the wake zones overlap at different TrPs, which is relevant for estimating the wake velocities, as described in the next subsection.

5.2.6.2. WAKE VELOCITY PROFILE

Interaction between the turbines and their wakes is modeled by correcting the velocities in the wake by means of a method that is again based on an extension of the FLORIS

model. If a TrP is situated in the wake zone of a turbine, the estimated effective wind speed $u_{t,z,p}$ at that TrP is found by correcting the free-stream velocity ‘passed on’ to that TrP, $U_{t,z,p}$, with a wake-induced reduction factor that is dependent on the ‘delayed’ turbine measurements and the downwind distance to the turbine. First, we consider the case of a single turbine influencing its own wake only, then we model the effect of overlapping wakes.

Velocities in a single wake In the FLORIS model, the velocity deficit in the wake zones decayed quadratically with the expansion of the wake (cf. Section 4.3.5), as in the Jensen model (Jensen, 1984; Katić et al., 1986). Because in the FLORIDYN model also the velocity profile closer to the rotor is relevant (in order to calculate the propagation delays), we extend this description with an arctangent function as a velocity correction factor. This factor models the gradual reduction of velocity as the flow passes the rotor, and it is based on the method in Torres et al. (2010) and the velocity profiles found in Annoni et al. (2014b). The factor goes to 1 further downstream. This leads to the following formulation for the velocity at a TrP p in the wake zone z of a turbine t :

$$r_{t,z,p} = 2a_{t,z,p} \overbrace{\left[\frac{1}{2} + \tan^{-1} \left(2 \frac{\Delta x_{t,z,p}}{\pi D_T(t)} \right) \right]}^{\text{near-rotor correction factor}} \left[\frac{D_T(t)}{D_T(t) + 2k_e m_{U,z}(\gamma_{t,z,p}) \Delta x_{t,z,p}} \right]^2 \quad (5.12)$$

$$u_{t,z,p} = U_{t,z,p} (1 - r_{t,z,p}) \quad (5.13)$$

The coefficients $m_{U,z}$ define how quickly the velocities recover to the free-stream velocity $U_{t,z,p}$ in different zones, as the distance to the rotor $\Delta x_{t,z,p}$ increases. Similar to the FLORIS model, these coefficients are adjusted for the rotor yaw angle as follows:

$$m_{U,z}(\gamma_{t,z,p}) = \frac{M_{U,z}}{\cos(\alpha_U + \beta_U \gamma_{t,z,p})} \quad (5.14)$$

where we set $M_{U,1} = M_{U,2} = M_{U,n}$, $M_{U,3} = M_{U,4} = M_{U,f}$ and $M_{U,5} = M_{U,6} = M_{U,m}$, so that the parameters $M_{U,n}$, $M_{U,f}$, $M_{U,m}$ define the wake recovery rate for the main wake zones (near, far, and mixing zone), relative to the k_e factor, and parameters α_U and β_U define how the recovery rates are adjusted for yaw offsets.

Combining wakes to find velocities in overlapping wakes For all the TrPs p in a wake zone z of a certain turbine t , we can calculate the velocity reduction factors induced by that turbine t , $r_{t,z,p}$. To consider the case in which wakes of multiple turbines overlap with a certain TrP, we have to combine the reduction factors of several turbines. Because the TrPs in one wake might be located in different downwind positions x than in another wake, we apply an interpolation to find all the effective reduction factors of the different turbine wakes overlapping with a certain TrP. The effective (interpolated) reduction factor at a TrP p in a wake zone z of a turbine t , of a different wake zone \tilde{z} of a different turbine \tilde{t} , is denoted as $r_{(\tilde{t}, \tilde{z} \rightarrow t, z, p)}$. We find $r_{(\tilde{t}, \tilde{z} \rightarrow t, z, p)}$ by interpolating the reduction factor of the two nearest TrPs (in the upwind and downwind direction) in the wake zone \tilde{z} of turbine \tilde{t} (cf. Figure 5.4). We use the notation $\text{upstr}(\tilde{t}, \tilde{z} \rightarrow t, z, p)$, for the index of the

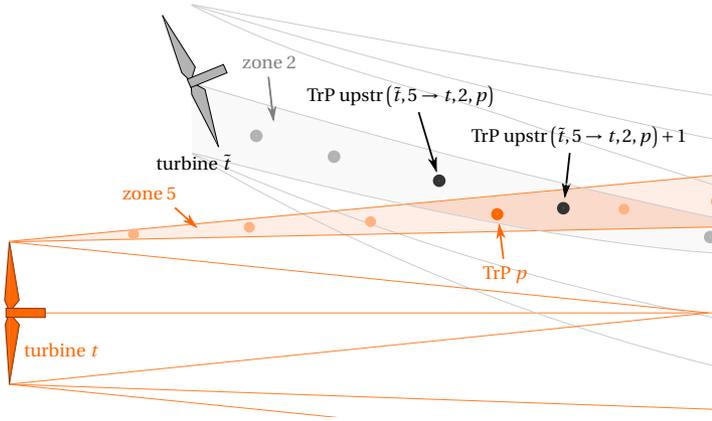


Figure 5.4: Wake zone mid-points of a tracking point (TrP) p , in zone 5 of a turbine t , and the upstream and downstream TrPs in an overlapping wake zone 2 of a turbine \tilde{t} , illustrating the interpolation mechanism in determining the velocities in overlapping wakes.

nearest upwind TrP belonging to the wake zone \tilde{z} of turbine \tilde{t} . The nearest downstream TrP in zone \tilde{z} of turbine \tilde{t} then is $\text{upstr}(\tilde{t}, \tilde{z} \rightarrow t, z, p) + 1$. This results in the following interpolation:

$$r(\tilde{t}, \tilde{z} \rightarrow t, z, p) = f_{\text{int}}(r; \tilde{t}, \tilde{z} \rightarrow t, z, p) \quad (5.15)$$

where f_{int} is a linear interpolation operator, which for some variable ξ defined at each TrP is given by:

$$f_{\text{int}}(\xi; \tilde{t}, \tilde{z} \rightarrow t, z, p) = \frac{x_{\tilde{t}, \tilde{z}, \text{upstr}(\tilde{t}, \tilde{z} \rightarrow t, z, p) + 1} - x_{t, z, p}}{x_{\tilde{t}, \tilde{z}, \text{upstr}(\tilde{t}, \tilde{z} \rightarrow t, z, p) + 1} - x_{\tilde{t}, \tilde{z}, \text{upstr}(\tilde{t}, \tilde{z} \rightarrow t, z, p)}} \xi_{\tilde{t}, \tilde{z}, \text{upstr}(\tilde{t}, \tilde{z} \rightarrow t, z, p)} + \dots \quad (5.16)$$

$$\frac{x_{t, z, p} - x_{\tilde{t}, \tilde{z}, \text{upstr}(\tilde{t}, \tilde{z} \rightarrow t, z, p)}}{x_{\tilde{t}, \tilde{z}, \text{upstr}(\tilde{t}, \tilde{z} \rightarrow t, z, p) + 1} - x_{\tilde{t}, \tilde{z}, \text{upstr}(\tilde{t}, \tilde{z} \rightarrow t, z, p)}} \xi_{\tilde{t}, \tilde{z}, \text{upstr}(\tilde{t}, \tilde{z} \rightarrow t, z, p) + 1}$$

In a similar way, we interpolate the crosswind positions of the wake zone boundaries ($y_{1, \tilde{t}, \tilde{z}, p}$ and $y_{0, \tilde{t}, \tilde{z}, p}$), and use those boundary positions to find out for each TrP, which wake zones of which turbines are overlapping with the wake zone mid-point $y_{M, t, z, p}$ of at that TrP. For brevity, here we omit the exact conditions which follow from simple geometry, and state that if a TrP p belonging to the wake zone z of turbine t is in the wake zone \tilde{z} of a turbine \tilde{t} , then the pair (\tilde{t}, \tilde{z}) belongs to the set $\mathcal{O}_{t, z, p}$. We can combine the effective reduction factors of each wake to find the velocity at a TrP, through:

$$u_{t, z, p} = U_{t, z, p} \prod_{(\tilde{t}, \tilde{z}) \in \mathcal{O}_{t, z, p}} (1 - r(\tilde{t}, \tilde{z} \rightarrow t, z, p)) \quad (5.17)$$

Further, we note that because the overlap conditions are checked at the wake zone mid-point $(x_{t, z, p}, y_{M, t, z, p})$ of each TrP, and because for simplicity it is assumed that the velocity $u_{t, z, p}$ determined at that mid-point, then is constant across crosswind direction, such that the boundary points $(x_{t, z, p}, y_{1, t, z, p})$ and $(x_{t, z, p}, y_{0, t, z, p})$ move parallel to the mid-point (cf. Figure 5.3).

5.2.6.3. CALCULATION OF EFFECTIVE VELOCITIES AT DOWNSTREAM TURBINES

Through the above modeling steps, a two-dimensional description of the velocity profile at the hub-height of the turbines is obtained, that is used to calculate the propagation of the wake zones through the wind field. This description includes the location of the wake zones at the hub-height. For the calculation of the effective wind velocity at the turbines, the overlap method of the FLORIS model (Section 4.3.6) is adopted, that is based on the cross-sectional profile of the wakes at the turbine rotors, rather than on the hub-height profile. Similar to the modeling steps above, the effective wind speeds at the rotor are estimated by combining the effect of the wake zones on the free-stream velocity. First the interpolation function f_{int} is used to find the set of delayed free-stream velocities U for each wake zone \tilde{z} of the turbines \tilde{t} upstream of turbine t , at the location of the rotor plane of t . Since the TrPs with index $p = 1$ are always located in the rotor plane, we can find these interpolated velocities, denoted as $U_{(\tilde{t}, \tilde{z} \rightarrow t)}$, as follows:

$$U_{(\tilde{t}, \tilde{z} \rightarrow t)} = f_{\text{int}}(U; \tilde{t}, \tilde{z} \rightarrow t, 1, 1) \quad (5.18)$$

To find the effective free-stream velocity for each turbine $t \in \mathcal{D}$ we weigh each of the delayed free-stream velocities by the overlap area of the corresponding wake zones \tilde{z} of other turbines \tilde{t} with the rotor of turbine t , denoted by $A_{\tilde{t}, \tilde{z} \rightarrow t}^{\text{overlap}}$ (see Figure 5.5). For the part of the rotor that is not overlapping with any wake (with area $A_t^{\text{noOverlap}}$), we use the delayed free-stream velocity in zone 7 of the upstream turbine that is closest to turbine t in terms of its y -location:

$$\text{closest}(t) = \arg \min_{\tilde{t} \in \mathcal{D}: x_T(\tilde{t}) < x_T(t)} (|y_T(t) - y_T(\tilde{t})|) \quad (5.19)$$

$$U_T(k) = \sum_{i \in \mathcal{D}: x_T(\tilde{i}) < x_T(t)} \sum_{\tilde{z}=1}^6 \left[\frac{A_{\tilde{i}, \tilde{z} \rightarrow t}^{\text{overlap}}}{A_T(t)} U_{(\tilde{i}, \tilde{z} \rightarrow t)} \right] + \frac{A_t^{\text{noOverlap}}}{A_T(t)} U_{(\text{closest}(t), 7 \rightarrow t)} \quad \forall t \in \mathcal{D} \quad (5.20)$$

5

Note that the free-stream speeds U_T for front turbines $t \in \mathcal{F}$ are estimated based on their own measurements (see Section 5.2.2), while the free-stream speeds for turbines $t \in \mathcal{D}$ estimated through equation (5.20) are fed back to the state-space delay model (eq. (5.5)), such that the free-stream wind speed measurements are essentially passed on from turbine to turbine through the wakes of the different turbines. Using the interpolation function f_{int} and a root-sum-square weighting by overlap area, we also can find the effective velocity reduction factor $r_T(t)$ for each turbine rotor:

$$r_T(t) = \sqrt{\sum_{\tilde{t} \in \mathcal{D}: x_T(\tilde{t}) < x_T(t)} \left[\sum_{\tilde{z}=1}^6 \frac{A_{\tilde{t}, \tilde{z} \rightarrow t}^{\text{overlap}}}{A_T(t)} r_{(\tilde{t}, \tilde{z} \rightarrow t)} \right]^2} \quad (5.21)$$

with $r_{(\tilde{t}, \tilde{z} \rightarrow t)} = f_{\text{int}}(r; \tilde{t}, \tilde{z} \rightarrow t, 1, 1)$. We then correct the free-stream velocity to find the effective velocity for the turbine:

$$u_T(t) = U_T(t) [1 - r_T(t)] \quad \forall t \in \mathcal{D} \quad (5.22)$$

The effective velocities $u_T(t)$ can then be used to calculate the powers of through equation (5.1).

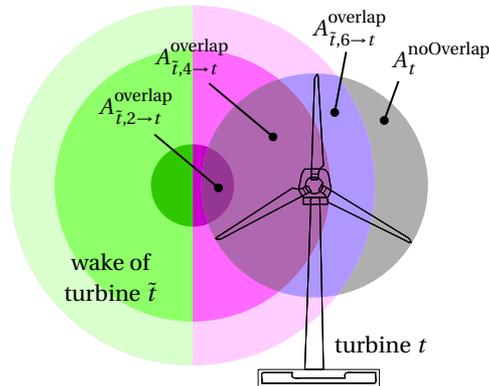


Figure 5.5: Cut-through of wake at a turbine rotor plane, showing the different overlap areas. The colors of the overlapping wake relate to the different wake zones as shown in Figure 5.2.

We have now described each part of the FLORIDYN model (gray parts in Figure 5.1). The model uses the measured electrical power production of the front turbines, in order to estimate the free-stream inflow velocity to the wind plant, and then uses control settings of each turbine, to estimate the hub-height flow field propagating through the plant, and the powers at the downstream turbines. In the next Section 5.3, we will perform a case study in which we compare the FLORIDYN model prediction with SOWFA simulation results, for a particular set of tuned model parameter values. Then, in Section 5.4, we will apply a Kalman filter to the FLORIDYN model (red parts in Figure 5.1) that also uses the power of the downstream wind turbines in order to correct the flow field.

5.3. SIMULATION CASE STUDY

Figure 5.6 and 5.7 show the results of a case study in which we simulate a small wind plant in three different ways:

- a high-fidelity 3D large-eddy simulation with the Simulator for On/Offshore Wind Farm Applications (SOWFA) (cf. Section 1.2.4.1, and Churchfield et al. (2012a); Fleming et al. (2013b))
- a simulation of the same wind plant with the FLORIDYN model, with the parameters of the model tuned to provide a good match with the SOWFA data (the model parameters are listed in Table 5.2)
- calculation of steady-state (SS) solution of the FLORIDYN model, which would be comparable with the predictions that the static FLORIS model would provide

The simulated wind plant consists of 2 rows with 3 NREL 5-MW baseline turbines (Jonkman et al., 2009) each, with a 5 rotor diameter (5D) spacing in the row direction, and 3D in the column direction. The turbine rows are rotated 5 degrees with respect to the wind direction. In the SOWFA simulations, an inflow with a 6% turbulence intensity and

ρ	p_p	η	N_p	ΔT	α_d	β_d	k_d	k_e
1.172	1.88	0.768	80	5 s	-4.5	0.01	0.15	0.0963
$m_{e,n}$	$m_{e,f}$	$m_{e,m}$	$M_{U,n}$	$M_{U,f}$	$M_{U,m}$	α_U	β_U	
-0.5	1	2.2	0.375	1	5.125	5°	1.66	

Table 5.2: FLORIDYN model parameter values used in the case study of Section 5.3

an 8 m/s mean velocity is used (see the case study in Section 4.5 for more details on the setup). After 400 s, the yaw angles of upstream turbines are misaligned with the incoming wind direction, in order to redirect the wakes away from the downstream turbines. The first 300 s of the simulation results are not shown in Figure 5.6 and 5.7, since in the time interval 0-300 s in the SOWFA simulation the wakes are developing, which is not representative of normal operation in a wind plant.

In the tuning of the FLORIDYN model, the time interval ΔT , was set to 5 s (a smaller time interval just increases calculation time without improving the fit with the SOWFA data). Then, the total number of TrPs was set at $N_p = 80$, large enough such that the last TrP in each wake zone is behind the last turbine in the row at all times during the simulation (note that the strings of TrPs get shorter if the velocity in the wake reduces). Further, the tuning parameters that are directly related to the turbine properties (η , p_p , α_U , β_U) and the wake deflection due to rotation (α_d and k_d), as well as the expansion coefficients used to subdivide the wake in wake zones $m_{e,n}$, $m_{e,f}$, $m_{e,m}$, were set to the same values as for the FLORIS model, because there is no clear reason that they should change in a model taking into account the dynamics of wake propagation, and because this also simplifies the tuning problem. The parameters related to the recovery of the wake velocity profile (k_e , $M_{U,n}$, $M_{U,f}$, $M_{U,m}$) were changed however, in order to find a better match between FLORIDYN and SOWFA. These parameters were tuned using a pattern search optimization algorithm (MathWorks, 2014) that minimized the total root-mean-square error between the turbine power signals, resulting in the settings listed in Table 5.2. The pattern search algorithm showed to be a robust way of optimizing of the parameters, but alternative methods of optimization may provide faster convergence.

5

In Figure 5.6, it can be seen that:

- The FLORIDYN model gives an overall better fit to the SOWFA-predicted turbine electrical power production data of the downstream turbines than the SS solution, with the exception of turbine 6. Specifically, in the 400-600 s time range, when the effects of the yaw angle changes propagate through the wind field, the FLORIDYN model will perform better. In Figure 5.7, it can be seen both in the SOWFA- and in the FLORIDYN-predicted flow-fields, that in response to the yaw change, first the part of the wake closer to the rotor will be redirected and then gradually the part further downstream. The SS solution does not include these transient effects.
- There are still significant errors between the FLORIDYN-predicted power output and the SOWFA-predicted output. This can at least partly be attributed to turbulence (and the resulting wake meandering of the wake), which is not fully represented by the FLORIDYN model. The mismatch is larger on the turbines further

model	type of software implementation	type of hardware	computation time
SOWFA	OpenFOAM	cluster, 512 CPUs	59h (distributed)
FLORIDYN	MATLAB code, partly compiled	single CPU on PC (1.6 GHz)	1.4s
FLORIS	MATLAB code	single CPU on PC (1.6 GHz)	0.6ms
FLORIS	compiled C code	single CPU on cluster	0.04ms

Table 5.3: Computational costs of simulating the 3x2 wind plant in the case study of Section 5.3, for different models.

downstream (5 and 6), where turbulence is more prevalent because of the wake overlap with two turbines.

These notions are also shown in the normalized root-mean square errors (RMSE) listed for each turbine in Figure 5.6. The normalized RMSE is defined as the root-mean square error between SOWFA-predicted power output and FLORIDYN-predicted power output, divided by the mean SOWFA-predicted power output.

As shown in Table 5.3, the computational complexity of the FLORIDYN model is increased significantly when compared to FLORIS because of the additional interpolation steps that are needed to describe the interaction effects between the wakes of different turbines that are propagating through the wind field, but the computational complexity is very small when compared to SOWFA. It should be noted that the FLORIDYN model used here consists of MATLAB code that is only partly compiled, and the computational efficiency could certainly be increased just by improving the software implementation (e.g. writing FLORIDYN in C-code).

5.4. USING A KALMAN FILTER TO CORRECT THE PREDICTED VELOCITY FIELD

The FLORIDYN model describes the velocity profile in the wakes of turbines. Although the ensemble-average effect of turbulence¹ on the wake recovery is included through a parametric model of the wake velocity profile, the local fluctuations as a result of turbulence are not included, and it is beyond the scope of this work (the development of a simplified, control-oriented dynamic wind plant model) to include a detailed turbulence model. In the case study in Section 5.3, we have seen that there is error between the FLORIDYN-predicted output, and the SOWFA-predicted output, which to some extent can be attributed to turbulence in the flow field. In control engineering, Kalman filters are widely used to reconstruct the state of models from measured data, under the influence of random noise influences, and model inaccuracies (Verhaegen et al., 2007). In this section, it is shown that in order to correct the errors between FLORIDYN and SOWFA, we can use Kalman filtering techniques to reconstruct the state of the FLORIDYN model, from measurements at the turbine (electrical power production and control settings). Since these are measurements available on a wind turbine, it is expected that

¹In turbulent flows generated under the same conditions, the basic behavior in terms of statistical properties will not change between repetitions of an experiment. The ensemble average is taken across the experiment repetitions to get the underlying mean value. Under the ergodic assumption, the ensemble average is the same as the time-average over a sufficiently long experiment. See also Carley (2011).

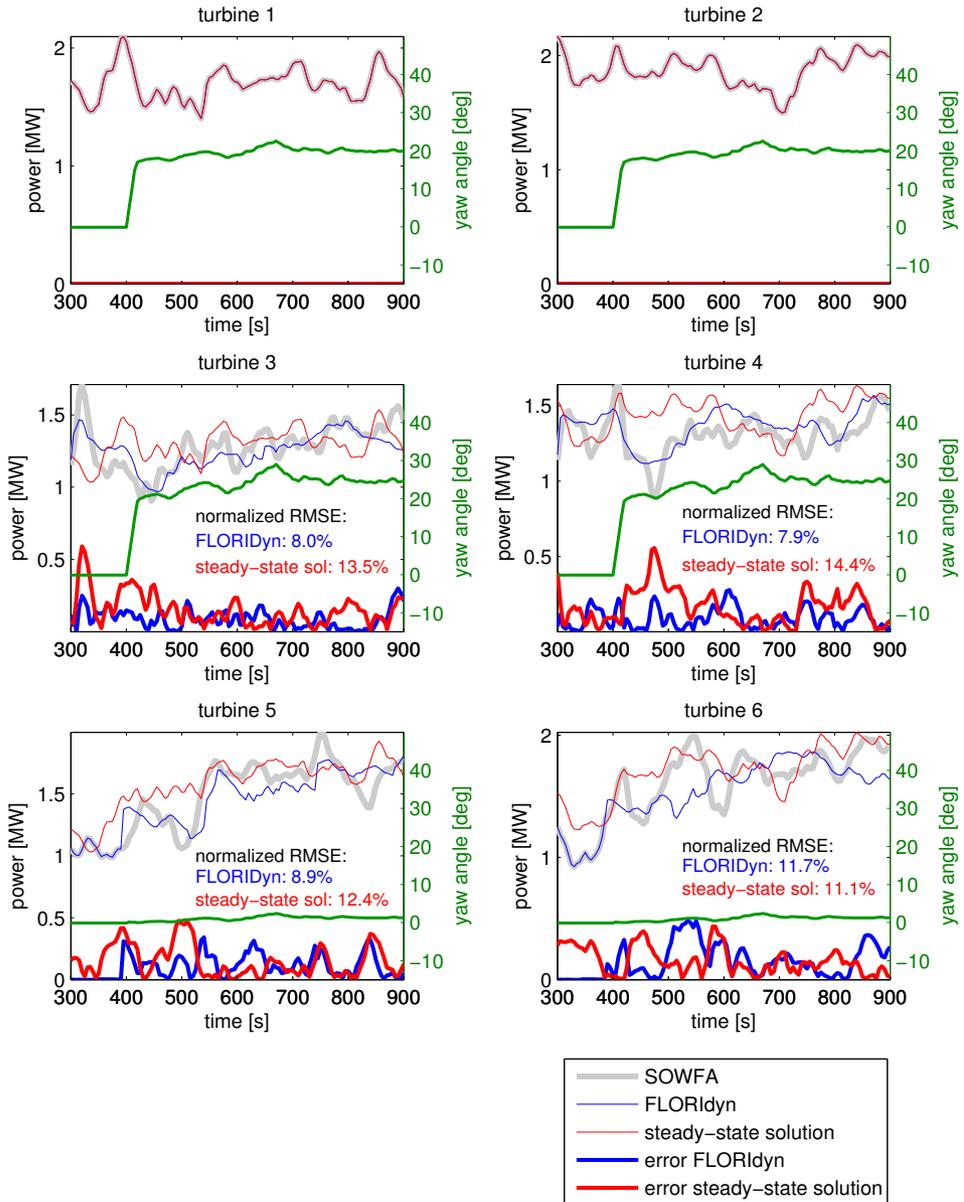


Figure 5.6: Predicted power productions and yaw angles in the case study of Section 5.3. After 400s, the yaw angles of the 4 upstream turbine are changed. Compared to the steady-state (SS) solutions, the dynamic FLORIDYN model is able to better predict the characteristics of the power production response of the downstream turbines in SOWFA (especially in the 400-600s time range), as is also seen in the root-mean-square error (RMSE) between the responses. This is caused by the fact that the wake effects have to propagate through the wind field (see also Figure 5.7), resulting in a delayed increase in power after the yaw offsets are applied. These delays are included in the FLORIDYN model, but not in the SS solutions.

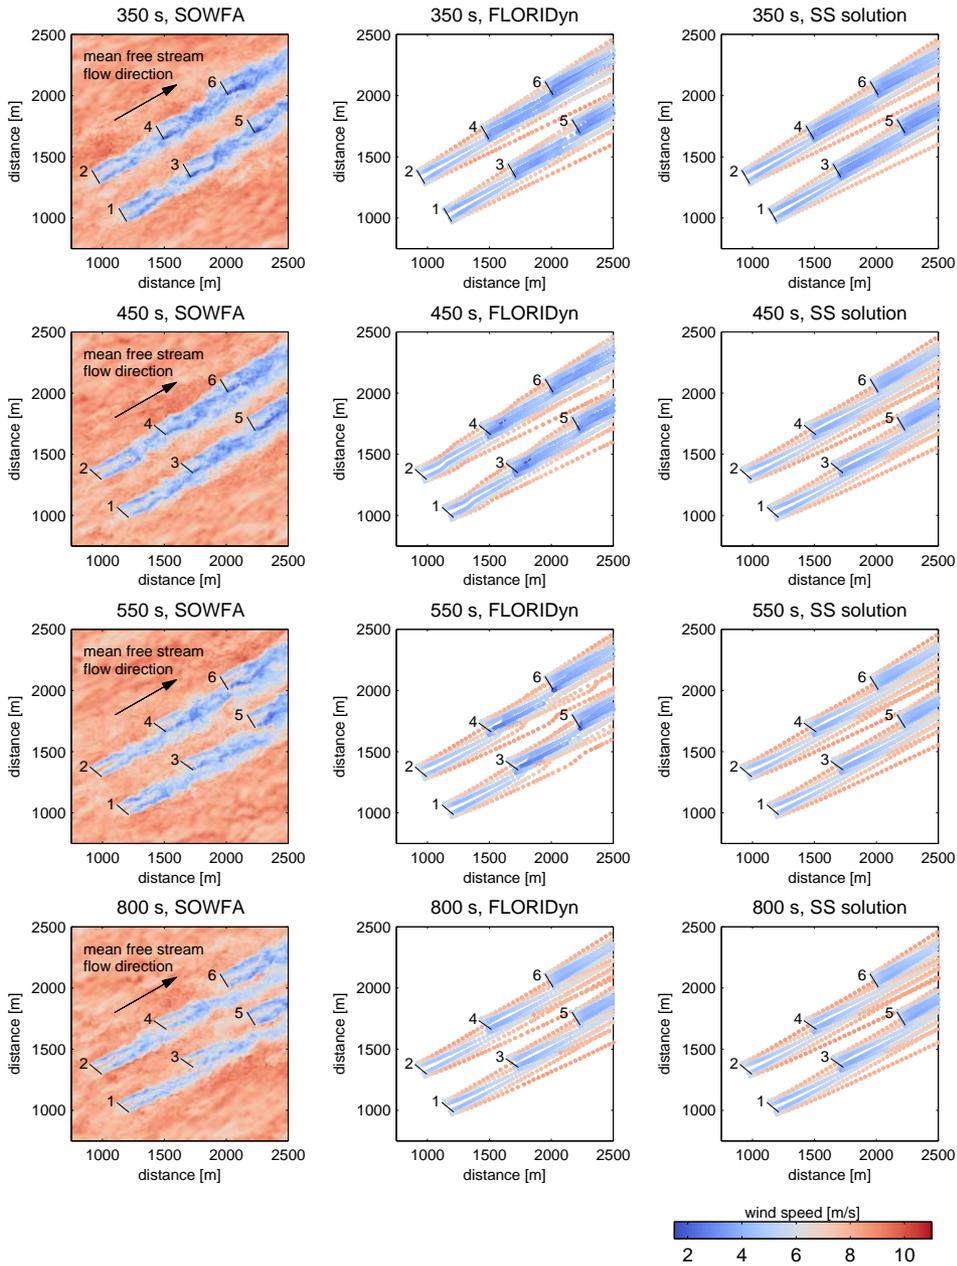


Figure 5.7: Predicted flow fields in the case study of Section 5.3, at different time instances. The dots in the FLORIDYN model and the steady-state (SS) results represent the location of the TrPs (mid-point of the wake zones), and the colors represent predicted the local wind speed. While both SOWFA and FLORIDYN predict a deflection of the wake that propagates through the wind field after the yaw change at 400s (the upstream part of the wake deflects first, then gradually the downstream part), the wake in the steady-state SS model responds instantaneously.

the techniques can be extended to application on real wind plants. Further, it is shown that if we focus on correcting the model-predicted velocity field by adjusting the free-stream velocities at different TrPs, the Kalman filter takes on a relatively simple form. This particular Kalman filter is shown as the red sub-block in the right-hand part of the FLORIDYN model overview in Figure 5.1. In Section 5.4.1, an explanation of this Kalman filter is provided. A simulation case study of applying this filter is presented in Section 5.4.2.

5.4.1. DESCRIPTION OF THE KALMAN FILTER FOR CORRECTING THE PREDICTED VELOCITY FIELD

Equations (5.18)-(5.20) define a time-varying mapping between the free-stream speeds effective at some of the TrPs and the effective free-stream speed at the turbines. We can write this mapping as:

$$\overrightarrow{U_{T,\mathcal{D}}}(k) = C_U(k) \overrightarrow{U}(k) \quad (5.23)$$

with $C_U(k)$ a sparse, time-varying matrix, $\overrightarrow{U_{T,\mathcal{D}}}(k)$ a vector in which all the turbine-effective free-stream speeds at the downstream turbines at time-step k , $\{U_T(t, k) \forall t \in \mathcal{D}\}$, are stacked, and $\overrightarrow{U}(k)$ a vector in which the free-stream velocities at each TrP in the wind plant, $\{U_{t,z,p}(k) \forall t \in \mathcal{P}, z \in \{1, \dots, 7\}, p \in \{1, \dots, N_p\}\}$, are stacked. Further, we can combine the update equations (5.5) for all wake zones of all downstream turbines $t \in \mathcal{D}$ in a system in state-space form:

$$\overrightarrow{U}(k+1) = \mathbf{A} \overrightarrow{U}(k) + \mathbf{B}_{\mathcal{D}} \overrightarrow{U_{T,\mathcal{D}}}(k) + \mathbf{B}_{\mathcal{F}} \overrightarrow{U_{T,\mathcal{F}}}(k) + \overrightarrow{w}(k) \quad (5.24)$$

where \mathbf{A} , $\mathbf{B}_{\mathcal{D}}$ and $\mathbf{B}_{\mathcal{F}}$ are matrices in which the matrices A and B are stacked in block-diagonal form. In the above state-space form, we make a distinction between the input vector consisting of the free-stream velocities effective at the downstream turbines, $\overrightarrow{U_{T,\mathcal{D}}}$, and the input vector consisting of the free-stream velocities effective at the front turbines, $\overrightarrow{U_{T,\mathcal{F}}}$. The vector w is a noise process describing the effect of turbulence and model inaccuracy. When combining equations (5.23) and (5.24) we can see the feedback of the free-stream velocities effective at the downstream turbines $\overrightarrow{U_{T,\mathcal{D}}}$ already mentioned before in Section 5.2.6.3. This leads to the fact that equations (5.23) and (5.24) can be combined in the following closed-loop form:

$$\overrightarrow{U}(k+1) = \tilde{A}(k) \overrightarrow{U}(k) + \mathbf{B}_{\mathcal{F}} \overrightarrow{U_{T,\mathcal{F}}}(k) + \overrightarrow{w}(k) \quad (5.25)$$

with:

$$\tilde{A}(k) = \mathbf{A} + \mathbf{B}_{\mathcal{D}} C_U(k) \quad (5.26)$$

We apply Kalman filtering techniques to make corrections on the model-predicted wake velocity profiles using the electrical power production measurements, and the control settings, at the turbines. We do this by correcting the free-stream velocities $U_{t,z,p}(k)$ at the different TrPs (the state of the above system). In the model, these free-stream velocities are then adjusted with the wake reduction factors to find the wake velocity profile, using equation (5.17), and using equation (5.22) the velocity effective at the turbines are found. Note that the combination of (5.22) and (5.23) can be written as:

$$\overrightarrow{w_{T,\mathcal{D}}}(k) = C_u(k) \overrightarrow{U}(k) + \overrightarrow{v}(k) \quad (5.27)$$

where $C_u(k)$ is a sparse, time-varying matrix, $\overrightarrow{u_{T,\mathcal{D}}}(k)$ is a vector in which all the turbine-effective wind speeds at the downstream turbines at timestep k , $\{u_T(t, k) \forall t \in \mathcal{D}\}$, are stacked, and where $v(k)$ represents output noise. From measurements at the downstream turbines (power, yaw and axial induction) we can estimate these effective velocities by inverting relation (5.1), and construct the vector $\overrightarrow{u_{T,\mathcal{D}}}(k)$. Then using this vector as input, we use a conventional Kalman filter for one-step-ahead prediction of the state of the system with state equation (5.25) and output equation (5.27). We follow the definition of the conventional Kalman filter provided in Verhaegen et al. (2007) (see also Kalman (1960) for the original derivation). When augmented with this Kalman filter, the system is:

$$\overrightarrow{U}(k+1) = \tilde{A}(k) \overrightarrow{U}(k) + \mathbf{B}_{\mathcal{F}} \overrightarrow{U_{T,\mathcal{F}}}(k) + K(k) \left[\overrightarrow{u_{T,\mathcal{D}}}(k) - \widehat{\overrightarrow{u_{T,\mathcal{D}}}}(k) \right] \quad (5.28)$$

where at each time-step the Kalman gain $K(k)$ is updated as:

$$P(k+1|k) = \tilde{A}(k) P(k|k-1) \tilde{A}(k)^T + Q - \tilde{A}(k) P(k|k-1) C_u(k)^T (R + C_u(k) P(k|k-1) C_u(k)^T)^{-1} \times (\tilde{A}(k) P(k|k-1) C_u(k)^T)^T \quad (5.29)$$

$$K(k) = \tilde{A}(k) P(k|k-1) C_u(k)^T (R + C_u(k) P(k|k-1) C_u(k)^T)^{-1} \quad (5.30)$$

where Q is the covariance matrix of the noise \overrightarrow{w} , and R the covariance matrix of the noise \overrightarrow{v} (for simplicity we assume the two noise sources \overrightarrow{w} and \overrightarrow{v} are not correlated). If we assume that the noise \overrightarrow{w} describes velocity variations at each TrP as a consequence of turbulence, we can interpret the covariance matrix of w as a matrix that describes the correlation between the turbulence at the different TrPs. Therefore, if we consider the matrix Q as a tuning parameter, it makes sense to simplify the tuning problem by parameterizing Q as follows:

$$Q = \begin{bmatrix} \bar{Q} & 0 & \cdots & 0 \\ 0 & \bar{Q} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \bar{Q} \end{bmatrix} \quad \text{with} \quad \bar{Q} = \begin{bmatrix} q_0 & q_1 & q_2 & \cdots & q_{N_p} \\ q_1 & q_0 & q_1 & \ddots & \ddots \\ q_2 & q_1 & \ddots & \ddots & \ddots \\ \vdots & \ddots & \ddots & \ddots & \ddots \\ \vdots & \ddots & \ddots & \ddots & q_2 \\ q_{N_p} & \cdots & q_2 & q_1 & q_0 \end{bmatrix} \quad (5.31)$$

Now, according to the above reasoning, scalar q_0 can be interpreted as the autocorrelation of the turbulence at each TrP, q_1 the cross-covariance of the turbulence of one TrP with its neighbouring TrP, scalar q_2 the correlation of the turbulence of two TrPs separated by one other TrP, etc. I.e., under the assumption that the Kalman filter mainly has to correct for homogeneous turbulence, for which the correlations are dependent on the relative distance between the points, rather than on the specific location (Carley, 2011), the matrix Q having the structure above, prescribes how well the turbulence remains correlated over distance. It has to be noted that with this interpretation, the fact that the Q matrix is defined constant, while the inter-TrP distances vary, is considered a further simplification.

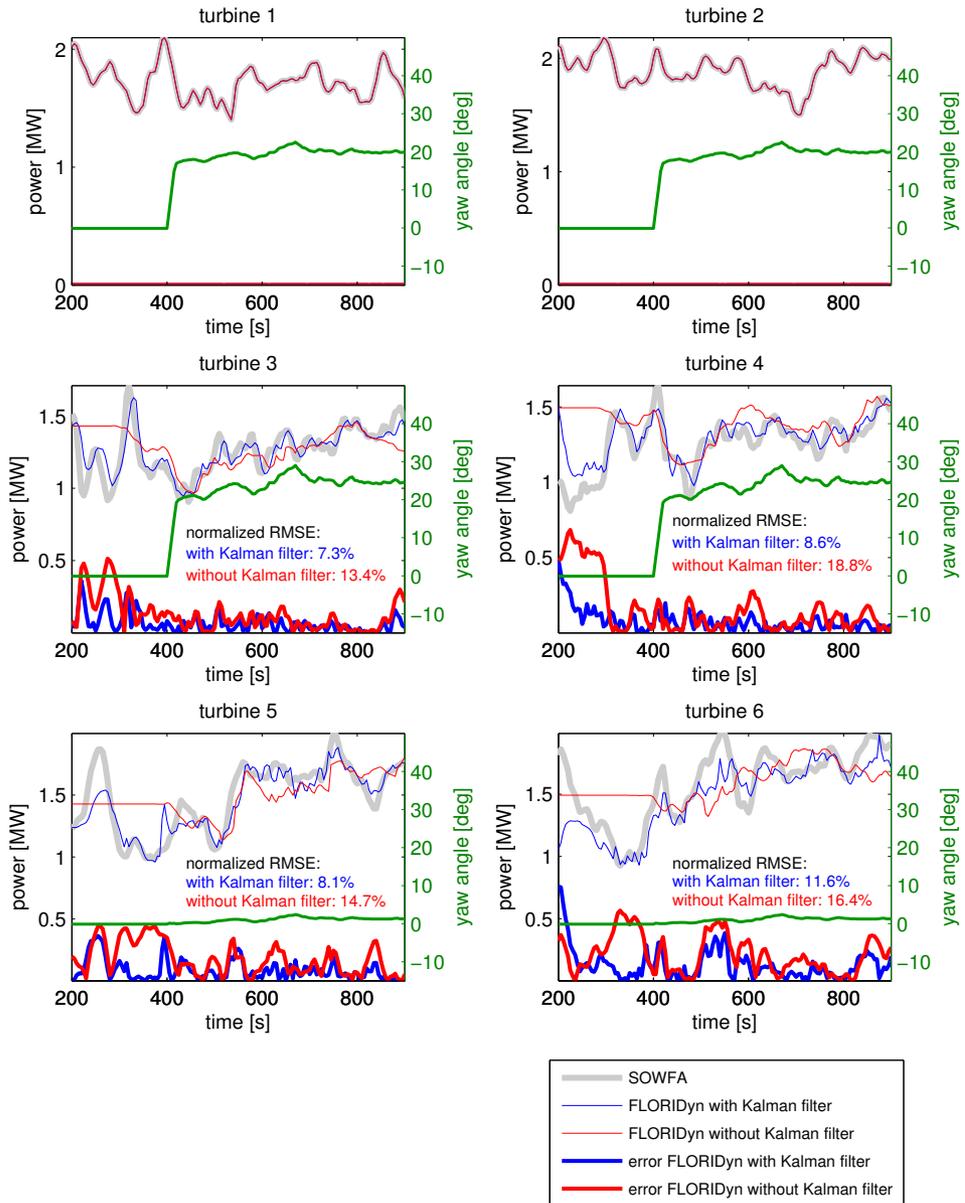


Figure 5.8: Result of the case study with the Kalman filter described in Section 5.4: predicted power productions, and yaw angles. With the Kalman filter enabled, corrections on the flow field can be performed based on the measured data (see also Figure 5.9), that result in a better fit the data when compared to the FLORIDYN model without these corrections. Also, the initial state of the flow field can be reconstructed from the measurements, which results in a much better fit in the initial response.

R	q_0	q_1	q_2	q_3	q_4	q_i for $5 \leq i < N_p$
$0.88I_{N_t \times N_t}$	0.14	0.060	0.052	0.041	0.0022	0

Table 5.4: Specification of the Q and R parameter values used in the Kalman filter in the case study of Section 5.4. The N_t by N_t identity matrix is denoted by $I_{N_t \times N_t}$, so matrix R is a diagonal matrix. Matrix Q is parameterized as in eq. 5.31, with q_\bullet values as listed above.

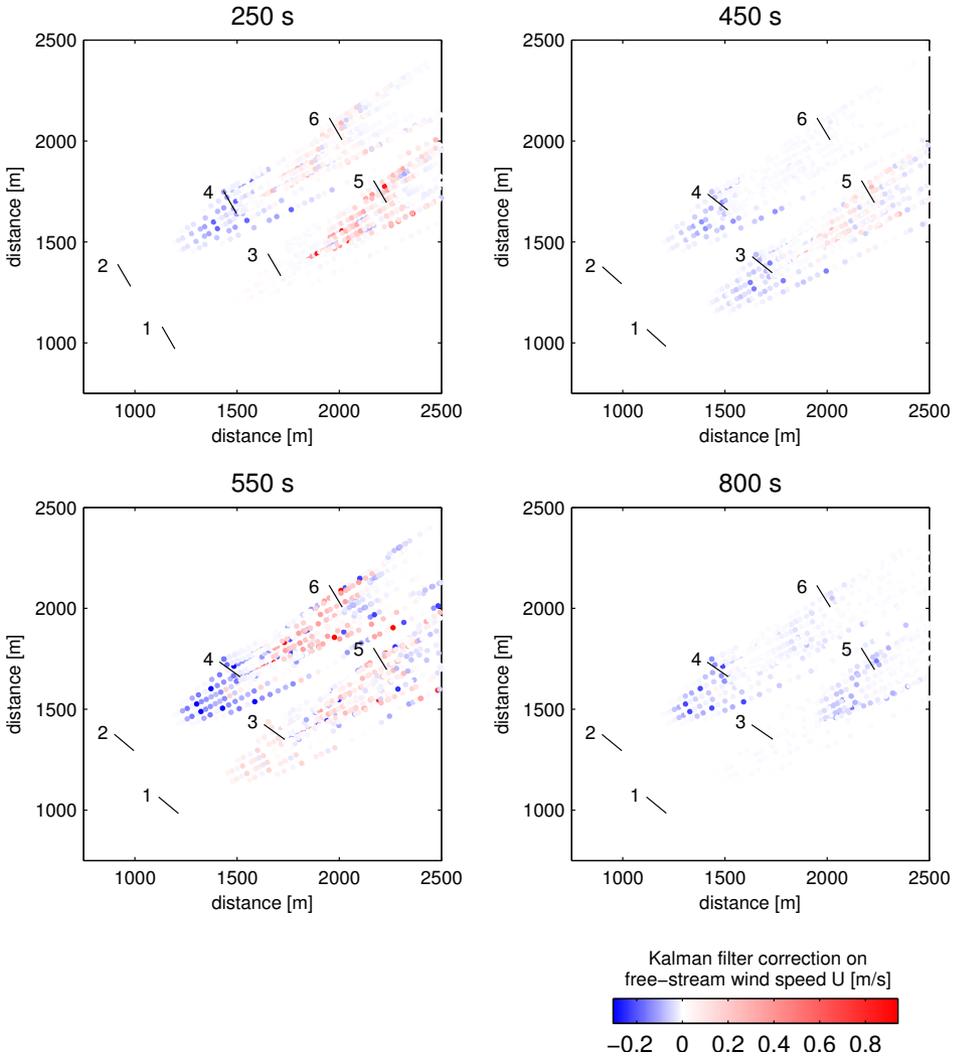


Figure 5.9: Result of the case study with the Kalman Filter described in Section 5.4: at different time instances in the simulation, it is shown for each TrP which correction on the free-stream velocity is applied by the Kalman filter

5.4.2. KALMAN FILTERING CASE STUDY

In Figure 5.8, the results are shown of applying a Kalman filter in the case study described in Section 5.3. The parameters Q and R used in this case are listed in Table 5.4, they were tuned to minimize the sum of RMSE for each turbine. Two effects can be seen in the results of this case study:

- An overall better fit of the responses can be expected by using the Kalman filter for one-step ahead prediction of the responses.
- Further, a feature of a Kalman filter is that it is able to correct the response to an initial guess of the unknown state of a system in the first part of a simulation by using measured data. The FLORIDYN model without the Kalman filter only uses measurements at the front turbines to estimate the free-stream velocity. Since wake traveling dynamics are slow, the transient response to the initial guess of the free-stream velocity at the downstream turbines are long, and therefore there is a large error in the initial response if the initial guess of the flowfield are not accurate. It can be seen in the responses that the Kalman filter can ensure a more quick mitigation of that error, by using the measured data on the downstream turbines to ‘reconstruct’ the velocity field in the downstream part of the wind plant.

As an illustration of the Kalman filtering process, in Figure 5.9 the corrections are shown that the filter performs on the free-stream velocity at the different TrPs upstream of the turbine.

5.5. APPLICATION OF THE FLORIDYN MODEL FOR OPTIMIZED YAW CONTROL

In this section, we give a case study of how a dynamic control-oriented model of the wake effects such as the FLORIDYN model can be used to optimize the control of the wind plant. By optimizing the control signals based on the predictions provided by the FLORIDYN dynamic model, rather than on a steady-state model, we account for delays in the wake effects.

In this case-study, we consider a scenario in which three turbines are placed in a line that is 5° offset from the the wind direction, with a 5 rotor diameter spacing. After 200s of simulated time, we simulate a short scheduled shut-down of the second turbine through blade feathering, by setting the axial induction of that turbine to zero for 200s. It has to be noted that this is an unrealistic reduction of axial induction, which just serves for a proof of the control concept. We aim to optimize the time sequence of yaw settings of the turbines for increased electrical energy production, by taking into account the wake traveling dynamics. To keep the optimization computationally efficient, we use physical reasoning to parameterize a solution for the optimization problem and reduce the parameter search space. We know that if the second turbine in the row will not extract power anymore, the front turbine can decrease its yaw angle to increase power extraction, since it does not have to steer away its wake from the next downstream turbine 2, but from turbine 3 that is standing further downstream. Further, we can assume that turbine 1 will adjust its yaw some time before turbine 2 pitches to feather, because the

wake effect of the yaw change of turbine 1 will reach the downstream turbines with some delay. We let turbine 1 respond to a reference by yawing with a maximum rate of $1^\circ/s$, and we restrict to positive yaw angles. If we prescribe that turbine 1 will adapt by making a step change on its yaw reference setting, and stepping back to its original yaw setting later, we can parameterize the solution using three parameters:

- $\Delta\gamma_T(1)$, the size of the yaw reference step change of turbine 1,
- k_1 , the time-step at which turbine 1 makes the yaw reference step change,
- k_2 , the time-step at which turbine 1 steps its yaw back to the initial settings.

Note that using the above notions, we can prescribe the following constraints on the optimization parameters $\Delta\gamma_T$, k_1 , k_2 :

$$\begin{aligned} -\gamma_T(1,0) < \Delta\gamma_T(1) < 0 \\ 0 < k_1\Delta T < 200s \\ k_1\Delta T < k_2\Delta T < 400s \end{aligned} \quad (5.32)$$

where $\gamma_T(1,0)$ is the initial yaw setting of a turbine 1, and ΔT the sample time. First, we search for the initial yaw settings for each turbine, which are the optimized yaw setting for steady-state operation with the axial induction of all the turbines at a constant $a_T = 1/3$. We perform an exhaustive grid search for these optimal initial yaw settings using the FLORIS model, with an incremental step of 0.5° for each yaw setting, and find that they are $\gamma_T(1,0) = 9.5^\circ$ for the front turbine, $\gamma_T(2,0) = 19.5^\circ$ for the middle turbine, and $\gamma_T(3,0) = 0^\circ$ for the back turbine. Then, in the next step we use the model to perform the optimization of parameters $\Delta\gamma_T(1)$, k_1 , k_2 of the adaptive yaw control sequence, with maximum total electrical energy production in the control horizon as the objective. We again perform an extensive grid search but in each parameter set evaluation we now simulate the dynamic system response to a particular yaw reference sequence with the FLORIDYN model. In each of these simulations, we use a time-step $\Delta T = 5s$, and a total simulated time of 600s. The grid search of the solution space prescribed by the inequalities (5.32), is first performed with an incremental step of 5 on k_1 and k_2 (i.e. 25s) and 2° on $\Delta\gamma_T(1)$, then the parameter search grid is refined around the optimal solution to incremental step of 1 on k_1 and k_2 (5s) and 1° on $\Delta\gamma_T(1)$. A total of 735 FLORIDYN simulations are needed to search the parameter space in this way. On average it takes 0.51s to evaluate a FLORIDYN simulation of the case in a MATLAB implementation on a 1.6 GHz PC, yielding a total calculation time of 375s to perform the parameter search. In Figure 5.10 the resulting optimized yaw sequence, as well as the resulting turbine power responses as predicted by the FLORIDYN model are shown (this case is referred to as ‘adaptive yaw’), and compared to the case where the yaw settings of each turbine are held constant throughout the simulation and thus does not adapt to the shutdown of turbine 2 (‘constant yaw’). Further the difference in cumulative electrical energy production between the adaptive yaw control and the constant control is plotted. It can be seen that a relatively small amount of energy production increase (0.19% of the total energy production) can be gained by the fact that turbine 1 reduces its yaw angle. Turbine 1 yaws at a time interval that is earlier than the interval of the shutdown of turbine 2, such that the wake traveling delays in the systems are accounted for. Figure 5.11 shows the wake velocity profiles predicted by FLORIDYN for the adaptive control case, with a short description of the different steps in the adaptive yaw control procedure.

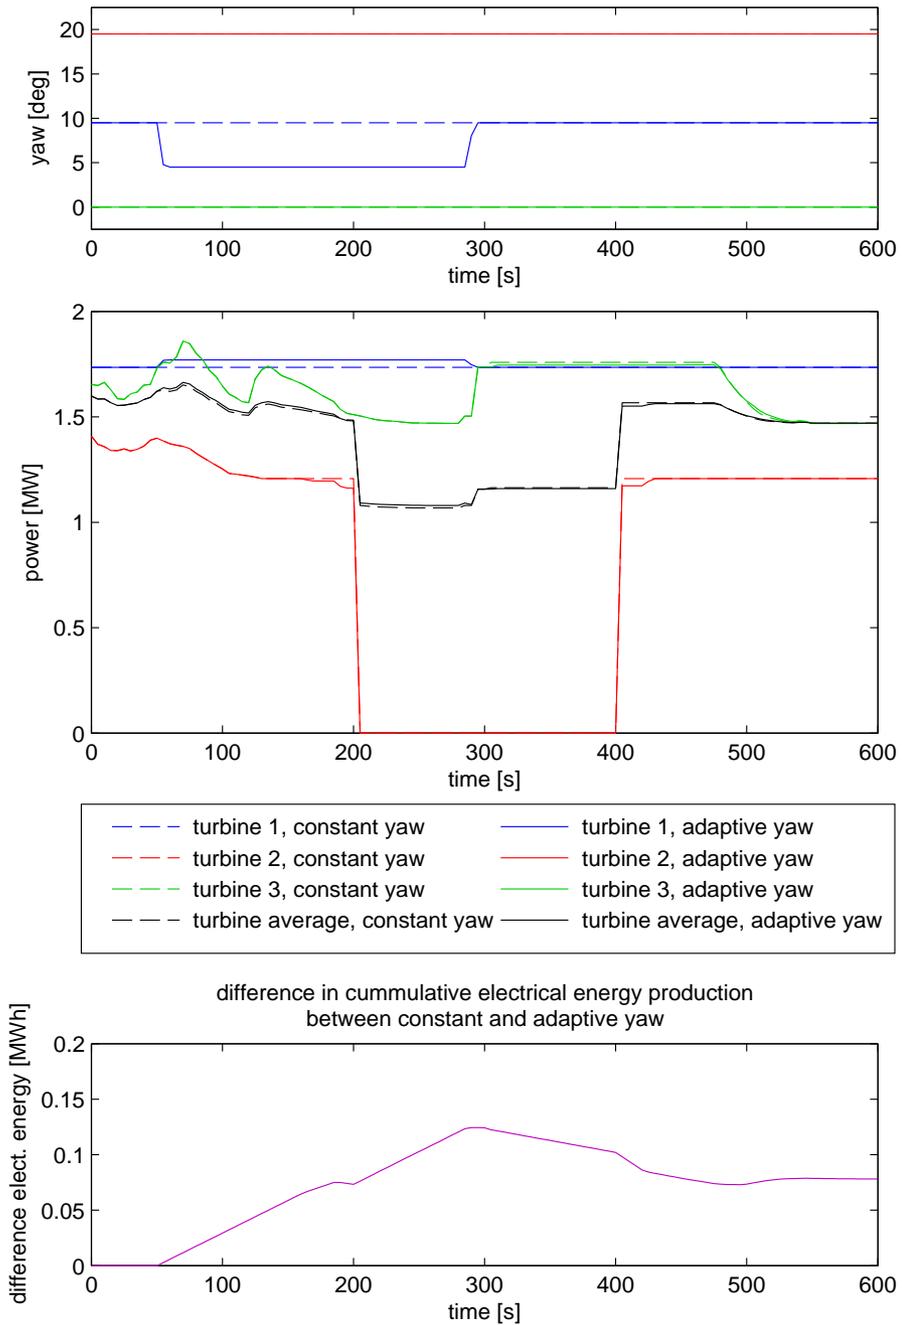
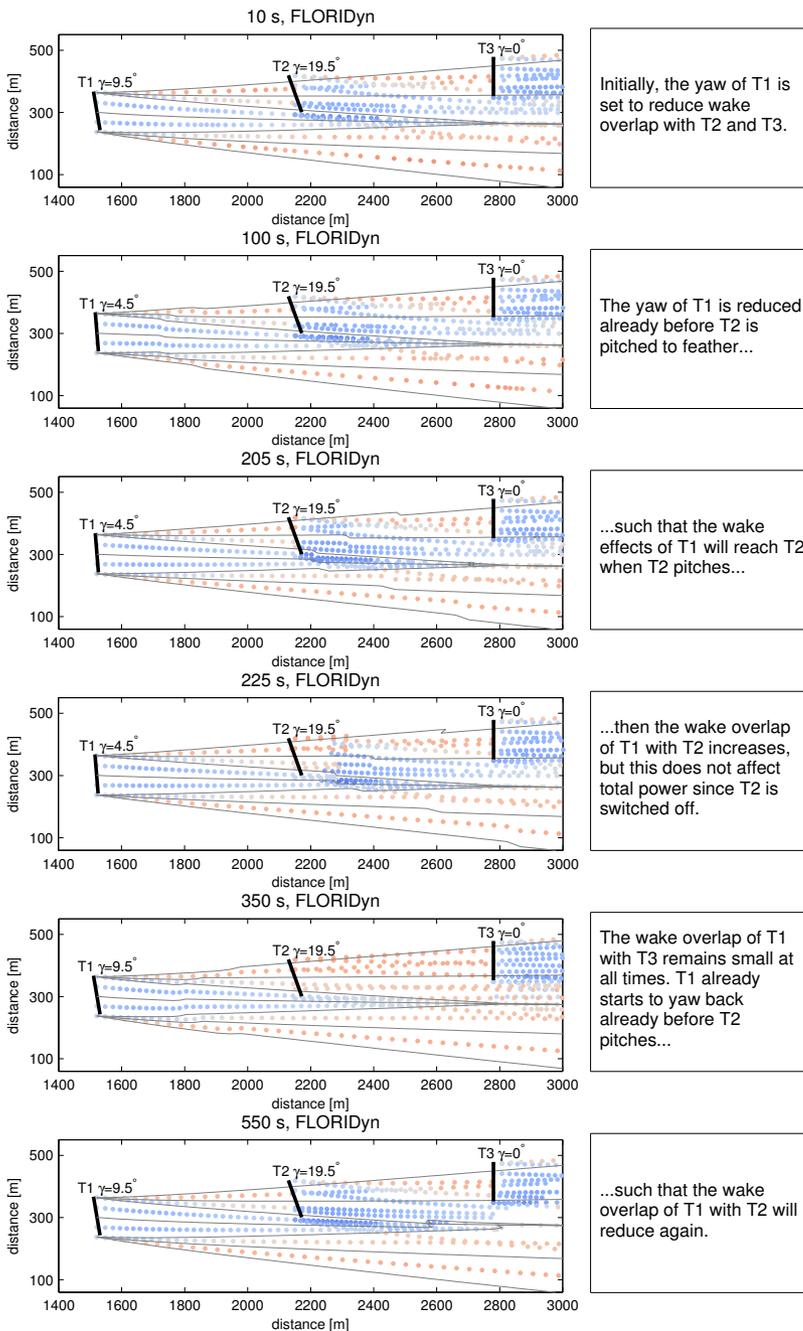


Figure 5.10: Result of the control case study described in Section 5.5: turbine yaw angles (top) and predicted power productions (middle) for the case where an optimized adaptive yaw control sequence is used, and a case where the yaw is held constant. Also shown is the cumulative difference in energy production that is gained by using the adaptive yaw control instead of constant yaw control in the simulated scenario.



Initially, the yaw of T1 is set to reduce wake overlap with T2 and T3.

The yaw of T1 is reduced already before T2 is pitched to feather...

...such that the wake effects of T1 will reach T2 when T2 pitches...

...then the wake overlap of T1 with T2 increases, but this does not affect total power since T2 is switched off.

The wake overlap of T1 with T3 remains small at all times. T1 already starts to yaw back already before T2 pitches...

...such that the wake overlap of T1 with T2 will reduce again.

Figure 5.11: Wake velocity profiles predicted by FLORIDYN for the adaptive yaw control case described in Section 5.5 (left), colored dots are TrPs, grey lines are wake zone boundaries. Also, a short description is provided of the different steps in the adaptive yaw control procedure (right).

5.6. DISCUSSION, CONCLUSIONS AND FUTURE WORK

The results from the case studies are promising: FLORIDYN is an extension of the steady-state FLORIS model to include the dynamics of the wake propagation, that is able to provide a prediction of the wind plant dynamic response with reasonable accuracy when compared to the results from SOWFA (a high-fidelity CFD model). An important feature is that no additional tuning parameters were introduced in the part of the model that describes the wake expansion and recovery properties, and only two new parameters were introduced in the state-space wake propagation model, such that the tuning process is not further complicated to a large extent.

The wake propagation effects are for a large part described by a linear state-space model. This allowed the development of a (linear) observer based on a Kalman filter using conventional techniques. The filter makes corrections on the velocities in the wind field to account for model inaccuracies and smaller-scale turbulence effects, based on measured power data from the turbines.

A control example was provided in which it is shown that the scheduling of the yaw settings can be optimized by taking into account dynamic effects in the wakes, mainly consisting of delays that are dependent on the wake velocity profile. However, we did not demonstrate a large increase of wind plant performance by going from control optimization based on a steady-state model, such as the FLORIS model, to a dynamic model, such as the FLORIDYN model.

Based on this experience, we expect the main benefit of using a control-oriented dynamic wake model is that the parameters of such a model can be optimized based on the dynamic response of the wind plant (like was done for the case study in Section 5.3), while a static model, such as the FLORIS model, is tuned based on time-averaged data of a period in which the wakes are fully propagated through the wind plant (see the procedures described in Section 4.2 and 4.3.7). This could be an important benefit when applying model-based control optimization on a real wind plant with continuously changing conditions.

A key requirement for enabling online optimization of the control, is a low computational cost of the internal model of the controller. In the control case study, the relatively low computational cost of the FLORIDYN model, allowed us to do an extensive search for the optimized yaw sequence. As mentioned before, still there is a relatively high computational cost when compared to FLORIS, because of the interpolation steps that are needed to evaluate the dynamics of the interaction between wakes. Therefore, more efficient optimization strategies should be explored in order to reduce the computation time of finding the optimized yaw settings. Also possibly the computational cost of FLORIDYN can be further reduced by using alternative ways to model the dynamics of the interaction, or simply by generating a more efficient implementation of FLORIDYN in compiled code.

That said, this work is only a first exploration into the use of a control-oriented parametric model of wind plant dynamics based on heuristic description of the wake dynamic properties. The simulation examples provided in this chapter provide a proof of concept, but they should be further validated. We recommend the following future de-

velopments for the FLORIDYN model (or similar control-oriented dynamic wake models):

- Gathering validation data from relevant and realistic scenarios with changing wind conditions from high-fidelity models like SOWFA, or with a real wind plant. One of these scenarios should be a changing wind direction. This is a substantial task, given the computational costs of high-fidelity models, and the uncontrollable nature of the conditions in real wind plants. It has to be investigated to which extend a simulation environment like SOWFA, in which the inflow conditions are generated in a highly-controlled fashion, is able to mimic the real-world changing inflow scenarios, so that repeated simulation experiments can be performed with different control strategies, under realistic conditions.
- Further validation of the FLORIDYN model, and developing extensions to include, for example, changing wind directions and wake meandering.
- The further development of observers for the model. Most importantly in our view, for further controls development, an observer should be developed that automatically updates the wake parameters of the model under changing flow conditions (turbulence, atmospheric stability) based on data measured at the turbines. Also, a possible extension is the development of an observer that corrects the centerline position of the wake to account for wake meandering effects. An example of the latter concept, consisting of a particle filter that estimates the wake centerline position based on measured turbine data, is provided in Fleming et al. (2014a).
- The exploitation of the model structure (a linear time-invariant state-space model combined with a nonlinear static feedback term) in the calculation of optimized control sequences, and/or the development of a more computationally efficient implementation of the model, to come to a control scheme for online optimization of the yaw-settings based on the predictions of the dynamic model.
- Exploring the possible benefits of an adaptive dynamic model-based control strategy in terms of its effect on energy yield of the wind plant, and loads on the individual turbines, in realistic wind scenarios and fault scenarios. As mentioned before in the case study provided in Section 5.5, there is only a small beneficial effect from using a dynamic model for model-based control, rather than a static model, but under time-varying conditions (e.g. a changing wind direction) the ability to perform dynamic control of the flow in the wind plant may be more essential.

6

CONCLUSIONS AND RECOMMENDATIONS

In this chapter, the main conclusions of the research presented in this thesis are summarized (Section 6.1), and recommended directions for future research are presented (Section 6.2).

6.1. CONCLUSIONS

In this thesis two research objectives have been addressed: the evaluation of the potential of the control degrees-of-freedom of the wind turbine to affect the interaction effects between the turbines, and the development of data-driven algorithms for the optimization of those control settings in order to improve wind plant performance. Below, the conclusions on each of the objectives are presented.

6.1.1. EVALUATION OF WIND TURBINE CONTROL DEGREES-OF-FREEDOM FOR WAKE CONTROL

As a general conclusion, high-fidelity simulation results show that there is more potential for yaw-based wake redirection wind plant control than for axial-induction-based wind plant control, to improve wind plant performance. Yaw-based control is predicted to yield a 1.1% electricity production increase on an annual basis, on a full-size offshore wind plant, which is considered a significant contribution towards the goal of reducing the costs of wind energy.

axial-induction-based control The results of SOWFA high-fidelity simulations of a two-turbine setup in conditions with low ambient turbulence conditions, in Chapter 2, show that axial-induction-based control techniques, using adaptation of generator torque and

collective blade pitch, have limited or no potential to increase electrical power production, because wake expansion makes that much of the energy added to the flow by the control changes, is lost to the atmosphere. In Chapter 3, a better potential of these techniques on a full-size wind plant (1.4% electricity production increase on an annual basis) is predicted using the Jensen engineering model, but it should be noted that further research (Annoni et al., 2014a) indicates that there is a significant discrepancy between the predictions of the SOWFA high-fidelity model, and engineering models in predicting the results of axial-induction-based control.

wake redirection control In Chapter 2 and 4, yaw control has been shown to be successful at inducing wake redirection and increasing the total power production of small wind plant setups, simulated with the low ambient turbulence conditions in SOWFA. In these simulations, both increases and reductions of fatigue loads result from yaw-based wake redirection, depending on the settings used, but the load increases can partly be mitigated by the use of standard load-reducing IPC. With optimized yaw settings, significant fatigue loads reductions could be achieved on a small wind plant setup. Based on the FLORIS data-driven model, a 1.1% electricity production increase is predicted on an annual basis, through using yaw-based wind plant control on a full-size offshore wind plant, the Princess Amalia Windpark.

6.1.2. DATA-DRIVEN METHODS FOR WIND PLANT CONTROL

The main conclusion with respect to the development of wind plant control algorithms, is that data-driven methods are needed in order to optimize the control settings for the specific ambient conditions in which the wind plant is operating, and that in developing these data-driven methods, the time-efficiency of the control method has to be taken into account in order to enable real-time implementation on the wind plant, in which special attention has to be paid to the wake propagation delays between the turbines.

Two main approaches for the development of time-efficient data-driven wind plant control were taken in this thesis: direct data-driven control development (control settings optimization directly based on measured data), and data-driven model-based control development (control optimization based on models for which the parameters are identified from measured data).

direct data-driven control In order to develop direct data-driven methods for real-time wind plant control, the wake propagation delays between turbines have to be taken into account in order to evaluate the time-efficiency of the optimization. Significant improvement of the time-efficiency can be achieved by using knowledge of the plant lay-out in the control strategy. This was demonstrated in Chapter 3, where it was shown that gradient-based axial-induction-based control optimization techniques with a distributed approach in which the effect on neighboring turbines is taken into account only, have a much faster convergence of the power optimization than is achieved with an existing method with full communication between the turbine.

data-driven model-based approaches Data-driven model-based wind plant control strategies are an effective way to optimize the yaw settings of each turbine in the plant with the objective of electricity production increase. In Chapter 4 a parametric model of steady-state wake deflection and velocity deficit was developed (FLORIS). This model has a relatively small number of parameters that can be identified using time-averaged electrical power measurements of the different turbines in the wind plant (a gray-box system identification approach). The model is computationally efficient enough to enable real-time online wind plant control optimization, using a game-theoretic search method.

The model can be further extended to include dynamics of the wake propagation delays between the turbines, which are dependent on the velocity field in the wake (Chapter 5). A feedback structure with a linear state-space model describes the propagation of changes in the wake through the wind field. This model structure has a relatively low complexity since only two tunable parameters are added to include the dynamics. This model structure enables the development of an observer that uses Kalman filtering to correct the flow field velocities on the basis of power measurements at the turbines. A relatively small increase of wind plant performance can be gained by going from yaw control optimization based on the steady-state model to yaw control based on the dynamic model. Therefore, the main practical benefit of including the dynamics, lies in the fact that the gray-box system identification can be performed directly on the basis of the dynamic responses of the wind plant, rather than based on time-averaged data, which is a benefit when applying data-driven model-based control on a real wind plant with continuously changing ambient conditions.

6.2. RECOMMENDATIONS

In this section, general recommendations are given for future research in the field of wind plant control. For more specific recommendations, we refer to the different chapters, in which more detailed suggestions for future work on the methods presented in this thesis are made.

With respect to the further evaluation of the control degrees-of-freedom to affect the wake properties, recommended future work aims at determining the influence of ambient atmospheric conditions (inflow turbulence, atmospheric stability, surface roughness) on the ability to affect the wake interaction effects. Because turbulence promotes wake recovery, it seems apparent that the level of ambient turbulence is negatively correlated with the potential of wind plant control to improve wind plant performance, but the exact influence is not quantified yet. Ongoing research in axial-induction-based control by Annoni et al. (2014a) also proposes extensions to the FLORIS model in order to take into account the effect of control settings themselves on turbulence and wake recovery, and this is shown to be crucial in evaluating the potential of axial-induction-based wind plant control.

With respect to direct data-driven methods, the main challenge for future work is to test these techniques in more realistic simulation scenarios, or on a real wind plant. Wind scenarios in which there is a change of wind direction, are especially interesting,

since in wind plants the wake interaction effects are sensitive with respect to wind direction. Special-purpose filtering techniques are to be developed that remove the distortions caused by small-scale turbulence in the measurements that the direct data-driven methods use.

For the further development of control-oriented data-driven wind plant model, three ‘paths’ can be chosen, that are outlined below. We give recommendations for each of the paths.

1. Extending parametric engineering wake models such that they emulate control-relevant features of the wake, as in Chapter 4 and 5. The parameters are to be found from high-fidelity, validated CFD models or measured data from a real wind plant (gray-box system identification), and should be scheduled with the specific ambient conditions.

A control-relevant feature of the wake that is not treated in detail in the wind plant models presented in this thesis, are the wake meandering motions. Although on large-scale wind turbines, yaw is a relatively slow control degree-of-freedom that deals with the large inertia of the rotor, and is therefore not likely to be suited for compensation of the wake meandering motions, still meandering could be taken into account as an uncertainty in the wake position, when developing a robust control framework. In the static FLORIS model, only the time-averaged effects of wake meandering are taken into account as an increased wake expansion. In a dynamic wake propagation model such as FLORIDYN however, the wake meandering should be taken into account, especially when designing an observer that estimates the wake state from measured data (Fleming et al. (2014a) presents an interesting particle filter approach for estimating the wake position from measured data).

Further, ongoing research by Annoni et al. (2014a) has aimed at including the effects of wake overlap in longer rows of turbines, on wake recovery.

For further controls development, an observer should be developed that automatically updates the wake parameters of parametric model under changing flow conditions (turbulence, atmospheric stability) based on data measured at the turbines. When it comes to the model-based control optimization techniques, the applied game-theoretic method has shown to be sufficiently efficient to enable real-time implementation on a small wind plant setup, but possibly more efficient (gradient-based) optimization schemes are needed to enable real-time control on larger wind plants. While in our simulation examples with yaw control it was shown that load reduction can be a side-effect of energy yield maximization, possibly a more direct approach can be developed in which (limitations on) estimated fatigue loads are taken into account in the model, and in the objective of the optimization.

2. Obtaining wake interaction models through black-box system identification The identification experiments should be repeated for different ambient conditions, and the resulting models should be combined (scheduled), to describe a full range of ambient conditions.

Because a lot of system identification experiments need be performed to develop a model for the full range of operational conditions, it is expected that the main value of black-box approaches is in supporting the other paths; black-box system identification could be used for finding specific unknown subsystems, for validation of other models, and possibly for reduction of CFD-based models (path 3).

In wind plant system identification, it is of interest to exploit knowledge on the layout of the wind plant and wind direction measurements, in order to find the topology of the network that describes the turbine-to-turbine wake interactions in the plant. Special-purpose system identification methods for networks with a known topology, such as those presented in Dankers (2014), can then be used. An alternative approach is to use system identification methods that deal with the topology reconstruction of interconnected systems from measured data, such as those presented in Torres Tapia (2014).

3. Developing wind plant models from a starting-point in which Navier-Stokes CFD models are converted to control-oriented models through simplification, projection, and order-reduction steps. Enabling real-time implementation is especially challenging for this path, since current high-fidelity CFD wind plant models, such as SOWFA, have a large computational complexity, and simplification may lead to significant changes in the prediction, as the study in Annoni et al. (2014b) shows. In this work, special-purpose numerical techniques that enable computationally efficient control synthesis for large-scale distributed systems, such as those presented in Rice (2010), are of interest.

For all of the above paths, we recommend that wind plant controls are taken into account in the design stage of wind plants. Ongoing research in this direction is presented in Fleming et al. (2014c), where optimization of control and wind plant topology is combined. In this work, it is shown that wind plant control is applied, the turbines in a wind plant may be placed closer together because the wake interaction effects are mitigated, enabling a further growth of wind plant electrical energy production per given land area.

The above recommendations show that this thesis presented pioneering work in wind plant controls, researching the feasibility of control concepts and solving practical issues to enable real-time implementation, but that further research is needed to deepen the knowledge on modeling and control solutions for wind plant optimization.

BIBLIOGRAPHY

- M. Abkar and F. Porté-Agel. The effect of atmospheric stability on wind-turbine wakes: A large-eddy simulation study. In *The Science of Making Torque from Wind*, Lyngby, Denmark, 2014.
- M. S. Adaramola and P. Å. Krogstad. Experimental investigation of wake effects on wind turbine performance. *Renewable Energy*, 36(8):2078–2086, 2011.
- M. A. Ahmad, S. Azuma, and T. Sugie. A model-free approach for maximizing power production of wind farm using multi-resolution simultaneous perturbation stochastic approximation. *Energies*, 7(9):5624–5646, 2014.
- J. Aho, A. Buckspan, J. Laks, P. Fleming, Y. Jeong, F. Dunne, M. Churchfield, L. Pao, and K. Johnson. Tutorial of wind turbine control for supporting grid frequency through active power control. In *Proceedings of the American Control Conference (ACC)*, pages 3120–3131, Montréal, Canada, 2012.
- J. F. Ainslie. Calculating the flowfield in the wake of wind turbines. *Journal of Wind Engineering and Industrial Aerodynamics*, 27(1–3):213–224, 1988.
- J. Annoni, P. M. O. Gebraad, A. K. Scholbrock, P. Fleming, and J. W. van Wingeren. Analysis of axial-induction-based wind plant control using low-order and high-order wind plant models. submitted, 2014a.
- J. Annoni, P. Seiler, K. Johnson, P. Fleming, and P. Gebraad. Evaluating wake models for wind farm control. In *Proceedings of the American Control Conference*, Portland, OR, USA, 2014b.
- S. Aubrun, T. F. Tchouaké, G. España, G. Larsen, J. Mann, and F. Bingöl. Comparison between wind tunnel and field experiments on wind turbine wake meandering. In *Progress in Turbulence and Wind Energy IV - Proceedings of the iTi Conference in Turbulence 2010*, Bertinoro, Italy, 2012.
- H. Bai, M. Arcak, and J. Wen. What is cooperative control? In *Cooperative Control Design*, volume 89 of *Communications and Control Engineering*, pages 1–4. Springer New York, 2011. ISBN 978-1-4614-0013-4.
- R. Barthelmie, S. Frandsen, K. Hansen, J. Schepers, K. Rados, W. Schlez, A. Neubert, L. Jensen, and S. Neckelmann. Modelling the impact of wakes on power output at Nysted and Horns Rev. In *Proceedings of the European Wind Energy Conference*, Marseille, France, 2009.
- R. J. Barthelmie and L. E. Jensen. Evaluation of wind farm efficiency and wind turbine wakes at the Nysted offshore wind farm. *Wind Energy*, 13(6):573–586, 2010.

- R. J. Barthelmie and S. C. Pryor. Potential contribution of wind energy to climate change mitigation. *Nature Climate Change*, 4:684–688, 2014.
- F. D. Bianchi, H. D. Battista, and R. J. Mantz. *Wind Turbine Control Systems; Principles, modelling and gain scheduling design*. Springer-Verlag, London, UK, 2007.
- E. Bitar and P. Seiler. Coordinated control of a wind turbine array for power maximization. In *Proceedings of the American Control Conference*, pages 2898–2904, Washington, D.C., USA, 2013.
- E. A. Bossanyi. Individual blade pitch control for load reduction. *Wind Energy*, 6(2): 119–128, 2003a.
- E. A. Bossanyi. Wind turbine control for load reduction. *Wind Energy*, 6(3):229–244, 2003b.
- E. A. Bossanyi. Controller for 5MW reference turbine. Technical Report 11593/BR/04, Garrad Hassan and Partners Limited, 2009.
- S. P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, UK, 2004.
- A. J. Brand and J. W. Wagenaar. A quasi-steady wind farm flow model in the context of distributed control of the wind farm. In *European Wind Energy Conference*, Warsaw, Poland, 2010.
- A. J. Brand, J. W. Wagenaar, P. J. Eecen, and M. C. Holtslag. Database of measurements on the Offshore Wind Farm Egmond aan Zee. In *Proceedings of the EWEA Annual Meeting*, Copenhagen, Denmark, 2013.
- M. Buhl Jr. MCrunch theory manual for version 1.00. Technical Report NREL/TP-500-43139, National Renewable Energy Laboratory, 2008.
- T. Burton, D. Sharpe, N. Jenkins, and E. A. Bossanyi. *Wind Energy Handbook*, chapter The Controller, pages 144–146. John Wiley and Sons, Ltd, 2002a. ISBN 9780470846063.
- T. Burton, D. Sharpe, N. Jenkins, and E. A. Bossanyi. *Wind Energy Handbook*, chapter Vortex cylinder model of the yawed actuator disk, pages 471–509. John Wiley and Sons, Ltd, 2002b. ISBN 9780470846063.
- M. Carley. Turbulence and noise. <http://people.bath.ac.uk/ensmjc/Notes/tnoise.pdf>, 2011.
- J. Choi and M. Shan. Advancement of Jensen (Park) wake model. In *Proceedings of the EWEA Annual Meeting*, Vienna, Austria, 2013.
- M. Churchfield, S. Lee, P. Moriarty, L. Martinez, S. Leonardi, G. Vijayakumar, and J. Bresseur. A large-eddy simulation of wind-plant aerodynamics. In *Proceedings of the AIAA Aerospace Sciences Meeting*, Nashville, Tennessee, USA, 2012a.
- 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. *Journal of Turbulence*, 2012b.

- A. Crespo, J. Hernández, and S. Frandsen. Survey of modelling methods for wind turbine wakes and wind farms. *Wind Energy*, 2(1):1–24, 1999.
- J. Dahlberg and D. Medici. Potential improvement of wind turbine array efficiency by active wake control. In *Proceedings of European Wind Energy Conference*, Madrid, Spain, 2003.
- A. G. Dankers. *System Identification in Dynamic Networks*. PhD thesis, Delft University of Technology, 2014.
- A. T. De Boer. The Netherlands. In *IEA Wind 2012 Annual Report*, pages 126–131. International Energy Agency, 2013.
- O. Edenhofer, R. Pichs-Madruga, Y. Sokona, K. Seyboth, P. Matschoss, S. Kadner, T. Zwickel, P. Eickemeier, G. Hansen, S. Schlömer, and C. von Stechow. Wind energy - resource potential. In *Renewable Energy Sources and Climate Change Mitigation*, Special Report of the Intergovernmental Panel on Climate Change, pages 543–550. Cambridge University Press, Cambridge, UK and New York, NY, USA, 2011.
- G. España, S. Aubrun, S. Loyer, and P. Devinant. Spatial study of the wake meandering using modelled wind turbines in a wind tunnel. *Wind Energy*, 14(7):923–937, 2011.
- European Commission. European energy security strategy (com/2014/0330). In *Communication from the Commission to the European Parliament and the Council*, Brussels, Belgium, 2014.
- P. Fleming, P. Gebraad, M. Churchfield, S. Lee, K. Johnson, J. Michalakes, J. W. van Wingerden, and P. Moriarty. SOWFA + super controller user’s manual. Technical Report NREL/TP-5000-59197, National Renewable Energy Laboratory, 2013a.
- P. Fleming, P. Gebraad, J. W. van Wingerden, S. Lee, M. Churchfield, A. Scholbrock, J. Michalakes, K. Johnson, and P. Moriarty. The SOWFA super-controller: A high-fidelity tool for evaluating wind plant control approaches. In *Proceedings of the EWEA Annual Meeting*, Vienna, Austria, 2013b.
- P. Fleming, P. Gebraad, M. Churchfield, J. W. van Wingerden, A. Scholbrock, and P. Moriarty. Using particle filters to track wind turbine wakes for improved wind plant controls. In *Proceedings of the American Control Conference*, Portland, OR, USA, 2014a.
- P. Fleming, P. M. O. Gebraad, S. Lee, J. W. van Wingerden, K. Johnson, M. Churchfield, J. Michalakes, P. Spalart, and P. Moriarty. Simulation comparison of wake mitigation control strategies for a two-turbine case. *Wind Energy*, 2014b. to appear.
- P. Fleming, A. Ning, P. M. O. Gebraad, and K. Dykes. Wind plant system engineering through optimization of layout and yaw control. in preparation, 2014c.
- P. A. Fleming, P. M. O. Gebraad, S. Lee, J. W. van Wingerden, K. Johnson, M. Churchfield, J. Michalakes, P. Spalart, and P. Moriarty. Evaluating techniques for redirecting turbine wakes using SOWFA. *Renewable Energy*, 70:211–218, 2014d.
- R. Fletcher and M. J. D. Powell. A rapidly convergent descent method for minimization. *Computer Journal*, 6:163–168, 1963.

- FLOW. Far and large offshore wind research program. <http://www.flow-offshore.nl/page/projects,2010a>.
- FLOW. Wind farm wake modelling, fatigue loads and control. <http://www.flow-offshore.nl/page/wind-farm-wake-modelling-fatigue-loads-and-control,2010b>.
- S. Frandsen, R. Barthelmie, S. Pryor, O. Rathmann, S. Larsen, J. Højstrup, and M. Thøgersen. Analytical modelling of wind speed deficit in large offshore wind farms. *Wind Energy*, 9(1-2):39–53, 2006.
- P. M. O. Gebraad and J. W. van Wingerden. A control-oriented dynamic model for wakes in wind plants. In *Proceedings of the Science of Making Torque from Wind*, Lyngby, Denmark, 2014a.
- P. M. O. Gebraad and J. W. van Wingerden. Maximum power-point tracking control for wind farms. *Wind Energy*, 2014b.
- P. M. O. Gebraad, J. W. van Wingerden, G. J. Van der Veen, and M. Verhaegen. LPV subspace identification using a novel nuclear norm regularization method. In *Proceedings of the American Control Conference*, pages 165–170. IEEE, 2011.
- P. M. O. Gebraad, J. W. van Wingerden, and M. Verhaegen. Sparse estimation for predictor-based subspace identification of LPV systems. In *Proceedings of System Identification*, volume 16, pages 1749–1754, 2012.
- P. M. O. Gebraad, J. W. van Wingerden, P. A. Fleming, and A. D. Wright. LPV identification of wind turbine rotor vibrational dynamics using periodic disturbance basis functions. *Control Systems Technology, IEEE Transactions on*, 21(4):1183–1190, 2013.
- P. M. O. Gebraad, F. W. Teeuwisse, J. W. van Wingerden, P. A. Fleming, S. D. Ruben, J. R. Marden, and L. Y. Pao. A data-driven model for wind plant power optimization by yaw control. In *Proceedings of the American Control Conference*, Portland, OR, USA, 2014c.
- 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*, 2014d. to appear.
- J. P. Goit and J. Meyers. Analysis of turbulent flow properties and energy fluxes in optimally controlled wind-farm boundary layers. In *The Science of Making Torque from Wind*, Lyngby, Denmark, 2014.
- J. S. Gonzalez, M. B. Payan, and J. R. Santos. Optimal control of wind turbines for minimizing overall wake effect losses in offshore wind farms. In *Proceedings of EUROCON*, pages 1129–1134, 2013.
- F. González-Longatt, P. Wall, and V. Terzija. Wake effect in wind farm performance: Steady-state and dynamic behavior. *Renewable Energy*, 39(1):329–338, 2012.
- S. Guntur, N. Troldborg, and M. Gaunaa. Application of engineering models to predict wake deflection due to a tilted wind turbine. In *Proceedings of the EWEA Annual Meeting*, Copenhagen, Denmark, 2012.

- Y. Hao, M. A. Lackner, R. E. Keck, S. Lee, M. J. Churchfield, and P. J. Moriarty. Implementing the dynamic wake meandering model in the NWTC design codes. In *Proceedings of the 32nd ASME Wind Energy Symposium*, National Harbor, Maryland, 2014.
- C. B. Hasager, L. Rasmussen, A. Peña, L. E. Jensen, and P.-E. Réthoré. Wind farm wake: The Horns Rev photo case. *Energies*, 6(2):696–716, 2013.
- E. Hau. Control systems and operational sequence. In *Wind Turbines - Fundamentals, Technologies, Application, Economics*, pages 429–465. Springer Berlin Heidelberg, 2013a. ISBN 978-3-642-27150-2.
- E. Hau. Basic concepts of wind energy converters. In *Wind Turbines - Fundamentals, Technologies, Application, Economics*, pages 65–78. Springer Berlin Heidelberg, 2013b. ISBN 978-3-642-27150-2.
- F. Heer, P. M. Esfahani, M. Kamgarpour, and J. Lygeros. Model based power optimisation of wind farms. In *Proceedings of the European Control Conference*, pages 1145–1150, Strasbourg, France, June 2014.
- T. Horvat, V. Spudic, and M. Baotic. Quasi-stationary optimal control for wind farm with closely spaced turbines. In *MIPRO, 2012 Proceedings of the 35th International Convention*, pages 829–834, May 2012.
- I. Houtzager. *Towards Data-Driven Control for Modern Wind Turbines*. PhD thesis, Delft University of Technology, 2011.
- Int. Energy Agency. Renewable power, levelised cost of energy. In *Tracking Clean Energy Progress 2014 - Energy Technology Perspectives 2014 Excerpt*, page 25. 2014.
- N. O. Jensen. A note on wind generator interaction. Technical Report Risø-M-2411, Risø National Laboratory, Roskilde, 1984.
- A. Jiménez, A. Crespo, and E. Migoya. Application of a LES technique to characterize the wake deflection of a wind turbine in yaw. *Wind Energy*, 13(6):559–572, 2010.
- K. Johnson and G. Fritsch. Assessment of extremum seeking control for wind farm energy production. *Wind engineering*, 36(6):701–716, 2012.
- K. E. Johnson and N. Thomas. Wind farm control: addressing the aerodynamic interaction among wind turbines. In *Proceedings of the American Control Conference*, pages 2104–2109, St. Louis, USA, 2009.
- J. Jonkman and M. L. Buhl. FAST user's guide. Technical Report NREL/EL-500-38230, National Renewable Energy Laboratory, 2005.
- J. Jonkman, S. Butterfield, W. Musial, and G. Scott. Definition of a 5-MW reference wind turbine for offshore system development. Technical Report NREL/TP-500-38060, National Renewable Energy Laboratory, 2009.
- R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME - Journal of Basic Engineering*, 82 (Series D):35–45, 1960.

- I. Katić, J. Højstrup, and N. O. Jensen. A simple model for cluster efficiency. In *Proceedings of the European Wind Energy Association Conference and Exhibition*, pages 407–410, Rome, Italy, 1986.
- T. Knudsen and T. Bak. Data driven modelling of the dynamic wake between two wind turbines. In *Proceedings of the 16th IFAC Symposium on System Identification*, Brussels, Belgium, 2012.
- T. Knudsen, T. Bak, and M. Soltani. Prediction models for wind speed at turbine locations in a wind farm. *Wind Energy*, 14(7):877–894, 2011.
- T. Knudsen, T. Bak, and M. Svenstrup. Survey of wind farm control – power and fatigue optimization. *Wind Energy*, 2014.
- F. Koch, M. Gresch, F. Shewarega, I. Erlich, and U. Bachmann. Consideration of wind farm wake effect in power system dynamic simulation. In *Proceedings of IEEE Power Tech*, St. Petersburg, Russia, 2005.
- E. Koutroulis and K. Kalaitzakis. Design of a maximum power tracking system for wind energy-conversion applications. *IEEE Transactions on Industrial Electronics*, 53(2): 486–494, 2006.
- K. A. Kragh and M. H. Hansen. Load alleviation of wind turbines by yaw misalignment. *Wind Energy*, 2013a.
- K. A. Kragh, P. A. Fleming, and A. K. Scholbrock. Increased power capture by rotor speed-dependent yaw control of wind turbines. *Journal of Solar Energy Engineering*, 135(3), 2013b.
- E. Lantz, R. Wiser, and M. Hand. The past and future cost of wind energy. In *World Renewable Energy Forum*, Denver, Colorado, 2012.
- G. C. Larsen, H. A. Madsen, F. Bingöl, J. Mann, S. Ott, J. N. Sørensen, V. Okulov, N. Troidborg, M. Nielsen, K. Thomsen, T. J. Larsen, and R. Mikkelsen. Dynamic wake meandering modeling. Technical Report Risø-R-1607(EN), Risø National Laboratory, 2007.
- G. C. Larsen, H. A. Madsen, K. Thomsen, and T. J. Larsen. Wake meandering: a pragmatic approach. *Wind Energy*, 11(4):377–395, 2008.
- S. Lee, M. Churchfield, P. Moriarty, J. Jonkman, and J. Michalakes. Atmospheric and wake turbulence impacts on wind turbine fatigue loadings (presentation slides). In *Proceedings of the AIAA Aerospace Sciences Meeting*, Nashville, Tennessee, USA, 2012.
- A. Littel. Windenergie op zee - ontwerp-structuurvisie nader beschouwd. *ODE Windnieuws*, 3:10–13, 2014. in Dutch.
- L. Machiels, S. Barth, E. Bot, H. Hendriks, and G. Schepers. Evaluation of “heat and flux” farm control, final report. Technical Report ECN-E-07-105, Energy Centre of the Netherlands, 2007.
- D. Madjidian, K. Mårtensson, and A. Rantzer. A distributed power coordination scheme for fatigue load reduction in wind farms. In *Proceedings of American Control Conference*, pages 5219–5224, San Francisco, CA, USA, 2011.

- J. R. Marden, S. D. Ruben, and L. Y. Pao. A model-free approach to wind farm control using game theoretic methods. *IEEE Transactions on Control Systems Technology*, 21(4):1207–1214, 2013.
- MathWorks. MATLAB direct search optimization algorithms. <http://www.mathworks.com/help/gads/direct-search.html>, 2014.
- D. Medici. *Experimental studies of wind turbine wakes: power optimisation and meandering*. PhD thesis, KTH Royal Institute of Technology, 2005.
- D. Medici and P. H. Alfredsson. Measurements on a wind turbine wake: 3D effects and bluff body vortex shedding. *Wind Energy*, 9(3):219–236, 2006.
- T. Mikkelsen, N. Angelou, K. Hansen, M. Sjøholm, M. Harris, C. Slinger, P. Hadley, R. Scullion, G. Ellis, and G. Vives. A spinner-integrated wind lidar for enhanced wind turbine control. *Wind Energy*, 16(4):625–643, 2013.
- Ministerie IenM. Ontwerp-rijksstructuurvisie windenergie op zee, Ministerie van Infrastructuur en Milieu, 2014.
- D. P. Molenaar. *Cost-effective design and operation of variable speed wind turbines*. PhD thesis, Delft University of Technology, 2003.
- P. Moriarty, J. S. Rodrigo, P. Gancarski, M. Churchfield, J. W. Naughton, K. S. Hansen, E. Macheaux, E. Maguire, F. Castellani, L. Terzi, S.-P. Breton, and Y. Ueda. IEA-task 31 WAKEBENCH: Towards a protocol for wind farm flow model evaluation. part 2: Wind farm wake models. In *The Science of Making Torque from Wind*, Lyngby, Denmark, 2014.
- G. Mosetti, C. Poloni, and B. Diviacco. Optimization of wind turbine positioning in large windfarms by means of a genetic algorithm. *Journal of Wind Engineering and Industrial Aerodynamics*, 51(1):105–116, 1994.
- National Renewable Energy Laboratory. NREL's high-performance computing capabilities in 2012. http://www.nrel.gov/energysciences/csc/high_performance_computing_capabilities, 2012.
- K. Nilsson, S. Ivanell, K. S. Hansen, R. Mikkelsen, J. N. Sørensen, S.-P. Breton, and D. Henningson. Large-eddy simulations of the Lillgrund wind farm. *Wind Energy*, 2014.
- NoordzeeWind B.V. Data of the NoordzeeWind monitoring and evaluation programme (NSW-MEP). <http://www.noordzeewind.nl/en/knowledge/reportsdata/>, 2013.
- K. Z. Østergaard, P. Brath, and J. Stoustrup. Estimation of effective wind speed. In *The Science of Making Torque from Wind*, Lyngby, Denmark, 2007.
- J. Park, S. Kwon, and K. H. Law. Wind farm power maximization based on a cooperative static game approach. In *Proceedings of the SPIE Active and Passive Smart Structures and Integrated Systems Conference*, San Diego, CA, USA, 2013.
- J. K. Rice. *Efficient algorithms for distributed control: a structured matrix approach*. PhD thesis, Delft University of Technology, 2010.

- M. Samorani. The wind farm layout optimization problem. In *Handbook of Wind Power Systems*, pages 21–38. Springer Berlin Heidelberg, 2013.
- B. Sanderse, S. van der Pijl, and B. Koren. Review of computational fluid dynamics for wind turbine wake aerodynamics. *Wind Energy*, 14(7):799–819, 2011. ISSN 1099-1824.
- P. Schaak. Heat and flux: enabling the wind turbine controller. Technical Report ECN-E-06-017, Energy Research Centre of the Netherlands, 2006.
- J. G. Schepers and S. P. van der Pijl. Improved modelling of wake aerodynamics and assessment of new farm control strategies. *Journal of Wind Engineering and Industrial Aerodynamics*, 75, 2007.
- E. Simley, L. Y. Pao, P. Gebraad, and M. Churchfield. Investigation of the impact of the upstream induction zone on lidar measurement accuracy for wind turbine control applications using large-eddy simulation. In *Proceedings of the Science of Making Torque from Wind*, Lyngby, Denmark, 2014.
- M. Soleimanzadeh and R. Wisniewski. Controller design for a wind farm, considering both power and load aspects. *Mechatronics*, 21(4):720–727, 2011.
- M. Soleimanzadeh, R. Wisniewski, and S. Kanev. An optimization framework for load and power distribution in wind farms. *Journal of Wind Engineering and Industrial Aerodynamics*, 107–108:256–262, 2012.
- M. Soleimanzadeh, R. Wisniewski, and K. Johnson. A distributed optimization framework for wind farms. *Journal of Wind Engineering and Industrial Aerodynamics*, 123: 88–98, 2013.
- M. Soleimanzadeh, R. Wisniewski, and A. Brand. State-space representation of the wind flow model in wind farms. *Wind Energy*, 17(4):627–639, 2014.
- M. Soltani, T. Knudsen, J. D. Grunnet, and T. Bak. AEOLUS toolbox for dynamic wind farm model, simulation, and control. In *European Wind Energy Conference*, pages 3119–3129, Warsaw, Poland, 2010.
- J. N. Sørensen and W. Z. Shen. Numerical modeling of wind turbine wakes. *Journal of fluids engineering*, 124(2):393–399, 2002.
- V. Spudić, M. Jelavic, M. Baotic, and N. Peric. Hierarchical wind farm control for power/load optimization. In *The Science of Making Torque from Wind*, pages 681–692, Heraklion, Greece, 2010.
- M. Steinbuch, W. W. de Boer, O. Bosgra, S. Peters, and J. Ploeg. Optimal control of wind power plants. *Journal of Wind Engineering and Industrial Aerodynamics*, 27:237–246, 1988.
- P. Torres, J. W. van Wingerden, and M. Verhaegen. Modeling of the flow in wind farms for total power optimization. In *Proceedings of the 9th IEEE International Conference on Control and Automation*, pages 1207–1214, Santiago, Chili, 2010.
- P. Torres Tapia. *Linear state-space identification of interconnected systems: a structured approach*. PhD thesis, Delft University of Technology, 2014.

- G. J. van der Veen. *Identification of Wind Energy Systems*. PhD thesis, Delft University of Technology, 2013.
- G. J. Van der Veen, I. J. Couchman, and R. O. Bowyer. Control of floating wind turbines. In *Proceedings of the American Control Conference (ACC)*, pages 3148–3153, Montréal, Canada, 2012.
- J. W. van Wingerden. Reconfigurable floating wind turbines, VENI project #11930, 2011.
- J. W. van Wingerden, A. W. Hulskamp, T. Barlas, B. Marrant, G. A. M. van Kuik, D. P. Moleenaar, and M. Verhaegen. On the proof of concept of a ‘smart’ wind turbine rotor blade for load alleviation. *Wind Energy*, 11(3):265–280, 2008. ISSN 1099-1824.
- J. W. van Wingerden, P. Gebraad, and M. Verhaegen. LPV identification of an aeroelastic flutter model. In *Proceedings of the IEEE Conference on Decision and Control*, pages 6839–6844. IEEE, 2010.
- M. Verhaegen and V. Verdult. *Filtering and System Identification, A Least Squares Approach*. Cambridge University Press, 2007.
- L. J. Vermeer, J. N. Sørensen, and A. Crespo. Wind turbine wake aerodynamics. *Progress in Aerospace Sciences*, 39(6–7):467–510, 2003.
- J. W. Wagenaar, L. Machiels, and J. Schepers. Controlling wind in ECN’s scaled wind farm. In *Proceedings of the EWEA Annual Meeting*, Copenhagen, Denmark, 2012.
- X. Yang and F. Sotiropoulos. On the predictive capabilities of LES actuator disk model in simulating turbulence past wind turbines and plants. In *Proceedings of the American Control Conference*, pages 2878–2883, Washington, D.C., USA, 2013a.
- Z. Yang, Y. Li, and Y. Seem. Improved individual pitch control for wind farm turbine load reduction via wake modeling. In *Proceedings of the AIAA Aerospace Sciences Meeting*, Orlando, Florida, USA, 2011.
- Z. Yang, Y. Li, and Y. Seem. Maximizing wind farm energy capture via nested-loop extremum seeking control. In *Proceedings of the ASME Dynamic Systems and Control Conference*, Palo Alto, California, USA, 2013b.

SUMMARY

Wind turbines that are clustered in a wind plant, have interaction with each other through the aerodynamics of the wind field in the wind plant. The aerodynamic interaction effects are caused by the turbine wakes, which are the flow structures that form behind each turbine. The wake is characterized by a reduced flow velocity caused by the extraction of energy from the flow by the turbine, and an increased turbulence intensity caused by the obstruction of the flow by the turbine. The velocity deficits will cause a decrease of electrical power production of turbines standing in the path of a wake of another turbine, and the increased turbulence may increase the fatigue loads on those downstream turbines. Wind plant control that takes into account wake interaction effects in the coordination of the control actions of the wind turbines, can enhance the performance of the wind plant, in terms of total electrical energy production, and the loads on the individual wind turbines. Enhancing wind plant performance in this way, will contribute to the reduction of the cost of offshore and onshore wind energy. In this thesis two research objectives have been addressed: one is the evaluation of the potential of the different control degrees-of-freedom of the wind turbine to affect the interaction effect between the turbines, and the other is the development of data-driven algorithms for the optimization of those control settings in order to improve wind plant performance.

The control degrees-of-freedom (DOFs) of a modern large-scale horizontal-axis wind turbine are generator torque, collective and individual blade pitch, and rotor yaw. For each of the control DOFs, we explored their ability to affect the wake interaction effects in the wind plant, through high-fidelity computational simulations of setups with one or more turbines. We researched two ways in which the DOFs can control the wake effects: by axial-induction-based wake control, in which the energy extraction of the wind turbine is changed in order to affect the velocity deficit in the wake, and by wake deflection, in which the flow direction in the wake is manipulated in order to steer them away from downstream turbines.

Axial-induction-based control can be performed by offsetting collective blade pitch or generator torque from their turbine-level optimal settings, in order to increase the power production of downstream turbines. The potential gain from using axial-induction-based plant-wide instead of turbine-level optimized control, are dependent on the particular atmospheric conditions, the wind plant configuration, and the turbine characteristics. Simulation cases presented in this thesis even show that there are circumstances in which axial-induction-based control gives no total power production increase on the wind plant, because wake expansion makes that much of the energy added to the flow by the control changes on upstream turbines, is lost to the atmosphere instead of being captured by downstream turbines.

A better potential is demonstrated for wake redirection control. Yaw control and modified individual pitch control (IPC) were shown to be able to induce significant wake

redirection. For the tested IPC-based techniques however, this wake redirection goes at the cost of loads increases, making it a less suitable candidate for practical application. Wake deflection through yaw offsets has been shown to be successful at increasing the total power production of small wind plant setups in high-fidelity simulations. Both increases and reductions of fatigue loads result from yaw-based wake redirection, depending on the settings used. The load increases can be partially mitigated using standard load-reducing IPC.

Two main approaches for data-driven wind plant control development were taken in this thesis: direct data-driven approaches, in which the control settings optimization is directly based on measured data, and data-driven model-based control, in which the optimization is based on models for which the parameters are identified from measured wind plant data.

Direct data-driven algorithms for axial-induction-based control are presented that optimize the axial-induction settings of each turbine in the plant with the objective of power production increase. A speed-up of the optimization is achieved by using gradient-based optimization techniques with a distributed approach in which the effect on neighboring turbines is taken into account only. Using information on the spatial configuration of the wind farm in this way, results in a much faster convergence of the power optimization than is achieved with an existing method with full communication between the turbines. This is because there are significant delays between control actions on turbines, and the response of turbines further downstream, because the wake effects need to propagate through the wind field.

Further, a data-driven model-based wind plant control strategy was presented to optimize the yaw settings of each turbine in the plant with the objective of electricity production increase. The optimization is based on predictions provided by a newly developed control-oriented model that predicts the effects of the yaw settings on the steady-state wake deflection and velocity deficit. The model has a relatively small number of parameters that can be identified using time-averaged electrical power measurements of the different turbines in the wind plant (gray-box system identification). The model is computationally efficient enough to enable real-time online wind plant control optimization, using a game-theoretic search method. The application of the wind plant optimization method in a high-fidelity CFD simulation of a small wind plant, demonstrated an increase in electrical power production, and a reduction of wind turbine fatigue loads. Based on the novel data-driven model, a 1.1% electricity production increase is predicted on an annual basis, through using yaw-based wind plant control on the Princess Amalia Windpark, a full-size offshore wind plant.

The control-oriented wind plant model was further extended to include dynamics of the wake propagation delays between the turbines. The dynamic model has a feedback structure in which a linear state-space model describes the propagation of changes in the wake through the wind field. Only two tunable model parameters were added to include these dynamics, such that the tuning process is not further complicated to a large extent. The model structure allowed the development of an observer that uses Kalman filtering to correct the predicted flow field velocities on the basis of power measurements at the turbines. A relatively small increase of wind plant performance was demonstrated

by going from yaw control optimization based on the steady-state model, to optimization based on the dynamic model. A practical benefit of including the dynamics is the fact that the gray-box system identification can be performed directly on the basis of the dynamic responses of the wind plant, rather than on time-averaged data. This is expected to be an important benefit when applying data-driven model-based control on a real wind plant with continuously changing ambient conditions.

This thesis presented pioneering work in data-driven wind plant controls, researching the feasibility of wake control concepts and solving practical issues to enable real-time implementation.

SAMENVATTING

Windturbines die zijn gegroepeerd in een windpark, hebben interactie met elkaar via de aerodynamica van het windveld in het windpark. De aerodynamische interactie-effecten worden veroorzaakt door de zoggen van de turbines. Het zog van de turbine is de stromingsstructuur die zich vormt achter de turbine. Het zog wordt gekarakteriseerd door een stromingsnelheidsvermindering, veroorzaakt door de extractie van energie uit de stroming door de turbine, en een toegenomen turbulentie-intensiteit, veroorzaakt door de versperring van de stroming door de turbine. De snelheidsvermindering veroorzaakt een afname van de elektriciteitsproductie bij turbines die in het pad staan van het zog van een andere turbine, en de toegenomen turbulentie kan de belastingen vermeerderen op die benedenstroomse turbines. Windparkregelsystemen die de zogsinteractie-effecten in beschouwing nemen bij de coördinatie van de regelacties van de windturbines, kunnen de prestatie van het windpark verbeteren, voor wat betreft de totale elektriciteitsproductie en de belastingen op de afzonderlijke windturbines. Door op die manier de windparkprestatie te verbeteren, kan worden bijgedragen aan de kostenvermindering van windenergie op zee en op land. In dit proefschrift worden twee onderzoeksdoelen behandeld: het eerste is de inschatting van het potentieel van de verschillende vrijheidsgraden van de turbine om de interactie-effecten tussen de turbines te beïnvloeden, en het tweede is de ontwikkeling van op data gebaseerde algoritmes voor de optimalisatie van de regelinstellingen om de windparkprestatie te verbeteren.

De vrijheidsgraden voor de regeling van een moderne grote windturbines met een horizontale as, zijn generatorkoppel, collectieve en individuele bladhoek en de gierhoek van de rotor (draaien om de verticale as, of 'kruien'). Voor elk van deze vrijheidsgraden werd onderzocht hoe ze de zogsinteractie-effecten in het windpark beïnvloeden. Daarbij werd gebruik gemaakt van nauwkeurige simulaties van opstellingen met een of meer turbines. We onderzochten twee manieren waarop de vrijheidsgraden de zogseffecten kunnen beïnvloeden: door op axiale-inductie gebaseerde zogsregeling, waarbij de energie-extractie van de windturbine wordt veranderd om de snelheidsafname in het zog te beïnvloeden, en door zogsafbuiging, waarbij de stromingsrichting in het zog wordt veranderd om het weg te sturen van benedenstroomse turbines.

Bij op axiale-inductie gebaseerde zogsregeling wordt afgeweken van de op turbine-niveau optimale collectieve bladhoek- of generatorkoppelinstellingen. De potentiële winst van het gebruik van axiale-inductieregeling die op windpark- in plaats van op turbineniveau optimaal is, hangt af van de specifieke atmosferische condities, de configuratie van het windpark, en de turbinekarakteristieken. De simulatievoorbeelden in dit proefschrift tonen zelfs aan dat er omstandigheden zijn waarin aanpassingen van de axiale inductie geen vermogenstoename oplevert, omdat expansie van het zog ervoor zorgt dat veel van de energie die aan de stroom wordt toegevoegd door middel van regelverandering, verloren gaat aan de atmosfeer.

Er is een beter potentieel aangetoond voor regelingen gebaseerd op zogsafbuiging. Het is aangetoond dat zowel aanpassing van de gierhoek, als van de individuele bladhoekregeling, significante zogsafbuiging veroorzaakt. Voor de individuele bladhoekregeltechnieken geldt echter dat de zogsafbuiging ten koste gaat van belastingtoenames, wat ze minder geschikt maakt voor praktische toepassing. Het is aangetoond in nauwkeurige simulatie, dat zogsafbuiging door middel van gierhoekafwijkingen de totale elektriciteitsproductie van kleine windparkopstellingen kan doen laten toenemen. Zowel toenames als afnames van vermoeiingsbelastingen kunnen het resultaat zijn van op gierhoek gebaseerde zogsafbuiging, afhankelijk van de instellingen die worden gebruikt. De belastingtoename kan deels worden gemitigeerd door gebruik te maken van standaard belastingverminderende individuele bladhoeksregeltechnieken.

In dit proefschrift zijn twee richtingen genomen in de ontwikkeling van op data gebaseerde windparkregelingen: direct datagestuurde methodes, waarin de optimalisatie van de regelinstellingen direct gebaseerd is op de gemeten data, en datagestuurde modelgebaseerde regelingen, waarin de optimalisatie gebaseerd is op modellen waarvoor de parameters worden geschat uit gemeten windparkdata.

In dit proefschrift zijn direct datagestuurde algoritmes voor op axiale-inductie gebaseerde regeling gepresenteerd, die de axiale-inductie instellingen optimaliseren voor elke turbine in het windpark met vermogensproductietoename als doel. Een versnelling van deze optimalisatie wordt bereikt door gebruik te maken van gradiëntgebaseerde optimalisatietechnieken met een gedistribueerde aanpak waarin alleen het effect van naburige turbines in beschouwing wordt genomen. Door op die manier gebruik te maken van de spatiale configuratie van het windpark, kan een veel snellere convergentie van de vermogensoptimalisatie worden bereikt dan met een bestaande methode die uitgaat van volledige communicatie tussen de turbines. Dit komt doordat er grote vertragingen zijn tussen de regelacties op de turbines, en de responsie van turbines verder stroomafwaarts, omdat de zogeeffecten zich eerst door het windveld dienen voort te planten.

Verder werden datagestuurde modelgebaseerde windparkregelingsstrategieën gepresenteerd, die de gierhoekinstellingen van de turbines in het windpark optimaliseren met vermogensproductietoename als doel. De optimalisatie is gebaseerd op voorspellingen die gegeven worden door een nieuw ontwikkeld, op regeltechniek gericht model dat de effecten van de gierhoekinstellingen op de stationaire zogsafbuigingen en -snelheden voorspelt. Het model heeft een relatief klein aantal afstempparameters. Deze parameters kunnen worden geïdentificeerd uit over de tijd gemiddelde elektrisch vermogensmetingen van de verschillende turbines in het windpark. Het model is rekenkundig efficiënt genoeg om realtime windparkregelingsoptimalisatie mogelijk te maken, gebruik makend van op speltheorie gebaseerde zoekmethode. De toepassing van de windparkregelingsoptimalisatie in een nauwkeurige rekenkundige simulatie van de stromingsdynamica in een klein windpark, resulteerde in een toename van de vermogensproductie van het windpark en een afname van de vermoeiingsbelastingen op de turbines. Het nieuwe datagestuurde model voorspelt een elektriciteitsproductietoename van het Prinses Amalia Windpark, een groot windpark op zee, van 1.1% op jaarbasis, als gebruik wordt gemaakt van de geoptimaliseerde gierhoekregeling in het windpark.

Het op regeltechniek gerichte windparkmodel is uitgebreid met de dynamica van

de voortplanting van het zog tussen the turbines. Het dynamisch model heeft een terugkoppelingsstructuur waarin een lineair toestandsruimtemodel de voortplanting van zogsveranderingen door het windveld beschrijft. Om deze dynamica toe te voegen, zijn slechte twee modelparameters toegevoegd, zodat proces van het afstemmen van de parameters niet veel is bemoeilijkt. Met deze modelstructuur kon een toestandsschatter (observer) worden ontwikkeld die Kalman-filtertechnieken gebruikt om snelheden in het voorspelde stromingsveld te corrigeren op basis van het gemeten elektrisch vermogen van de turbines. Een relatief kleine toename van de windparkprestatie werd behaald door te gaan van optimalisatie van de gierhoekregeling op basis van het stationaire model, naar optimalisatie op basis van het dynamisch model. Een praktisch voordeel van het toevoegen van de dynamica is het feit dat het afstemmen van de modelparameters direct kan worden gedaan op basis van de dynamische responsie van het windpark, in plaats van op data die gemiddeld is over de tijd. Naar verwachting is dit een belangrijk voordeel wanneer datagestuurde modelgebaseerde regeling wordt toegepast in een echt windpark met continu veranderende omgevingscondities.

Dit proefschrift bevat verkennend werk op het gebied van datagestuurde windparkregelingen. In dit onderzoek verkenden we de haalbaarheid van concepten voor zogregeling in het windpark en droegen we oplossingen aan voor praktische problemen bij de realtime implementatie van deze concepten.

LIST OF PUBLICATIONS

JOURNAL ARTICLES

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*, 2014d. to appear

P. M. O. Gebraad and J. W. van Wingerden. Maximum power-point tracking control for wind farms. *Wind Energy*, 2014b

P. M. O. Gebraad, J. W. van Wingerden, P. A. Fleming, and A. D. Wright. LPV identification of wind turbine rotor vibrational dynamics using periodic disturbance basis functions. *Control Systems Technology, IEEE Transactions on*, 21(4):1183–1190, 2013

P. A. Fleming, P. M. O. Gebraad, S. Lee, J. W. van Wingerden, K. Johnson, M. Churchfield, J. Michalakes, P. Spalart, and P. Moriarty. Evaluating techniques for redirecting turbine wakes using SOWFA. *Renewable Energy*, 70:211–218, 2014d

P. Fleming, P. M. O. Gebraad, S. Lee, J. W. van Wingerden, K. Johnson, M. Churchfield, J. Michalakes, P. Spalart, and P. Moriarty. Simulation comparison of wake mitigation control strategies for a two-turbine case. *Wind Energy*, 2014b. to appear

J. Annoni, P. M. O. Gebraad, A. K. Scholbrock, P. Fleming, and J. W. van Wingerden. Analysis of axial-induction-based wind plant control using low-order and high-order wind plant models. submitted, 2014a

P. Fleming, A. Ning, P. M. O. Gebraad, and K. Dykes. Wind plant system engineering through optimization of layout and yaw control. in preparation, 2014c

CONFERENCE PAPERS

(Abbreviated list, not including journal-published work.)

P. M. O. Gebraad and J. W. van Wingerden. A control-oriented dynamic model for wakes in wind plants. In *Proceedings of the Science of Making Torque from Wind*, Lyngby, Denmark, 2014a

P. M. O. Gebraad, F. W. Teeuwisse, J. W. van Wingerden, P. A. Fleming, S. D. Ruben, J. R. Marden, and L. Y. Pao. A data-driven model for wind plant power optimization by yaw control. In *Proceedings of the American Control Conference*, Portland, OR, USA, 2014c

E. Simley, L. Y. Pao, P. Gebraad, and M. Churchfield. Investigation of the impact of the upstream induction zone on lidar measurement accuracy for wind turbine control ap-

plications using large-eddy simulation. In *Proceedings of the Science of Making Torque from Wind*, Lyngby, Denmark, 2014

P. Fleming, P. Gebraad, M. Churchfield, J. W. van Wingerden, A. Scholbrock, and P. Moriarty. Using particle filters to track wind turbine wakes for improved wind plant controls. In *Proceedings of the American Control Conference*, Portland, OR, USA, 2014a

J. Annoni, P. Seiler, K. Johnson, P. Fleming, and P. Gebraad. Evaluating wake models for wind farm control. In *Proceedings of the American Control Conference*, Portland, OR, USA, 2014b

P. Fleming, P. Gebraad, J. W. van Wingerden, S. Lee, M. Churchfield, A. Scholbrock, J. Michalak, K. Johnson, and P. Moriarty. The SOWFA super-controller: A high-fidelity tool for evaluating wind plant control approaches. In *Proceedings of the EWEA Annual Meeting*, Vienna, Austria, 2013b

P. M. O. Gebraad, J. W. van Wingerden, and M. Verhaegen. Sparse estimation for predictor-based subspace identification of LPV systems. In *Proceedings of System Identification*, volume 16, pages 1749–1754, 2012

P. M. O. Gebraad, J. W. van Wingerden, G. J. Van der Veen, and M. Verhaegen. LPV subspace identification using a novel nuclear norm regularization method. In *Proceedings of the American Control Conference*, pages 165–170. IEEE, 2011

J. W. van Wingerden, P. Gebraad, and M. Verhaegen. LPV identification of an aeroelastic flutter model. In *Proceedings of the IEEE Conference on Decision and Control*, pages 6839–6844. IEEE, 2010

CURRICULUM VITAE

Pieter Gebraad was born on May 3rd, 1986 in Capelle aan den IJssel, The Netherlands.

Between 1998 and 2004, he followed secondary education at the IJsselcollege in Capelle aan den IJssel. In 2004 he started studies at Delft University of Technology, where in September 2010, he received the Master of Science degree (cum laude) in mechanical engineering, with a specialization in controls engineering. His M.Sc. graduation project was titled “LPV Subspace Identification: Improving the Performance and Application to Wind Turbines”.

In November 2010, he started his Ph.D. project at the Delft Center for Systems and Control and the Delft University Wind Research Institute (DUWIND), titled “Data-Driven Wind Plant Control”. This project was funded by the Far and Large Offshore Wind (FLOW) research program. During his Ph.D. research he supervised several M.Sc. students, and he participated in giving education at the Aerospace and Mechanical Engineering faculties, and in organizing the 2011 EAWE PhD Seminar on Wind Energy in Europe. During the summers of 2012 and 2013, he was a visiting researcher at the National Wind Technology Center of the National Renewable Energy Laboratory (NREL), in Boulder, Colorado, USA.

Since November 2014, he works as an advanced controls developer at NREL, performing postdoctoral research in the field of wind plant control.



ACKNOWLEDGMENTS

I would like to thank Jan-Willem van Wingerden, Michel Verhaegen and the Far and Large Offshore Wind (FLOW) research consortium for giving me the opportunity to perform wind energy research at the Delft Center of Systems and Control (DCSC), and for providing me with valuable support and feedback during the project. Further, I would like to thank my colleagues at the National Renewable Energy Laboratory (NREL), Colorado University (CU) and Minnesota University, and DCSC alumni Floris Teeuwisse and Filip van Dam, for their very significant contribution to the research presented this thesis. I thank the committee members for providing feedback and participating in the defense.

I would like to thank Jan-Willem for giving me the opportunity to meet Gijs, Sachin, Edwin, Noortje, and other DCSC colleagues, Giuseppe, Maxim, and other DUWIND colleagues, Paul, Kat & Ellianne, Andy & Dana, Jason, Katherine, Mohit & Alicia, Katie, and other NREL colleagues, David, Lorenz, and Steffen from Stuttgart, Gillis from Eindhoven, Lucy and her group at CU, Shara and Shirin from Boulder, Kara & Colin from Denver (then), and Jen from White Bear Lake. I thoroughly enjoyed working, sharing an office, organizing a PhD seminar, traveling, sharing a house, playing sports and games, going places, biking, hiking, dancing, singing, and/or simply hanging out with them. Together, they have made the past four years a very memorable and enjoyable time for me.

Finally, I would like to thank my family and friends for their support and showing an interest in my work, and for continuing to make time to be together.

*Pieter Gebraad
Delft, October 2014*

<i>A felicidade é como a pluma</i>	<i>Happiness is like a feather</i>
<i>Que o vento vai levando pelo ar</i>	<i>That the wind carries through the air</i>
<i>Voa tão leve</i>	<i>It flies so lightly</i>
<i>Mas tem a vida breve</i>	<i>But it has a brief life</i>
<i>Precisa que haja vento sem parar</i>	<i>It always needs the wind to be there</i>

Freely translated from the song
'A Felicidade' by Antônio Carlos Jobim and Vinicius de Moraes